

OFFICIAL MICROSOFT LEARNING PRODUCT

20414C

Implementing an Advanced Server
Infrastructure

NOT USE ONLY. STUDENT USE PROHIBITED

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

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

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

© 2014 Microsoft Corporation. All rights reserved.

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

Product Number: 20414C

Part Number: X19-30977

Released: 4/2014

MCT USE ONLY. STUDENT USE PROHIBITED

MICROSOFT LICENSE TERMS MICROSOFT INSTRUCTOR-LED COURSEWARE

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

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

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

1. DEFINITIONS.

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

- l. "Personal Device" means one (1) personal computer, device, workstation or other digital electronic device that you personally own or control that meets or exceeds the hardware level specified for the particular Microsoft Instructor-Led Courseware.
 - m. "Private Training Session" means the instructor-led training classes provided by MPN Members for corporate customers to teach a predefined learning objective using Microsoft Instructor-Led Courseware. These classes are not advertised or promoted to the general public and class attendance is restricted to individuals employed by or contracted by the corporate customer.
 - n. "Trainer" means (i) an academically accredited educator engaged by a Microsoft IT Academy Program Member to teach an Authorized Training Session, and/or (ii) a MCT.
 - o. "Trainer Content" means the trainer version of the Microsoft Instructor-Led Courseware and additional supplemental content designated solely for Trainers' use to teach a training session using the Microsoft Instructor-Led Courseware. Trainer Content may include Microsoft PowerPoint presentations, trainer preparation guide, train the trainer materials, Microsoft One Note packs, classroom setup guide and Pre-release course feedback form. To clarify, Trainer Content does not include any software, virtual hard disks or virtual machines.
- 2. USE RIGHTS.** The Licensed Content is licensed not sold. The Licensed Content is licensed on a **one copy per user basis**, such that you must acquire a license for each individual that accesses or uses the Licensed Content.
- 2.1 Below are five separate sets of use rights. Only one set of rights apply to you.
- a. **If you are a Microsoft IT Academy Program Member:**
 - i. Each license acquired on behalf of yourself may only be used to review one (1) copy of the Microsoft Instructor-Led Courseware in the form provided to you. If the Microsoft Instructor-Led Courseware is in digital format, you may install one (1) copy on up to three (3) Personal Devices. You may not install the Microsoft Instructor-Led Courseware on a device you do not own or control.
 - ii. For each license you acquire on behalf of an End User or Trainer, you may either:
 - 1. distribute one (1) hard copy version of the Microsoft Instructor-Led Courseware to one (1) End User who is enrolled in the Authorized Training Session, and only immediately prior to the commencement of the Authorized Training Session that is the subject matter of the Microsoft Instructor-Led Courseware being provided, **or**
 - 2. provide one (1) End User with the unique redemption code and instructions on how they can access one (1) digital version of the Microsoft Instructor-Led Courseware, **or**
 - 3. provide one (1) Trainer with the unique redemption code and instructions on how they can access one (1) Trainer Content,**provided you comply with the following:**
 - iii. you will only provide access to the Licensed Content to those individuals who have acquired a valid license to the Licensed Content,
 - iv. you will ensure each End User attending an Authorized Training Session has their own valid licensed copy of the Microsoft Instructor-Led Courseware that is the subject of the Authorized Training Session,
 - v. you will ensure that each End User provided with the hard-copy version of the Microsoft Instructor-Led Courseware will be presented with a copy of this agreement and each End User will agree that their use of the Microsoft Instructor-Led Courseware will be subject to the terms in this agreement prior to providing them with the Microsoft Instructor-Led Courseware. Each individual will be required to denote their acceptance of this agreement in a manner that is enforceable under local law prior to their accessing the Microsoft Instructor-Led Courseware,
 - vi. you will ensure that each Trainer teaching an Authorized Training Session has their own valid licensed copy of the Trainer Content that is the subject of the Authorized Training Session,

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

b. If you are a Microsoft Learning Competency Member:

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

c. **If you are a MPN Member:**

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

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

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

e. **If you are a Trainer.**

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

- ii. You may customize the written portions of the Trainer Content that are logically associated with instruction of a training session in accordance with the most recent version of the MCT agreement. If you elect to exercise the foregoing rights, you agree to comply with the following: (i) customizations may only be used for teaching Authorized Training Sessions and Private Training Sessions, and (ii) all customizations will comply with this agreement. For clarity, any use of “*customize*” refers only to changing the order of slides and content, and/or not using all the slides or content, it does not mean changing or modifying any slide or content.

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

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

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

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

3. **LICENSED CONTENT BASED ON PRE-RELEASE TECHNOLOGY.** If the Licensed Content’s subject matter is based on a pre-release version of Microsoft technology (“**Pre-release**”), then in addition to the other provisions in this agreement, these terms also apply:
 - a. **Pre-Release Licensed Content.** This Licensed Content subject matter is on the Pre-release version of the Microsoft technology. The technology may not work the way a final version of the technology will and we may change the technology for the final version. We also may not release a final version. Licensed Content based on the final version of the technology may not contain the same information as the Licensed Content based on the Pre-release version. Microsoft is under no obligation to provide you with any further content, including any Licensed Content based on the final version of the technology.
 - b. **Feedback.** If you agree to give feedback about the Licensed Content to Microsoft, either directly or through its third party designee, you give to Microsoft without charge, the right to use, share and commercialize your feedback in any way and for any purpose. You also give to third parties, without charge, any patent rights needed for their products, technologies and services to use or interface with any specific parts of a Microsoft software, Microsoft product, or service that includes the feedback. You will not give feedback that is subject to a license that requires Microsoft to license its software, technologies, or products to third parties because we include your feedback in them. These rights survive this agreement.
 - c. **Pre-release Term.** If you are an Microsoft IT Academy Program Member, Microsoft Learning Competency Member, MPN Member or Trainer, you will cease using all copies of the Licensed Content on the Pre-release technology upon (i) the date which Microsoft informs you is the end date for using the Licensed Content on the Pre-release technology, or (ii) sixty (60) days after the commercial release of the technology that is the subject of the Licensed Content, whichever is earliest (“**Pre-release term**”). Upon expiration or termination of the Pre-release term, you will irretrievably delete and destroy all copies of the Licensed Content in your possession or under your control.

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

 - access or allow any individual to access the Licensed Content if they have not acquired a valid license for the Licensed Content,
 - alter, remove or obscure any copyright or other protective notices (including watermarks), branding or identifications contained in the Licensed Content,
 - modify or create a derivative work of any Licensed Content,
 - publicly display, or make the Licensed Content available for others to access or use,
 - copy, print, install, sell, publish, transmit, lend, adapt, reuse, link to or post, make available or distribute the Licensed Content to any third party,
 - work around any technical limitations in the Licensed Content, or
 - reverse engineer, decompile, remove or otherwise thwart any protections or disassemble the Licensed Content except and only to the extent that applicable law expressly permits, despite this limitation.
- 5. RESERVATION OF RIGHTS AND OWNERSHIP.** Microsoft reserves all rights not expressly granted to you in this agreement. The Licensed Content is protected by copyright and other intellectual property laws and treaties. Microsoft or its suppliers own the title, copyright, and other intellectual property rights in the Licensed Content.
- 6. EXPORT RESTRICTIONS.** The Licensed Content is subject to United States export laws and regulations. You must comply with all domestic and international export laws and regulations that apply to the Licensed Content. These laws include restrictions on destinations, end users and end use. For additional information, see www.microsoft.com/exporting.
- 7. SUPPORT SERVICES.** Because the Licensed Content is “as is”, we may not provide support services for it.
- 8. TERMINATION.** Without prejudice to any other rights, Microsoft may terminate this agreement if you fail to comply with the terms and conditions of this agreement. Upon termination of this agreement for any reason, you will immediately stop all use of and delete and destroy all copies of the Licensed Content in your possession or under your control.
- 9. LINKS TO THIRD PARTY SITES.** You may link to third party sites through the use of the Licensed Content. The third party sites are not under the control of Microsoft, and Microsoft is not responsible for the contents of any third party sites, any links contained in third party sites, or any changes or updates to third party sites. Microsoft is not responsible for webcasting or any other form of transmission received from any third party sites. Microsoft is providing these links to third party sites to you only as a convenience, and the inclusion of any link does not imply an endorsement by Microsoft of the third party site.
- 10. ENTIRE AGREEMENT.** This agreement, and any additional terms for the Trainer Content, updates and supplements are the entire agreement for the Licensed Content, updates and supplements.
- 11. APPLICABLE LAW.**

 - a. United States. If you acquired the Licensed Content in the United States, Washington state law governs the interpretation of this agreement and applies to claims for breach of it, regardless of conflict of laws principles. The laws of the state where you live govern all other claims, including claims under state consumer protection laws, unfair competition laws, and in tort.

b. Outside the United States. If you acquired the Licensed Content in any other country, the laws of that country apply.

- 12. LEGAL EFFECT.** This agreement describes certain legal rights. You may have other rights under the laws of your country. You may also have rights with respect to the party from whom you acquired the Licensed Content. This agreement does not change your rights under the laws of your country if the laws of your country do not permit it to do so.
- 13. DISCLAIMER OF WARRANTY. THE LICENSED CONTENT IS LICENSED "AS-IS" AND "AS AVAILABLE." YOU BEAR THE RISK OF USING IT. MICROSOFT AND ITS RESPECTIVE AFFILIATES GIVES NO EXPRESS WARRANTIES, GUARANTEES, OR CONDITIONS. YOU MAY HAVE ADDITIONAL CONSUMER RIGHTS UNDER YOUR LOCAL LAWS WHICH THIS AGREEMENT CANNOT CHANGE. TO THE EXTENT PERMITTED UNDER YOUR LOCAL LAWS, MICROSOFT AND ITS RESPECTIVE AFFILIATES EXCLUDES ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT.**
- 14. LIMITATION ON AND EXCLUSION OF REMEDIES AND DAMAGES. YOU CAN RECOVER FROM MICROSOFT, ITS RESPECTIVE AFFILIATES AND ITS SUPPLIERS ONLY DIRECT DAMAGES UP TO US\$5.00. YOU CANNOT RECOVER ANY OTHER DAMAGES, INCLUDING CONSEQUENTIAL, LOST PROFITS, SPECIAL, INDIRECT OR INCIDENTAL DAMAGES.**

This limitation applies to

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

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

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

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

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

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

Cette limitation concerne:

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

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

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

Revised September 2012

Welcome!

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

- **Microsoft Certified Trainers and Instructors**—Your instructor is a technical and instructional expert who meets ongoing certification requirements. And, if instructors are delivering training at one of our Certified Partners for Learning Solutions, they are also evaluated throughout the year by students and by Microsoft.
- **Certification Exam Benefits**—After training, consider taking a Microsoft Certification exam. Microsoft Certifications validate your skills on Microsoft technologies and can help differentiate you when finding a job or boosting your career. In fact, independent research by IDC concluded that 75% of managers believe certifications are important to team 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.

Vladimir Meloski-Content Developer

Vladimir is an MCT, a Most Valuable Professional (MVP) on Exchange Server, and a consultant who provides unified communications and infrastructure solutions based on Microsoft Exchange Server, Microsoft Lync® Server, and Microsoft 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 also has been involved as a subject matter expert and technical reviewer for several Microsoft Official Curriculum (MOC) courses.

Dave Franklyn-Content Developer

Dave Franklyn, MCT, MCSE, Microsoft Certified IT Professional (MCITP), Microsoft MVP Windows Expert-It Pro, is a Senior Information Technology Trainer and Consultant at Auburn University in Montgomery, Alabama. He's been teaching for 15 years. He's been working with computers since 1976, when he started out in the main-frame world, and then moved early into the networking arena. Before joining Auburn University, he spent 22 years in the U.S. Air Force as an electronic communications and computer systems specialist, retiring in 1998. He is president of the Montgomery Windows IT Professional Group.

Ronnie Isherwood-Technical Reviewer

Ronnie Isherwood, MCITP, MCT, Member British Computer Society (MBCS), has been working in the IT industry for more than 15 years. He has 13 years of experience working as a systems engineer and as a consultant delivering server-based solutions. Additionally, he has eight years of experience working mainly with virtualization technologies. He has contributed to multiple System Center community education programs, has presented at Microsoft Private Cloud community events, participated in Microsoft Learning betas, and has presented at the E2E Virtualization conference. His technical passion is virtualization and virtualization management. Ronnie is cofounder of a local Windows user group and a committee member of the Chartered Institute for IT. Prior to becoming a systems engineer, he learned to program in basic and Microsoft Visual Basic®, delivering solutions to the financial services industry.

Telmo Sampaio-Content Developer

Telmo Sampaio is the Chief Geek at MCTrainer.NET in Miami, FL specializing in System Center, Microsoft SharePoint®, Microsoft SQL, and .NET. Telmo wrote his first application in 1984, which he with the intent of demonstrating physics concepts to his fellow classmates. His passion for technology and teaching made him a self-taught developer from early age. In 1989, he moved to Wellesley, MA when his father was transferred to work in Boston for a year. He kept developing applications to demonstrate science and math concepts, and decided to remain in the U.S. after his family left. In 1991, he moved back to Brazil and studied Systems Analysis at Pontifícia Universidade Católica do Rio de Janeiro. When Microsoft extended their Microsoft Certified Professional program to Brazil, Telmo was one of the first in the country to become certified. In 1994, he started delivering Microsoft classes. Soon he was managing the largest training center in Latin America. To date, he has been certified in over 20 different Microsoft products, passing more than 80 exams. After moving back to the U.S. in 2003, Telmo became a contributor to several Microsoft certification exams, an author for official courseware and books, and a speaker at events such as TechEd, PASS, and MMS. Telmo lives in Miami, FL with his wife Joanne, and spends his weekends with his three boys: Marco, Rafael, and Enzo. That is, when he is not traveling the world delivering training.

Damir Dizdarevic- Content Developer

Damir Dizdarevic is an MCT, MCSE, 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 MOC courses, and has published more than 400 articles in various IT magazines, such as Windows ITPro and INFO Magazine. He's also a frequent and highly rated speaker on most of Microsoft conferences in Eastern Europe. Additionally, he is a Microsoft MVP for Windows Server Infrastructure Management.

Contents

Module 1: Overview of Management in an Enterprise Data Center

Lesson 1: Overview of the Enterprise Data Center	1-2
Lesson 2: Overview of the Microsoft System Center 2012 R2 Components	1-12
Lab: Considerations for Implementing an Enterprise Data Center	1-24

Module 2: Planning and Implementing a Server Virtualization Strategy

Lesson 1: Planning a VMM Deployment	2-2
Lesson 2: Planning and Implementing a Server Virtualization Host Environment	2-10
Lab: Planning and Implementing a Server Virtualization Strategy	2-19

Module 3: Planning and Implementing Networks and Storage for Virtualization

Lesson 1: Planning a Storage Infrastructure for Virtualization	3-2
Lesson 2: Implementing a Storage Infrastructure for Virtualization	3-10
Lesson 3: Planning and Implementing a Network Infrastructure for Virtualization	3-16
Lesson 4: Planning and Implementing Network Virtualization	3-30
Lab: Planning and Implementing Virtualization Networks and Storage	3-38

Module 4: Planning and Deploying Virtual Machines

Lesson 1: Planning a Virtual Machine Configuration	4-2
Lesson 2: Preparing for Virtual Machine Deployments with VMM	4-10
Lesson 3: Deploying Virtual Machines	4-21
Lesson 4: Planning and Implementing Hyper-V Replica	4-25
Lab: Planning and Implementing a Virtual Machine Deployment and Management Strategy	4-31

Module 5: Planning and Implementing a Virtualization Administration Solution

Lesson 1: Planning and Implementing Automation with System Center 2012	5-2
Lesson 2: Planning and Implementing System Center 2012 Administration	5-7
Lesson 3: Planning and Implementing Self-Service Options in System Center 2012	5-13
Lesson 4: Planning and Implementing Updates in a Server Virtualization Infrastructure	5-21
Lab: Planning and Implementing an Administration Solution for Virtualization	5-26

Module 6: Planning and Implementing a Server Monitoring Strategy

Lesson 1: Planning Monitoring in Windows Server 2012	6-2
Lesson 2: Overview of Operations Manager	6-10
Lesson 3: Planning and Configuring Monitoring Components	6-25
Lesson 4: Configuring Integration with VMM	6-32
Lab: Implementing a Server Monitoring Strategy	6-37

Module 7: Planning and Implementing High Availability for File Services and Applications

Lesson 1: Planning and Implementing Storage Spaces	7-2
Lesson 2: Planning and Implementing DFS	7-7
Lesson 3: Planning and Implementing NLB	7-14
Lab: Planning and Implementing High Availability for File Services and Applications	7-19

Module 8: Planning and Implementing a High Availability Infrastructure Using Failover Clustering

Lesson 1: Planning an Infrastructure for Failover Clustering	8-2
Lesson 2: Implementing Failover Clustering	8-16
Lesson 3: Planning and Implementing Updates for Failover Clustering	8-22
Lesson 4: Integrating Failover Clustering with Server Virtualization	8-24
Lesson 5: Planning a Multisite Failover Cluster	8-32
Lab: Planning and Implementing a Highly Available Infrastructure by Using Failover Clustering	8-37

Module 9: Planning and Implementing a Business Continuity Strategy

Lesson 1: Overview of Business Continuity Planning	9-2
Lesson 2: Planning and Implementing Backup Strategies	9-10
Lesson 3: Planning and Implementing Recovery	9-19
Lesson 4: Planning and Implementing Backup and Recovery of Virtual Machines	9-27
Lab: Implementing a Virtual Machine Backup Strategy with DPM	9-31

Module 10: Planning and Implementing a Public Key Infrastructure

Lesson 1: Planning and Implementing Deployment of a Certification Authority	10-2
Lesson 2: Planning and Implementing Certificate Templates	10-16
Lesson 3: Planning and Implementing Certificate Distribution and Revocation	10-22
Lesson 4: Planning and Implementing Key Archival and Recovery	10-32
Lab: Planning and Implementing an Active Directory Certificate Services Infrastructure	10-36

Module 11: Planning and Implementing an Identity Federation Infrastructure

Lesson 1: Planning and Implementing an AD FS Server Infrastructure	11-2
Lesson 2: Planning and Implementing AD FS Claims Providers and Relying Parties	11-15
Lesson 3: Planning and Implementing AD FS Claims and Claim Rules	11-20
Lesson 4: Planning and Implementing Web Application Proxy	11-26
Lab: Planning and Implementing AD FS Infrastructure	11-30

Module 12: Planning and Implementing Data Access for Users and Devices

Lesson 1: Planning and Implementing DAC	12-2
Lab A: Implementing DAC and Access-Denied Assistance	12-13
Lesson 2: Planning Workplace Join	12-22
Lesson 3: Planning Work Folders	12-27
Lab B: Implementing Work Folders	12-33

Module 13: Planning and Implementing an Information Rights Management Infrastructure

Lesson 1: AD RMS Overview	13-2
Lesson 2: Planning and Implementing an AD RMS Cluster	13-8
Lesson 3: Planning and Implementing AD RMS Templates and Policies	13-19
Lesson 4: Planning and Implementing External Access to AD RMS Services	13-29
Lesson 5: Planning and Implementing AD RMS Integration with DAC	13-43
Lab: Planning and Implementing an AD RMS Infrastructure	13-46

Lab Answer Keys

Module 1 Lab: Considerations for Implementing an Enterprise Data Center	L1-1
Module 2 Lab: Planning and Implementing a Server Virtualization Strategy	L2-3
Module 3 Lab: Planning and Implementing Virtualization Networks and Storage	L3-9
Module 4 Lab: Planning and Implementing a Virtual Machine Deployment and Management Strategy	L4-19
Module 5 Lab: Planning and Implementing an Administration Solution for Virtualization	L5-25
Module 6 Lab: Implementing a Server Monitoring Strategy	L6-35
Module 7 Lab: Planning and Implementing High Availability for File Services and Applications	L7-45
Module 8 Lab: Planning and Implementing a Highly Available Infrastructure by Using Failover Clustering	L8-55
Module 9 Lab: Implementing a Virtual Machine Backup Strategy with DPM	L9-71
Module 10 Lab: Planning and Implementing an AD CS Infrastructure	L10-77
Module 11 Lab: Planning and Implementing AD FS Infrastructure	L11-91
Module 12 Lab A: Implementing DAC and Access-Denied Assistance	L12-109
Module 12 Lab B: Implementing Work Folders	L12-120
Module 13 Lab: Planning and Implementing an AD RMS Infrastructure	L13-129

About This Course

This section provides a brief description of the course-20414C: Implementing an Advanced Server Infrastructure, and includes details about the audience, suggested prerequisites, and course objectives.

Course Description

Audience

This course is intended for Information Technology (IT) professionals who are responsible for planning, designing and deploying a physical and logical Windows Server® 2012 enterprise and Active Directory® Domain Services (AD DS) infrastructure that includes network services. Candidates typically would have experience with previous Windows Server operating systems and have Windows Server 2012 certification Microsoft Certified Solutions Associate (MCSA) or equivalent skills.

The secondary audience for this course will be candidates are IT professionals who are looking to take the exam 70-414: Implementing an Advanced Enterprise Server Infrastructure, as a standalone certification, or as part of the requirement for the Microsoft Certified Solutions Expert (MCSE) certification.

Student Prerequisites

In addition to their professional experience, students who attend this training should have technical knowledge that includes an understanding of:

- TCP/IP and networking concepts.
- Windows Server 2012 and AD DS, including planning, designing and deploying AD DS and network infrastructure.
- Using scripts and batch files.
- Security concepts, such as authentication and authorization.
- Deployment, packaging, and imaging tools.
- Working on a team or with a virtual team.
- Creating proposals and making budget recommendations.
- Students should have achieved the Windows Server 2012 MCSA certification, as well as completed Course 20413B: Designing and Implementing an Enterprise Server Infrastructure, or have equivalent knowledge.

Students who attend this training can meet the prerequisites by attending the following courses, or obtaining equivalent knowledge and skills:

- 20410C: Installing and Configuring Windows Server 2012
- 20411C: Administering Windows Server 2012
- 20412C: Configuring Advanced Windows Server 2012 Services
- 20413C: Designing and Implementing an Enterprise Server Infrastructure

Course Objectives

After completing this course, students will be able to:

- Describe the considerations for managing an enterprise data center.
- Plan and implement a server virtualization strategy using System Center 2012.
- Plan and implement networks and storage for virtualization.
- Plan and deploy virtual machines.
- Manage a virtual machine deployment.
- Plan and implement a server monitoring strategy.
- Plan and implement high availability for file services and applications.
- Plan and implement a highly available infrastructure by using Failover Clustering.
- Plan and implement a business continuity strategy.
- Plan and implement a public key infrastructure (PKI).
- Plan and implement an Identity Federation infrastructure.
- Plan and implement secure data access for users and devices.
- Plan and implement an Information Rights Management (IRM) infrastructure.

Course Outline

The course outline is as follows:

Module 1, Overview of Management in an Enterprise Data Center

Module 2, Planning and Implementing a Server Virtualization Strategy

Module 3, Planning and Implementing Networks and Storage for Virtualization

Module 4, Planning and Deploying Virtual Machines

Module 5, Planning and Implementing a Virtualization Administration Solution

Module 6, Planning and Implementing a Server Monitoring Strategy

Module 7, Planning and Implementing High Availability for File Services and Applications

Module 8, Planning and Implementing a Highly Available Infrastructure by Using Failover Clustering

Module 9, Planning and Implementing a Business Continuity Strategy

Module 10, Planning and Implementing a Public Key Infrastructure

Module 11, Planning and Implementing an Identity Federation Infrastructure

Module 12, Planning and Implementing Data Access for Users and Devices

Module 13, Planning and Implementing an Information Rights Management Infrastructure

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 that is essential for an effective in-class learning experience.
- **Lessons:** These sections guide students through the learning objectives, and provide the key points that are critical to the success of their in-class learning experience.
- **Labs:** These provide a real-world, hands-on platform for students in which they can apply the knowledge and skills they learn in the module.
- **Module Reviews and Takeaways:** These provide on-the-job reference material to boost knowledge and skills retention.
- **Lab Answer Keys:** These provide step-by-step guidance for lab solutions.



Course Companion Content on the <http://www.microsoft.com/learning/companionmoc>

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

- **Modules:** These include companion content, such as questions and answers, detailed demonstration 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:** These 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 (D) version of this course is released, and students who have taken this course will be able to download the Companion Content at that time from the <http://www.microsoft.com/learning/en/us/companion-moc.aspx> site. Please check with your instructor when the 'D' version of this course is scheduled to release to learn when you can access Companion Content for this course.
- **Course evaluation:** At the end of the course, students 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, students can send an email to support@mscourseware.com. To inquire about the Microsoft Certification Program, send an email to mcphelp@microsoft.com.

Virtual Machine Environment

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

Virtual Machine Configuration

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

Important: At the end of most of the labs, you must revert the virtual machines to a snapshot. You can find the instructions for this procedure at the end of each lab. Some labs are dependent on previous labs being completed; for these labs, you will not revert the virtual machines. Ensure that you follow the steps at the end of each lab carefully.

The following table shows the role of each virtual machine that this course uses:

Virtual machine	Role
20414C-LON-Host1, LON-Host2	A Windows Server 2012 host machines (boot to vhd file)
20414C-LON-DC1	A domain controller in the Adatum.com domain
20414C-LON-SVR1, LON-SVR2, LON-SVR3	Member servers in the Adatum.com domain
20414C-LON-VMM1	A System Center Virtual Machine Manager (VMM) 2012 server
20414C-TOR-SVR1	A member server in the Adatum.com domain, in a branch office location
20414C-TOR-SS1	A Windows Server 2012 with iSCSI targets preconfigured
20414C-LON-OM1	A System Center Operations Manager 2012 (Operations Manager) server
20414C-LON-OR1	A System Center Orchestrator (Orchestrator) server
20414C-LON-DM1	A System Center Data Protection Manager 2012 (Data Protection Manager) server
20414C-LON-WSUS	A Windows Server Update Services server.
20414C-LON-CA1	A standalone server
20414C-LON-CL1	A client computer with Microsoft® Office 2013 in the Adatum.com domain
20414C-LON-CL2	A standalone client computer with Microsoft® Office 2013

Virtual machine	Role
20414C-LON-CORE	A standalone server running the Windows Server 2012 Server Core installation option.
20414C-TREY-DC1	A domain controller in the TreyResearch.net domain
20414C-TREY-CL1	A client computer with Microsoft® Office 2013 in the TreyResearch.net domain

Software Configuration

The following software is installed on each virtual machine:

- Windows Server 2012 R2
- Microsoft Windows 8.1
- Microsoft System Center 2012 R2

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.

These courses will require Hardware level 7. The additional resources are required because of the inclusion of Windows System Center 2012. All virtual machines will be built and run in Hyper-V in Windows Server 2012. The host machine for this course is provided in a Boot From VHD (Native Boot) configuration.

Hardware Level 7 is as follows. Note the changes for the Hard Disk configuration.

- 64 bit Intel Virtualization Technology (Intel VT) or AMD Virtualization (AMD-V) processor (2.8 gigahertz [Ghz] dual core or better recommended).
- Hard Disk: Dual 500 GB hard disks 7200 RPM SATA labeled C drive and D drive.
- 16 GB random access member (RAM).
- DVD: dual layer recommended.
- Network adapter
- Sound card
- Video adapter: Supports 1440X900 resolution
- Monitor: Dual SVGA monitors 17" or larger that supports 1440X900 minimum resolution

Module 1

Overview of Management in an Enterprise Data Center

Contents:

Module Overview	1-1
Lesson 1: Overview of the Enterprise Data Center	1-2
Lesson 2: Overview of the Microsoft System Center 2012 R2 Components	1-12
Lab: Considerations for Implementing an Enterprise Data Center	1-24
Module Review and Takeaways	1-26

Module Overview

In many large companies, recent years have brought dramatic changes to the enterprise data center hosting many of the IT services. One of the most significant changes is the introduction of virtualization. Currently, the default deployment for new servers is often a virtual machine rather than a physical machine. Data centers have changed in other ways, as organizations need to deal with changing business requirements.

These changes mean that the enterprise data center is now more complex, and needs to be much more responsive to changes. Managing the data center has become more complex as organizations seek to optimize performance and provide features like self-service. These requirements drive the need for a new set of management tools to manage the data center.

This module describes some of the changes and new requirements that organizations are experiencing in their data centers. This module also describes how you can use Microsoft® System Center 2012 R2 to manage this environment.

Objectives

After completing this module, you will be able to:

- Describe the enterprise data center.
- Describe how you can use System Center 2012 R2 to manage the enterprise data center.

Lesson 1

Overview of the Enterprise Data Center

Many large organizations have deployed one or more central data centers to provide most of their information technology (IT) services. These data centers are adapting to address changing business requirements. This lesson describes some of the changes that organizations have made to their data centers, and how new requirements for IT processes and tools have evolved from these changes.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe an enterprise data center.
- Describe the importance of virtualization in the data center.
- Explain extending the enterprise data center to the public cloud.
- Describe the tools needed to manage and monitor the data center.
- Describe the requirements and options for providing security in the data center.
- Describe the importance of providing monitoring and reporting in the data center.
- Describe options for implementing high availability and business continuity in the data center.
- Describe options for automating the data center.

What Is an Enterprise Data Center?

An enterprise data center is a centralized location from which organizations provide the necessary IT services for accomplishing their business requirements and goals. Most enterprise data centers provide a similar set of core services, including:

- Infrastructure services. Enterprise data centers provide a core set of services that are required for all other services to operate. These include the core network infrastructure and network services such as Dynamic Host Configuration Protocol (DHCP) and Domain Name System (DNS).
- Authentication and security services. Most organizations use Active Directory® Domain Services (AD DS) to provide a central security database for user authentication and authorization. Organizations may need to provide additional security services for authenticating users to applications that are not AD DS–integrated or for authenticating users who do not have AD DS accounts.
- Central storage of data. Almost all enterprise data centers provide centralized storage of data. This data may be stored on shared folders on file servers, in databases, or on websites. Users throughout the organization need secure access to this data.

- Enterprise data centers provide core services for organizations, including:
 - Infrastructure services
 - Authentication and security services
 - Central storage of data
 - Messaging and collaboration services
 - Server-based applications
- Enterprise data centers must adapt so they can:
 - Provide services to a wide variety of clients
 - Enhance security
 - Optimize resources
 - Provide rapid deployment of new services
 - Ensure that services are always available

- Messaging and collaboration services. Email is an essential part of most business processes. In some organizations, other collaboration tools, such as instant messaging, desktop conferencing, and other interactive web-based tools, are becoming more important. The enterprise data center provides access to these services.
- Server-based applications. The enterprise data center provides access to a wide variety of server-based applications, including web-based applications, database applications, remote desktop-based applications, or even virtual machines that users can access.

The core set of services that enterprise data centers provide has not changed significantly. However, the way in which IT organizations provide these services has changed radically. Typically, the services that enterprise data centers provide, and the clients who consume the services, have been easy to predict and manage. Initially, most data center services were for internal users, who connected to the services by using desktop computers connected to a wired network. Organizations deployed services on physical servers and IT departments understood how to deploy and manage new services that the organization required. The data center's boundary and the clients that it supported were easily definable and fairly static.

However, in the last 10 years, data centers have changed significantly in the way they provide their services. Now data centers need to provide services to clients that connect to the data center from internal networks but also from locations and networks outside the organization. Data center clients may include mobile devices. This means that organizations now need to provide access to internal services from the Internet, while providing much higher levels of security. Organizations have become much more conscious of the cost and the environmental impact of running large data centers, so they ask their IT departments to optimize the performance and utilization of all data center components. Successful organizations can adapt rapidly to changing business requirements, and IT departments must be able to roll out necessary new services efficiently. Organizations are allowing a larger and more diverse group of users to access their IT services. Therefore, it is becoming increasingly important that these services are always available. In addition, organizations must develop highly effective business continuity plans.

Considerations for Virtualizing an Enterprise Data Center

One of the central features of an enterprise data center is virtualization, especially server virtualization. Most organizations, except for the smallest, deploy almost all new servers as virtual machines. An environment that is almost completely virtualized provides many benefits:

- Increased utilization of hardware. The primary benefit that most organizations experience with virtualization is that they can fully utilize server hardware. Organizations can deploy a smaller number of virtualization hosts as physical machines and then deploy multiple virtual machines to each host. This enables organizations to utilize their hardware resources fully, instead of barely utilizing multiple physical machines. Because they deploy fewer physical machines, organizations experience significant decreases in data center power and cooling requirements.
- Business agility. Another important benefit of virtualization is the ability to respond rapidly to changing business requirements. In a virtual environment, an IT department can address a request from a business group to have a new server, set of servers, or application deployed in hours rather than days or weeks. If you need to move a server to a different location, it is a simple matter of copying files across a network.

- Virtualization in the data center provides the following benefits:
 - Proper utilization of hardware
 - Business agility
 - Options for high availability
 - Options for administrative flexibility
- Virtualizing the enterprise data center adds management complexity and requires optimized processes and tools

- Options for high availability. Server virtualization also provides important new options for implementing high availability. You can use the same high availability options in a virtual environment as in a physical environment. For example, you can deploy virtual machines that use network load balancing (NLB), failover clustering, or application high availability. Virtualization provides the additional option of making the virtual machine highly available so that you can provide high availability for services and applications that do not natively support high availability.
- Options for administrative delegation, self-service, and automation. Virtual environments provide many options for enhancing the management processes within a data center. In a virtual environment, you can easily give business users access to only one or two virtual machines. With the right management tools and processes, you can enable these users to manage their small part of the data center, including options for creating additional servers or services.

While deploying a virtual environment can provide significant benefits, it also adds a level of complexity to the management processes. Managing a virtual environment requires that you manage the physical or host layer as well as the virtual layer. Largely, you can manage the layers separately, but the layers are still dependent on each other. If you mismanage the physical layer and a physical host shuts down accidentally, you may affect many business groups whose virtual machines were running on that host. If you mismanage virtual machines and one virtual machine begins to consume too many host resources, other virtual machines may be affected.

Virtualization provides very significant benefits for almost all organizations. To manage the virtual environment properly, you need to implement processes and tools that maximize the benefits of virtualization while minimizing the management complexity.

Considerations for Extending the Enterprise Data Center to the Public Cloud

Organizations utilize a private cloud in the data center to provide a flexible infrastructure that can scale on demand. Using a private cloud allows you to manage resources in an agile and effective way. Extending your data center into the public cloud is called the *hybrid cloud* model. By using a hybrid cloud, you can take advantage of external resources when it makes sense for your business. With Windows Azure™, Windows Server®, and System Center 2012 R2 you have a hybrid cloud that allows you to manage both the private cloud and public cloud with the same tools.

- Extending your private cloud into Windows Azure provides many benefits, including:
 - Scale on demand: You add or remove virtual machines as necessary
 - Hybrid applications: You develop applications that use both the public and private clouds
 - Development and testing: Provides an environment for development and testing
 - Backup: Provides a backup location that you can access from anywhere
 - Consistent toolset: You use the same toolset to manage both Windows Azure and your private cloud

Extending your private cloud into Windows Azure allows you to move workloads easily between the private and public clouds, while maintaining a complete view of the infrastructure. By using Windows Azure, you can take advantage of cloud resources and experience these benefits:

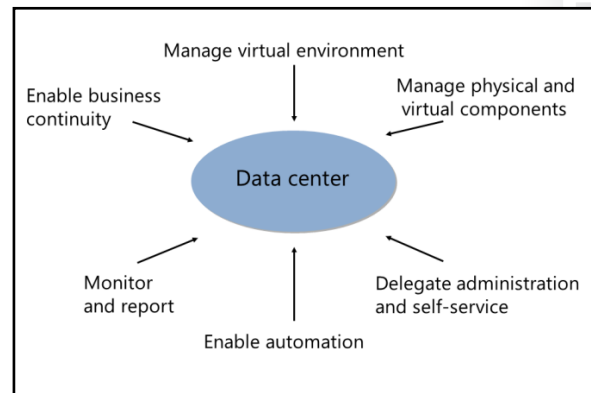
- Scale on demand. Windows Azure infrastructure as a service (IaaS) scales up or down depending on your needs. You can base your virtual machines on the Windows Azure templates or create your own images and deploy them when necessary.
- Hybrid applications. You can develop applications that use resources in both Windows Azure and your private cloud. For example, you can develop a web-based application in Windows Azure that references a Microsoft SQL Server® database in the private cloud.

- Development and testing. You can use the Windows Azure platform to deploy applications temporarily in Windows Azure for development and testing. You can build Windows Azure virtual machines quickly and take them down easily as you work through the development phase of an application.
- Backup. Using Windows Azure as a backup location can improve your disaster recovery efforts because your backups will be available anywhere.
- Consistent toolset. You use the same tools to manage both Windows Azure and your private cloud. You can use System Center 2012 R2 and Windows PowerShell® to manage both environments.


Choosing Tools for Managing and Monitoring the Data Center

As the enterprise data center has changed, so have requirements for tools to manage the data center. In the old data center model, you could manage most components by using specialized tools, with little concern for how those management tools interacted with other components in the data center.

As data centers have become more complex, the need for better-integrated management tools has grown. Some examples of how tool requirements have changed include:



- Tools must be optimized for managing virtual environments. In today's data center, almost all new servers are deployed as virtual machines. Many applications are also deployed as virtual applications by using presentation virtualization, application virtualization, or desktop virtualization. Any tool that you use to manage a data center must be able to manage all components of the virtualized data center.
- Tools must manage both virtual and physical components of the data center. Although most new servers and applications are virtualized, all data centers still require a physical computing layer. This might include servers and applications running on physical hardware and will definitely include the host layer for the virtualized environment. The management tools should be able to manage both the physical layer and the virtual layer by using the same administrative consoles and management processes.

 **Note:** Choosing the right operating system to run within virtual machines and on the physical hosts is an important part of the planning process. Windows Server 2012 R2 optimizes integration with management solutions like Microsoft System Center 2012 R2.

- Tools must manage both the private cloud and public cloud. As more companies extend their data centers into the public cloud, administrators should be able to manage both environments by using the same toolset. Windows Azure and Windows Server 2012 allow you to use the same tools to manage both environments.
- Tools must provide options for delegating administrative functions, including self-service. One of the key features in a data center is the ability to respond quickly and effectively to changing business requirements and situations. In this environment, it is inefficient to have a single group of enterprise administrators manage all components. To enable a rapid response to changing situations, the management tools should provide the efficient and secure delegation of tasks to other

administrators. Alternatively, the tools should enable the delegation of self-service to business units that can address changing requirements without interacting with the enterprise administrators who configured the tools.

- Tools must provide means for automating processes. To enable features like self-service, and to enable efficient and error-free management of the data center components, the management tools need to provide ways to automate many of the management tasks. The tools should enable these tasks to run automatically on a scheduled basis, when a user action initiates them, or in response to a monitoring alert. The tasks that an automated process triggers may interact with one or many data center components.
- Tools must enable monitoring and reporting for all components in the data center. As the complexity of the data center increases, so does the importance of monitoring all its components and providing intelligent reporting of their current state. Administrators cannot monitor all components individually; they require tools that capture monitoring information from all components and then report only the relevant information. Many components will have dependencies on other components and the monitoring tools need to recognize these relationships.
- Tools must provide the means to ensure business continuity. In the enterprise data center, some services and components need to be available all the time. Any unscheduled downtime in these components may have a significant business impact. The management tools should provide the means to apply updates to services and components without affecting the service availability. They should also make it possible to recover a service in the event of a service failure.

Considerations for Providing Secure Services in the Data Center

In a typical data center, AD DS provides most of the functionality that is necessary for internal users to sign in to the network and access network services and applications securely. If all users have AD DS accounts and sign in to the network using only internal computers that are members of the same AD DS environment, it is easy to provide secure access to data center services and applications. However, many organizations now require that many different types of users access resources in the data center by using a wide variety of devices. Users who need access to data center resources may include employees who are traveling or working from home. They may also include users who are not employees, but who are customers or who work for partner organizations. Users may be using managed laptop computers or unmanaged computers and mobile devices to access resources. The organizations they work for may have security requirements such as:

- Data center security requirements may include:
 - Two-factor authentication
 - Computers that are authenticated for VPN and application access
 - Access to internal applications for external contractors
 - Encrypted website traffic
 - Secured external-facing applications
 - Protection of information after it leaves the organization
- You can meet these requirements by deploying:
 - AD CS
 - AD FS
 - AD RMS
 - AD LDS
 - Windows Azure AD

- The organization's security group may require that administrators who perform domain administrator or enterprise administrator tasks on the network provide two factors of authentication. That is, providing only a user name and password may not provide an adequate level of security.
- The organization may require users to use laptop computers that are members of the internal Active Directory domain if they want to sign in to certain highly secured websites or connect to the internal network through a virtual private network (VPN).
- The organization may allow only employees to have user accounts in the AD DS domain. However, the organization needs to provide access to websites and applications so that contractors can work both internally and externally.

- The organization may need to ensure encryption of all authentication and web access traffic to Internet-facing websites.
- The organization may have deployed a product-ordering application for external customers. This application must be secure, so that only authorized users can access the application and place orders.
- The organization may have implemented security rules about how to share specific types of information with users inside and outside the company. The organization needs a method for enforcing these security requirements.

It can be difficult to meet all these security requirements with AD DS alone. AD DS may not provide the required functionality. Even if it does, most organizations do not expose their internal AD DS domain controllers directly to the Internet because of security concerns.

Microsoft has developed additional Active Directory services to address the business and security requirements that AD DS alone cannot address. These services include:

- Active Directory Certificate Services (AD CS), which provides a public key infrastructure (PKI) that you can use for issuing and managing certificates that provide additional levels of security.
- Active Directory Federation Services (AD FS), which enables partner organizations to establish federated trusts. Then these organizations can manage their own digital identities and access applications that the other organization hosts.
- Active Directory Rights Management Services (AD RMS), which provides protection and access control to sensitive information, such as documents and email messages. You can apply this protection even after the content has left the organization.
- Active Directory Lightweight Directory Services (AD LDS), which is a Lightweight Directory Access Protocol (LDAP) directory that provides organizations with flexible support for directory-enabled applications. This is useful if organizations do not want or require the management overhead of a full AD DS deployment.
- Windows Azure Active Directory (Windows Azure AD) is a Windows Azure–based implementation of AD DS. It is an Active Directory service provided and managed by Microsoft to provide web-based application authentication. Customers manage only their users in Windows Azure AD. This is useful for organizations that want to provide authentication for public cloud–based resources without exposing their internal AD DS infrastructure to the Internet.



Active Directory Lightweight Directory Services Overview

<http://go.microsoft.com/fwlink/?LinkID=286067>

These additional Active Directory services extend AD DS functionality and features, but you can also integrate them with AD DS. For example, you can configure AD CS to issue computer certificates automatically to all computers that are members of an AD DS domain. You can configure AD LDS to synchronize some information with AD DS. Users can take advantage of AD FS and AD RMS services only after signing in to their AD DS accounts.

Considerations for Monitoring and Reporting on Data Center Services

An effective monitoring and reporting tool is critical in an enterprise data center. The data center will have many different services and applications deployed, and the services change frequently. In this environment, it is essential to understand what you should monitor and to provide the right monitoring tools. The key areas of the enterprise data center that require monitoring are:

- Physical and virtual machines. Almost all data center services and applications run on physical or virtual machines, so it is critical that enterprise administrators are aware of the performance of all computers in the data center. You must monitor physical hardware on the hosts and virtual hardware on the guests for performance and availability. Physical hardware includes the disk subsystem, devices that are attached to storage area networks (SANs) or network-attached storage (NAS), and other hardware devices.
- Core network services. Almost all data center services require a core set of services in order to function. These core services include the network infrastructure, DNS, and AD DS. If these services are not functional, or are overloaded, then no other services will function as expected. The enterprise monitoring system must be able to monitor the core services for performance and availability and provide alerts when issues arise.
- Core applications. Many data center applications depend on a core set of provided applications. A central team of administrators may deploy these applications, including messaging servers, database servers, and web servers. These applications are accessible to all users, and all business applications may share the same infrastructure. Because business processes and business applications depend on these core applications being available and providing adequate services, you must monitor these applications.
- Security. As the boundaries of the data center have become less defined, and the categories of clients who access the data center have become more varied, the role of security monitoring has increased greatly. Many different types of clients, including unmanaged clients, may be connecting to the data center. The data center is now designed to provide services for users who are anywhere, not just users inside the organization's physical building. As a result, it is critical that tools be available to monitor for security breaches and to ensure that any security breach raises an instant alert.
- Compliance with service level agreements (SLAs) or standards. Many organizations have defined SLAs that specify availability and performance requirements for data center components. Data center management tools should be able to monitor and report on compliance with these SLAs and possibly other organizational and industry standards.
- Business applications. In addition to monitoring the core shared infrastructure and applications, it is important to monitor the business applications individually. This ensures that the business application is providing the required services, but also that a single business application is not consuming more than its share of the shared components. Monitoring a single business application may require monitoring multiple items. For example, a single application may use a hardware load balancer, several web servers, a clustered database server, and file servers. Monitoring this application will require a tool that can monitor each component and understand the dependencies between each component.

Monitoring tools must be able to monitor:

- Physical and virtual machines
- Core network services
- Core applications
- Security
- Compliance
- Business applications

Enterprise data center monitoring tools must be able to:

- Collect data from multiple systems
- Apply intelligence to collected data
- Provide historical data

Many hardware vendors provide monitoring tools to monitor the physical hardware and operating system components. Windows Server provides monitoring tools, such as Performance Monitor, that can gather

detailed information about the performance on any operating system and many application components running on the server. However, typically these tools can monitor single servers only, and can only collect, and not intelligently analyze, the gathered information.

The enterprise data center requires tools that can do the following:

- Collect detailed monitoring information from multiple systems and applications. Monitoring tools must be able to collect information about all the components, services, and applications that are deployed in the data center. This may include hardware from many different vendors, different types of operating systems, and different types of applications.
- Apply intelligence to the collected information. Collecting the information alone is not enough, because the raw data is not meaningful, generally. You must analyze and correlate data from different sources to extract meaning. The monitoring tools must provide this intelligence and be configurable to ensure that the information is meaningful for each application.
- Provide historical data. Monitoring tools should provide the means to store and report on past performance information. This is necessary for understanding how the performance characteristics of an application may have changed over time, and for providing information for future planning.

Considerations for High Availability and Business Continuity

The enterprise data center provides services to a wide variety of potential clients, including organization employees who are internal and external, customers, and business partners. These clients may access the data center services at any time and from any place. Because of this, it is critical that the services are highly available and very responsive. In the case of a service outage, it is crucial that the service be restored as quickly as possible.

- To implement high availability:
 - Deploy redundant components
 - Enable automatic recovery
 - Implement Windows Server 2012 features
- To implement business continuity:
 - Create a backup plan
 - Create a recovery plan
 - Understand the dependencies between the two plans

Implementing High Availability

In order to provide high availability, a system must be able to survive the failure of one or more components without causing a disruption in a service or application. In order to ensure high availability, you should:

- Deploy redundant components. If a system is to survive the failure of individual components, you need to deploy multiple components that provide the same function. Some components can provide redundancy at the individual server level. For example, you can deploy servers with multiple network adapters and multiple power supplies. You can deploy other components, such as multiple network switches, Internet or wide area network (WAN) connections, and power sources, at the data center level. In some cases, you can provide redundancy by deploying multiple physical or virtual servers that provide the same functionality. In some cases, a single level of redundancy may be sufficient; in other cases, you may want to protect against multiple failures of the same component.
- Implement automatic failover or recovery. To provide a highly available solution, it is important that the system can fail over or recover from the loss of any single component. The system should be able to detect when a component fails and recover automatically from that failure. Usually, this recovery involves shifting the services that the failed component provides to a redundant component. It is important that this failover happens automatically and that it does not require administrative attention or effort.

In many cases, a highly available solution will also provide load balancing. This means that, when all redundant components are functional, the service requests from clients will spread across all components. For example, when all of the web servers in a highly available deployment are functional, client requests for a website on the web servers will spread across all available servers.

There are many different ways to provide high availability. Within an enterprise data center, it is very common to see redundant power sources, redundant cooling systems, redundant network connections, and a redundant storage infrastructure.

Windows Server 2012 provides several options for implementing high availability including:

- NIC Teaming, which you can use to manage redundant network connections.
- Storage spaces, which you can use to provide redundancy and automatic failover for storage devices attached to the server.
- Network Load Balancing, which provides high availability and load balancing for web-based services.
- Failover clustering, which provides high availability for many services, including Windows Server 2012 Hyper-V® virtual machines.
- Distributed File System (DFS), which you can use to create multiple copies of shared folders and files to optimize performance and provide improved availability.

Implementing Business Continuity

Ideally, the implementation of high availability in the data center will reduce the effects of any potential disruption in services. However, it is critical that organizations have the means to recover from a failure beyond the scope of the high availability solution. For example, multiple components or even an entire data center may fail.

There are many different ways to provide business continuity, depending on the type of disasters against which the organization is protecting. Many large organizations have deployed disaster recovery data centers and implemented plans for a complete data center failover in the event of a data center failure. Almost all organizations perform regular backups of their data centers and store the backups offsite to ensure that they can recover from the loss of the data center. For example, you could use Windows Backup to restore backups at an offsite recovery center in the case of a disaster. If you use virtualization, you can create Hyper-V replicas that would allow quick recovery of servers in a disaster.

When creating a business continuity plan, you should include:

- A backup plan. You must back up all data that is important to the organization. In addition, you should back up any computer configuration information that is important in restoring enterprise data center services. For example, it is critical to back up AD DS domain controllers, especially if all domain controllers are located in the same data center.
- A recovery plan. The recovery plan details how you will recover data or services in the event of data loss or service failure. The recovery plan should provide detailed information on what is required to perform the recovery and how to perform the recovery.

When creating the backup and recovery plan, you must consider the dependencies between the two plans. Frequently, business continuity plans include SLAs for the recovery point objects and recovery time objectives. Recovery point objects define the amount of accessible data loss in the event of a catastrophic failure. For example, the recovery point object might state that no more than four hours of data should ever be lost if a file server hosting file shares fails. With this recovery point object, you must make sure that you back up the data at least every four hours. Recovery time objectives define how much time it will take to recover data or a service. If you have a one-hour recovery time objectives for the file servers, then you may need to perform the backup to disk rather than to tape, in order to restore data quickly enough to meet the recovery time objective.

Considerations for Automating Data Center Solutions

As the enterprise data center grows and becomes more complex, it can become inefficient to manage all components manually within the data center. To make your management more efficient, you should explore options for automating as many repetitive tasks as possible. Some options for automation include:

- Automating management processes. Network administrators can automate many of the most common tasks that they perform. Most commonly, you might automate tasks in which you provision new objects or components. For example, in a large corporation, the process of creating new user accounts can be a time-consuming and complex task. This is because you may need to assign the user to multiple groups, configure multiple attributes, and create the account on multiple systems. You can automate this task easily by deploying a tool such as Microsoft Forefront® Identity Manager, which you can configure to perform all of these tasks based on a single request. Another task that administrators automate frequently is the provisioning of new virtual machines. By precreating templates and using features like intelligent placement in System Center 2012 R2 Virtual Machine Manager (VMM), you can trigger the creation of new virtual machines with minimal effort.
- Automating service delivery and change management. Many organizations use automation to provide different ways to start management processes. Rather than have an administrator initiate the process, you can provide tools for initiating service requests to end users or business groups. These service requests will then initiate automated management processes automatically. For example, you might give business groups access to a tool like System Center 2012 R2 Orchestrator or System Center 2012 R2 App Controller. Then you could allow business group administrators to initiate the creation of a new set of virtual machines that they need to deploy a new application.
- Automating incident and problem management. Another option for automating data center tasks is to automate problem resolution, which may require integration of several different data center components. For example, a monitoring system may detect that a virtual machine or application is consuming more resources than normal. The monitoring system may trigger an alert, which the virtual machine management system will consume. Then the virtual machine management system may move that virtual machine to another host with more resources available. If an application is consuming more resources than usual, or if the application has stopped responding, the alert may prompt an application restart. System Center 2012 R2 Service Manager provides this type of functionality.

Automating data center solutions can include:

- Automating management processes
- Automating service delivery and change management
- Automating incident and problem management

Automation requires a high level of integration between tools

Automating business processes almost always requires the integration of multiple components and tools within the data center. This sort of integration is one of the key features of System Center 2012 R2.

Lesson 2

Overview of the Microsoft System Center 2012 R2 Components

Within a data center, there is a high level of interdependence between the different components. Therefore, it is critical that the tools you use to manage the data center are also integrated. You must be able to use these tools to manage all components and provide all required services.

System Center 2012 R2 is designed to be a complete management suite for the enterprise data center. With System Center 2012 R2, Microsoft has brought all of the System Center components into a single product with increased integration between the components. This lesson introduces the various System Center 2012 R2 components and describes how you can use these tools to manage the enterprise data center.

Lesson Objectives

After completing this lesson, you will be able to:

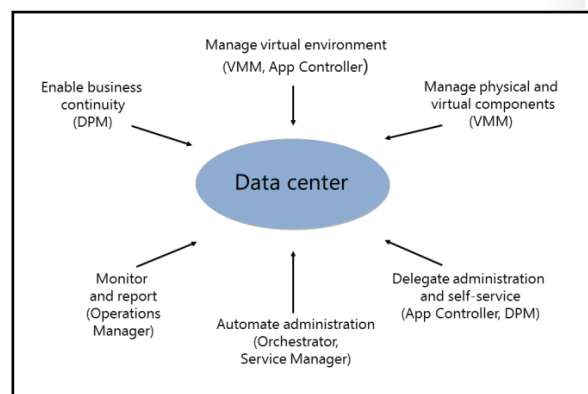
- Describe how you can use System Center 2012 R2 to manage the enterprise data center.
- Describe System Center 2012 R2 Configuration Manager.
- Describe System Center 2012 R2 Virtual Machine Manager.
- Describe System Center 2012 R2 App Controller.
- Describe System Center 2012 R2 Service Manager.
- Describe System Center 2012 R2 Orchestrator.
- Describe System Center 2012 R2 Operations Manager.
- Describe System Center 2012 R2 Data Protection Manager.
- Describe additional tools for managing the data center.

Using Microsoft System Center 2012 R2 to Manage the Enterprise Data Center

System Center 2012 R2 provides the tools to deploy, manage, operate, automate, and monitor data centers.

The following is a list of the System Center 2012 R2 components:

- **Configuration Manager.** Configuration Manager provides software management, operating system deployment, and mobile device management. You can use Configuration Manager for advanced update and third-party update management. In addition, it includes System Center 2012 R2 Endpoint Protection. Endpoint Protection is the enterprise anti-malware solution for client devices. It gives administrators the ability to manage client security from a single console.



- Virtual Machine Manager (VMM). VMM provides administrators with a single administrative tool for deploying and managing a virtualization infrastructure, including components such as hosts, storage, networks, libraries, and update servers. This infrastructure provides the foundation for managing the configuration and deployment of virtual machines.
- App Controller. App Controller provides a self-service portal for administrators who are deploying and managing applications and services across one or more sites. App Controller enables you to access and manage resources from one or more VMM management servers and from multiple Windows Azure subscriptions.
- Service Manager. Service Manager offers service management, process automation, asset tracking, and a self-service portal to access resources defined in a service catalog. Service Manager offers an easy-to-build configuration management database, which pulls data from AD DS and System Center components. This allows companies to establish and use controls and operations based on the guidelines of the Information Technology Infrastructure Library or the Microsoft Operations Framework.
- Orchestrator. Orchestrator is a runbook automation component that allows administrators to integrate and automate their data centers. Orchestrator utilizes integration packs, including many out-of-box authored packs that allow administrators to connect different systems.
- Operations Manager. Operations Manager is the management component for application and performance monitoring. You can integrate Operations Manager with VMM, Service Manager, Orchestrator, and Data Protection Manager. Operations Manager utilizes vendor-authored management packs that provide deep application insight and health state monitoring.
- Data Protection Manager (DPM). DPM is an enterprise backup component that performs application-aware block-level backups. It utilizes Volume Shadow Copy Service (VSS) writers to help protect and recover applications such as SQL Server, Exchange Server, Microsoft SharePoint® Server 2012, and AD DS. Additionally, it provides specific VSS writers for System Center 2012 R2 components.

Overview of Configuration Manager

Configuration Manager provides inventory management, software management, operating system deployment, and mobile device management. Configuration Manager provides several key functionalities and benefits, including:

- Automation of operating system deployment. This includes providing administrators with tools to create and deploy operating system images to computers using a Pre-Boot EXecution Environment (PXE) or portable media, such as universal serial bus (USB) flash drives or DVDs.
- Deployment of software applications. This enables administrators to manage deployment of applications to users across devices, such as desktops, servers, laptops, and mobile devices. Furthermore, administrators can create an application catalog and publish it to a self-service website.
- Management of software updates. This includes advanced management, monitoring, and deployment of software updates.
- Advanced grouping of resources and remote control tools. This allows administrators to control clients remotely from the Configuration Manager Console.

Configuration Manager features include:

- Automation of operating system deployment
- Deployment of software applications
- Management of software updates
- Remote control
- Endpoint Protection
- Compliance settings management
- Asset Intelligence and inventory

- Endpoint Protection. Endpoint Protection in Configuration Manager enables the management of device security, Windows Firewall, and the detection and removal of malicious software.
- Management of compliance settings. The compliance settings feature allows you to track, assess, and remediate the configuration compliance of devices such as servers, workstations, laptops, and mobile devices.
- Asset Intelligence and inventory. Configuration Manager provides the ability to collect information about the hardware and software in your organization. You can also monitor software license usage.

Other features within Configuration Manager include:

- Administration of power features and plans.
- Management of client health and monitoring.
- Management of virtual desktops.
- Management of Exchange ActiveSync®-enabled mobile devices.

System Center 2012 R2 Configuration Manager includes new features and capabilities, including support for:

- Windows® 8.1
- Deployment of Windows 8.1 apps
- Real-time administration for Endpoint Protection-related tasks
- Integration with Windows Intune™



For more information about Configuration Manager capabilities, go to:

<http://go.microsoft.com/fwlink/?LinkID=393715>

Overview of Virtual Machine Manager

VMM provides administrators with a single administrative tool for deploying and managing a virtualization infrastructure, including components such as hosts, storage, networks, libraries, and update servers. This infrastructure provides the foundation for managing the configuration and deployment of virtual machines. You can use VMM for managing a single virtual machine host computer, or as many as 400 hosts and 8,000 guests.

VMM consists of a VMM management server, VMM database, and Virtual Machine Manager console. A deployment requires these core components. You can deploy them to a single server or to multiple servers.

The following are some of the key VMM features:

- Bare-metal deployment of hosts. You can automate deployment of Windows Server host machines that have the Hyper-V server role installed on physical servers with an installed baseboard management controller (BMC) and that meet discovery and deployment prerequisites. System Center 2012 R2 enables administrators to discover more information about a target host's resources and configure more networking settings, such as logical switches.

VMM features include:

- Bare-metal deployment of hosts
- Host and cluster creation
- Host groups
- Cross-platform management
- Storage configuration/network configuration
- Intelligent placement/dynamic optimization
- Power optimization
- PRO
- Usage metering & reporting

- Host and cluster creation. You can create Hyper-V hosts and clusters easily by using the Virtual Machine Manager console, which simplifies manual deployment and reduces the possibility of configuration errors.
- Host groups. You can group hosts for logical separation, such as business use, performance, and geographical location, and you can apply changes to multiple hosts.
- Cross-platform management. VMM supports the management of Citrix XenServer hosts and pools and supports VMware ESX hosts through integration with VMware vSphere.



Note: System Center 2012 R2 VMM can manage only Windows Server 2008 R2, Windows Server 2012, or Windows Server 2012 R2 hosts. You cannot manage Windows Server 2008 host machines by using System Center 2012 R2 VMM. System Center 2012 R2 VMM supports only ESX or ESXi 4.1 or newer versions. You can also manage Citrix XenServer 6.0 or newer versions by using System Center 2012 R2 VMM. However, first you must install the System Center Integration Pack on the Citrix XenServer hosts and then add the hosts in VMM.

- Storage configuration. VMM supports the discovery, classification, and provisioning of storage for your Hyper-V hosts, including thin provisioning capabilities. VMM storage discovery works with Storage Management Initiative Specification, Common Information Model (CIM) XML, and symmetric multiprocessing (SMP) storage providers. Additionally, VMM 2012 SP1 supports the new Windows standards-based Storage Management Service and the Server Message Block (SMB) 3.0 protocol.
- Network configuration. Network configuration enables the creation of logical networks, media access control (MAC) address pools, and supported load balancers. Additionally, VMM supports the new Windows Server 2012 network virtualization features, including the ability to run overlapping addresses on the same physical network.
- Intelligent placement. This helps you select an appropriate host, based on the virtual machine that you are deploying, and includes ratings of hosts against expected utilization thresholds, such as percentage of central processing unit (CPU), I/O, and network throughput.
- Power optimization. You can configure VMM to use thresholds that you specify. This enables evacuation of underutilized hosts that belong to a cluster. Then you can turn off these hosts to conserve energy. As demand increases, you can turn on these hosts again.
- Performance and Resource Optimization (PRO). PRO tips, which are a feature of integrating Operations Manager with VMM, can offer preset remediation based on alerts. For example, you can use PRO tips to initiate the live migration of a virtual machine from a heavily utilized host machine to a host machine with more capacity.
- Microsoft Server Application Virtualization (Server App-V). Server App-V allows the virtualization of server-based applications. VMM has a built-in service and template designer that can help you construct single and multitier applications. Then you can deploy these applications as services, which you can scale through automation.
- Usage metering & reporting. VMM integrates with Operation Manager to enable reporting on all virtual machines that VMM manages. When you deploy the Windows Azure Pack for System Center, you can also report on usage information in Windows Azure.



For more information about new features in System Center 2012 R2 VMM, go to:

<http://go.microsoft.com/fwlink/?LinkId=253224>

Overview of App Controller

App Controller is a self-service portal that enables administrators and end users to control, deploy, and configure applications and virtual machines across VMM deployments and public clouds.

App Controller is the replacement for the VMM self-service portal, which System Center 2012 SP1 and newer versions no longer include.

When you build a virtualization environment, you do not want to give everyone in your IT department access to your systems management and virtualization host servers. Furthermore, you do not want to deploy and maintain management consoles for every vendor or IT employee that may require occasional interaction with a virtual machine. If employees access a subset of a virtual machine deployment regularly, you should grant them permissions only for the tasks that they need to perform, and make visible only the systems that they need. App Controller is a web portal that provides that access, in addition to self-service capabilities that enable administrators to deploy and administer resources across multiple VMM management servers. Administrators can also deploy and administer Windows Azure and service-provider data center resources through App Controller.

App Controller

- Replaces the deprecated VMM self-service portal
- Provide delegated access to private and public cloud resources, such as:
 - Virtual machines
 - Services
 - Templates, images
- Allows administrators to migrate between VMM, Windows Azure, and service provider data centers

Managing VMM Instances

You can configure App Controller to use up to five VMM management servers and their resources. App Controller provides web-based access through which you can control applications, virtual machines, and their resources, including libraries and shares.

Managing Public Cloud Resources

App Controller can control as many as 20 Windows Azure subscriptions. It allows you to upload VHDs and images to Windows Azure from a library or network shares and add virtual machines to deployed services in Windows Azure. Additionally, you can manipulate and migrate virtual machines to and from Windows Azure.

In System Center 2012 R2, App Controller allows you to add a Service Provider Framework hosting-provider connection. By using this connection, you can connect your App Controller instance to a service configured by a hosting provider that uses the System Center 2012 R2 Service Provider Foundation. Service Provider Foundation enables service providers to offer IaaS to their clients.

App Controller allows you to:

- Use role-based access to VMM resources.
- Use role-based access to Windows Azure subscription resources.
- Upload VHD files and images from the VMM library to Windows Azure.
- Migrate a virtual machine from VMM to Windows Azure.

Overview of Service Manager

You can use the Service Manager component for automating business processes and implementing service management as defined in the Information Technology Infrastructure Library and the Microsoft Operations Framework. There are many out-of-box processes for change, release, lifecycle, and incident and problem management.

Virtualization environments are dynamic by nature. Therefore, ideally, you should govern them by using documented processes and procedures that are based on the Information Technology Infrastructure Library or Microsoft Operations

Frameworks. Service Manager can help you govern virtualization or private cloud computing with the following functionality:

- Management of incidents, problems, changes, and releases. Service Manager offers application and infrastructure owners, administrators, service analysts, and end users a single location from which to govern and manage deployment changes and administrate a complex virtualization environment. It provides a SharePoint-based portal that you can customize and configure easily with a software or service catalog that you can link to self-service request offerings. You can configure request offerings to trigger business approval processes and system processes that deliver the request. This provides a level of automation that increases efficiency significantly.
- Management Packs. Management packs, such as the System Center Cloud Services Process Pack, enable you to extend the capability of Service Manager quickly.



Note: Management packs extend System Center 2012 R2 functionality and enable the integration between System Center components. You can download and install a wide variety of management packs for most System Center components.

- An integrated platform. Service Manager has several available connectors to take advantage of Service Manager's full integration capabilities. You can use these connectors to import data into the Server Manager configuration management database from AD DS, comma-separated values (CSV) files, and other System Center components.

These are a few of the available connectors:

- AD DS. AD DS adds information about users, groups, printers, and computers.
- Operations Manager. There are two connectors for Operations Manager. The first imports discovered configuration items. The second is an alert connector that can create incidents based on alerts. If Operations Manager uses the Windows Azure management pack, alerts from Windows Azure will be available also.
- Orchestrator. Orchestrator provides the ability to invoke runbooks. You can fulfill service catalog requests automatically when you use runbooks with the self-service portal.
- Configuration Manager. Configuration Manager allows you to import manageable hardware and software configuration items.
- VMM. VMM imports template objects and storage classifications, which you can use to create service offerings.

With Service Manager, you can:

- Implement service management, as defined in the Information Technology Infrastructure Library and the Microsoft Operations Framework
- Provide processes for:
 - Incident and problem management
 - Change control and lifecycle management
 - Enforcing compliance
 - Managing integrated systems using connectors and solution accelerators

Microsoft also provides a solution accelerator, the Cloud Services Process Pack. This enables you to use predefined forms and automation that utilizes Service Manager, Orchestrator, and VMM to provide a complete IaaS offering with self-service functionality. The pack allows administrators to populate the service catalog with request offerings that they can automate to deliver virtual machine and cloud resources to end users, while following that business approval processes.



For more information on System Center 2012 - Service Manager Parts, go to:

<http://go.microsoft.com/fwlink/?LinkId=253994>

Overview of Orchestrator

Orchestrator is an IT process automation solution that you can use to automate the creation, monitoring, and deployment of key resources in your environment. Data center administrators perform many critical daily tasks to ensure that their infrastructure is highly available and reliable. They also strive to reduce the time it takes to provision new infrastructure while providing self-service capabilities to end users. In addition, the administrators must maintain quality standards and system efficiency.

By using Orchestrator, you can:

- Automate processes in your data center, regardless of hardware or platform.
- Automate your data center operations and standardize best practices to improve operational efficiency.
- Connect different systems from different vendors.

Orchestrator uses runbooks to automate tasks that administrators perform frequently in a data center. You can create runbooks by using the Runbook Designer within Orchestrator. The Runbook Designer is a simple drag-and-drop interface that makes it easy to design processes to accomplish complex tasks.

Some examples of runbooks:

- Creating Active Directory users.
- Deploying virtual machines.
- Adding users to groups.
- Deploying backup or monitoring agents.
- Sending notification emails to indicate status or completion.

Integrating Orchestrator with Other System Center Components

Orchestrator has a number of built-in runbook activities that perform a wide range of functions that you can extend with integration packs. Integration packs contain runbook activities and objects that provide Orchestrator with the ability to extend its capabilities to other Microsoft and third-party components. The integration pack for Service Manager includes activities that enable Orchestrator to obtain details of incidents and problems that the Service Manager environment has generated. This integration provides a useful mechanism in automating tasks with Service Manager. For example, you can create a runbook in Orchestrator that creates a problem record in Service Manager when the number of related incidents reaches a specified number.

Orchestrator provides the ability to:

- Automate processes across systems and platforms
- Automate best practices
- Implement end-to-end automation across multiple System Center products
- Implement out-of-box integration packs

The integration pack for DPM includes activities that enable you to automate tasks within the DPM environment. For example, you can create a runbook that automates the protection of a data source, such as SQL Server, based on a new service that you are provisioning in VMM.

Building runbooks from the activities in the integration packs can ensure that repetitive tasks are performed quickly and with improved accuracy. You can deploy the Windows Azure integration pack for Orchestrator to allow automation of Windows Azure operations related to certificates, deployments, cloud services, storage, and virtual machines.

Furthermore, you can build your own integration packs to build workflows and processes as necessary.

Updated integration packs for System Center 2012 R2 are available in the Microsoft Download Center.

Orchestrator has several new features in System Center 2012 SP1, including:

- New integration packs, including some third-party integration packs.
- Management capabilities for VMM self-service user roles.
- Management capabilities for multiple VMM stamps (scale units). Additionally, you can aggregate results from multiple stamps.
- Integration with App Controller to manage hosted environments.

Microsoft System Center 2012 R2 Orchestrator has added additional features that can be helpful in the data center:

- New integration pack for SharePoint.
- Updated Windows Azure integration pack.
- Updated integration pack for VMM.

Overview of Operations Manager

Operations Manager is a cross-platform monitoring and alerting solution that provides application and infrastructure monitoring. You can integrate Operations Manager with VMM, Service Manager, and Orchestrator to provide automated remediation in response to errors, performance issues, and outages. Operations Manager also provides management packs to monitor other systems, including many third-party hardware and software components.

Operations Manager includes the following features:

- Network monitoring. Operations Manager supports the discovery of network routers and switches. This provides a platform with which you can monitor networks, from the desktop to the servers.
- Application code monitoring. Operations Manager provides detailed monitoring information for applications, including .NET and Java Enterprise Edition applications, and the ability to identify and pinpoint problems with applications.
- End-to-end monitoring. Operations Manager can monitor applications from end to end. This means that it can monitor the application, the operating system that it runs on, the hardware that the operating system relies on, and the network devices that provide access to the application. If you

Operations Manager provides:

- Network monitoring
- Application code monitoring
- End-to-end monitoring
- Dashboards
- Heterogeneous platform monitoring

By integrating Operation Manager and VMM, you can monitor the entire virtualized environment

distribute the application across multiple systems, you can configure Operations Manager to show the topology of the application in a single pane so that operators can see instantly where the problem is in the topology.

- **Dashboards.** Operations Manager offers predefined and easily customizable dashboards for monitoring key statistics, alerts, and issues from a single management console.
- **Heterogeneous platform monitoring.** Operations Manager monitors Windows servers and applications, but it also monitors Linux and UNIX systems for health and performance issues.

Operations Manager uses management packs that contain information about the objects that you monitor. Application vendors usually develop these management packs. For example, Microsoft authors management packs for each version of its operating systems and server application products such as SQL Server and Exchange Server. Additionally, you can import the System Center Management Pack for Windows Azure to monitor the Windows Azure environment.



To search for and to download management packs for System Center, go to:

<http://go.microsoft.com/fwlink/?LinkID=286068>

Integrating Operations Manager and VMM

You can integrate VMM and Operations Manager to provide complete monitoring of both physical host machines and virtual machines. This integration allows you to:

- Monitor the health and availability of virtual machines, hosts, the VMM management server, the VMM database server, and the library servers. In addition, you can monitor a VMware-based virtual environment.
- See the diagram views of your virtualized environment from within the Operations Manager console.
- Implement PRO tips, which collect performance data from host machines, virtual machines, and applications. PRO tips enable you to automate changes to the VMM and host environment, based on the performance information that Operations Manager provides. For example, if a physical hard disk fails, an alert in Operations Manager could trigger the evacuation of the host with a degraded disk subsystem. Another example is using performance information to scale out a web farm automatically, in response to increased transactions in VMM. The reports are available in the Virtual Machine Manager console, but you retrieve display data from Operations Manager.
- Enable maintenance mode integration. When you place hosts in maintenance mode, VMM attempts to put them in maintenance mode in Operations Manager.
- Integrate SQL Server Analysis Services (SSAS), which allows you to run forecasting reports that can predict host activity based on history of disk space, memory, network I/O, disk I/O, and CPU usage. This also supports usage of a SAN for usage forecasting.

System Center 2012 R2 Operations Manager introduces several improvements to data center management:

- **Fabric monitoring.** By integrating with VMM, you can monitor the fabric of your private cloud. The monitoring consists of two new features, the Fabric Health Dashboard and the Fabric Monitoring Diagram View.
- **Support for IPv6.** System Center 2012 R2 Operations Manager can now use IPv6 addresses as input for discovery and display of the IPv6 address information.
- **System Center Advisor.** The System Center Advisor is an online service that monitors the installations of Microsoft server software. The alerts that the System Center Advisor generates are now integrated into the Operations Manager console.

- Java Application Performance Monitoring. You can now monitor exceptions and performance of the Java applications in your environment using the Operations Manager Application Advisor.



For more information, go to How to Connect VMM with Operations Manager:

<http://go.microsoft.com/fwlink/?LinkID=286069>

Overview of Data Protection Manager

DPM is a data backup and recovery solution that works for disk-to-disk and disk-to-tape backups. It enables you to back up and restore Windows servers and application servers such as SQL Server, Exchange Server, Hyper-V, file servers, AD DS, and SharePoint Server. In addition, DPM includes support for system state and bare-metal recovery, offers protection for Windows desktop clients, and provides some elements of self-service.

When planning a virtualization environment, you must implement a backup system that can back up:

DPM provides:

- VSS backups
- Hyper-V item level backup and recovery
- Hyper-V host and guest support
- Integration with Operations Manager
- Integration with other System Center 2012 components
- Self-service functionality
- Cloud-based backups

New features in System Center 2012 R2 DPM include:

- SQL Server cluster support
- Virtualized deployment
- Linux virtual machine backup

- Virtual machines. Virtual machines can provide a challenge to older backup software products, as they may not be virtualization-aware. Additionally, older backup solutions may not be application-aware. You need to consider how to back up both the systems and the applications on the systems. For example, Microsoft Exchange backups should protect Exchange components, such as stores and mailboxes. Additionally, if you want to protect your entire server structure, you should perform a system state backup and include the data drives. If you must recover your whole server, you have to recover from a full backup of all components.
- Host server backup. Not to be confused with backing up the host itself, a host-level backup is a Hyper-V-aware backup designed to protect the virtualization files that comprise a virtual machine, such as the virtual machine configuration files, VHDs, and snapshots. DPM uses VSS to back up files while they run. You can use this form of backup to recover an entire virtual machine or one of its disks in place to the same virtualization host server or to an alternate virtualization host server.

DPM provides some important features as a data center backup system:

- VSS backups. DPM uses VSS to provide protection of data sources while the data source continues to run. This means that you do not have to take applications and servers offline while DPM provides the protection for them. After an initial full backup is complete, DPM can back up individual block changes incrementally, allowing for fast and efficient backup and recovery.
- Hyper-V item level recovery support. DPM can recover specific files, folders, volumes, and VHDs from a host-level backup of Hyper-V virtual machines.
- Hyper-V host and guest support. DPM supports host-based protection when the agent is installed on the host computer and guest-based protection when the agent is installed on the virtual machine. For guests running Windows Server 2003 and newer versions, DPM provides online backups that ensure that DPM does not bring down the protected virtual machine while providing protection.

- Integration with Operations Manager, which provides monitoring of the DPM environment by using the DPM Management Pack. The Data Protection Manager Central Console, which is built on Operations Manager, allows you to monitor all Data Protection Manager servers from a central computer. You can use the central console to open a Data Protection Manager Administrator Console to manage DPM remotely.
- Integration with other System Center 2012 components. With DPM's integration with Orchestrator, you can automate functions such as protection and recovery of data. By using Service Manager and the self-service portal together with DPM and Orchestrator, you can offer these functions as services to the private cloud.
- Self-service functionality. DPM has a self-service function that allows administrators to configure and delegate the restore functionality to user self-service users. You can grant permission to restore to the same server or to an alternate server, which you can choose.
- Cloud-based backups. You can integrate DPM with Windows Azure Backup to allow the management of cloud-based backups through DPM.

System Center 2012 R2 DPM has introduced new features, including:

- SQL Server cluster support. DPM now supports the use of SQL Server clustered server nodes for its database.
- Virtualized deployment. Now you can install DPM on a virtual machine. DPM can use .vhd storage pool disks through the VMM library.
- Linux virtual machine backup. DPM provides a mechanism for backing up Linux virtual machines by using file consistent snapshots.

Additional Tools for Managing the Data Center

Aside from its core components, System Center provides several additional services and tools, including two cloud-based System Center services, Windows Intune and Microsoft System Center Advisor, and several downloadable tools. The following list provides details about these services and tools:

- Windows Intune. Windows Intune is a cloud-based solution for computer management and security. You can use it to deploy software, support customers via remote control, create policies for update management, and ensure that your Endpoint Protection is current. Windows Intune offers many of the features available with Configuration Manager without the need to have an on-premise deployment. The current release of Windows Intune integrates with System Center 2012 SP1 Configuration Manager to allow administrators to manage devices registered with both products from a single console.
- System Center Advisor. System Center Advisor is a cloud-based optimization service that proactively scans your servers to check for known issues with Windows Server, AD DS, SQL Server, Exchange Server, SharePoint Server, and Hyper-V. It uses an agent that can communicate with a single gateway to reduce bandwidth.

System Center provides the following services and tools:

- Windows Intune
- System Center Advisor
- Virtual Machine Servicing Tool 2012
- Plug-in for Microsoft Virtual Machine Converter for VMware vSphere client
- Microsoft Baseline Configuration Analyzer
- VMMCA
- Migration Automation Toolkit

System Center Advisor and Windows Intune are cloud-based services, while the following tools in this list are task-specific tools.

- Microsoft Virtual Machine Converter Plug-In for VMware vSphere Client. Administrators who want to convert VMware virtual machines to Hyper-V-based virtual machines can use the plug-in for the virtual machine converter for the VMware vSphere client. It extends the vSphere client to allow conversions from the virtual machine context menu.
- Microsoft Baseline Configuration Analyzer 2.0. Microsoft Baseline Configuration Analyzer 2.0 can help you maintain optimal system configuration by analyzing configurations of your computers against a predefined set of best practices, and then reporting results. This is a prerequisite installation for the Virtual Machine Manager Configuration Analyzer.
- Virtual Machine Manager Configuration Analyzer (VMMCA). Virtual Machine Manager Configuration Analyzer is a diagnostic tool that you can use to evaluate configuration settings for servers that are running VMM management server roles or are acting as virtual machine hosts. The Virtual Machine Manager Configuration Analyzer scans the specified system's hardware and software configuration, and then evaluates them against a set of predefined rules, reporting on any configurations that are not optimal. To run the Virtual Machine Manager Configuration Analyzer, you must first install the Microsoft Baseline Configuration Analyzer.
- Migration Automation Toolkit. The Migration Automation Toolkit is a collection of Windows PowerShell scripts that automate virtual machine conversions by using the Microsoft Virtual Machine Converter. You can use the toolkit to migrate one server or several servers at the same time.



How to Connect VMM with Operations Manager

<http://go.microsoft.com/fwlink/?LinkID=392380>



Microsoft Baseline Configuration Analyzer 2.0

<http://go.microsoft.com/fwlink/?LinkID=286071>

Lab: Considerations for Implementing an Enterprise Data Center

Scenario

A. Datum Corporation is an engineering and manufacturing company. The organization's base is in London, England, and it has branch offices throughout Europe, Australia, and North America. The company is expanding its network of business partnerships and providing additional services to customers. As the company expands, it is becoming apparent that some of the business requirements are also changing.

Faced with these changes, IT management at A. Datum is launching a project to evaluate the current IT infrastructure. They want to identify the components within the infrastructure that require upgrades or enhancements to provide the required services. IT management has identified the following requirements as the core goals for the server infrastructure upgrade at A. Datum:

- Whenever possible, IT management wants to deploy all servers, services, and applications in a virtual machine environment. The management of the virtual environment must be optimized for virtualization, monitoring, and automation.
- All services and applications must be available during regular business hours. Business-critical services and applications must be available 24 hours per day, seven days per week.
- The upgrade must address the security department requirements.
- Several business groups have identified additional business requirements regarding external-facing websites and applications that key business partners and customers use. The upgrade must address these requirements also.

Objectives

After completing this lab, you will be able to plan for the implementation of services within an enterprise data center.

Estimated Time: 30 minutes

No virtual machines are required for this lab.

Exercise 1: Planning the Secure Implementation of Services Within an Enterprise Data Center

The main tasks for this exercise are as follows:

- Review the IT and business requirements
- Identify how to address the IT and business requirements
- Discuss your proposed solution with the class, as guided by your instructor

► Task 1: Review the IT and business requirements

Read the lab scenario

► **Task 2: Identify how to address the IT and business requirements**

Answer the following questions:

1. How will you meet the requirement to provide a single tool for managing virtualization hosts and virtual machines?
2. How will you meet the requirement for business unit administrators to manage their applications and virtual machines?
3. How will you meet the requirement to provide detailed information about the performance of all data center components?
4. How will you meet the requirement to automate processes within the data center?
5. How will you meet the requirement to provide high availability for the required applications?
6. How will you meet the security department requirements?
7. How will you meet the requirement to improve the performance of the sales website?
8. How will you meet the requirement to enable integration with Trey Research?
9. How will you meet the requirement to provide the required access to the partner website?

► **Task 3: Discuss your proposed solution with the class, as guided by your instructor**

Be prepared to discuss your answers with the class

Results: After completing this exercise, you will have identified the components that you will need to include in the data center design.

Question: How do the requirements at A. Datum compare to your organization's requirements? What requirements are similar? What additional requirements do you have?

Question: What services and tools are you using to manage your data center? How well integrated are the tools?

Module Review and Takeaways

Review Questions

Question: Explain how System Center components are integrated, and list the benefits of integration.

Question: Compared to physical machines, what additional high availability option can you use when you deploy virtual machines?

Module 2

Planning and Implementing a Server Virtualization Strategy

Contents:

Module Overview	2-1
Lesson 1: Planning a VMM Deployment	2-2
Lesson 2: Planning and Implementing a Server Virtualization Host Environment	2-10
Lab: Planning and Implementing a Server Virtualization Strategy	2-19
Module Review and Takeaways	2-26

Module Overview

This module introduces the Microsoft® System Center 2012 R2 Virtual Machine Manager (VMM) components. You will see how they integrate to enable you to configure, deploy, and manage a server virtualization environment. Later, you will review the planning steps and consideration for a System Center 2012 R2 Virtual Machine Manager deployment.

Objectives

After this module, you will be able to:

- Determine an appropriate topology for a VMM deployment.
- Plan the deployment of the VMM database, management server, and libraries.
- Deploy and add virtualization hosts.
- Create and manage host groups.
- Manage VMM libraries.

Lesson 1

Planning a VMM Deployment

It is possible to deploy all VMM components on a single server. However, in medium to large enterprises, you might spread the components that comprise a VMM deployment across multiple servers. In this lesson, you will learn about the properties of key VMM components and the factors that you should consider when choosing how to deploy these components in your environment.


Lesson Objectives

After completing this lesson, you will be able to:

- Determine an appropriate topology for a VMM deployment.
- Plan the deployment of the VMM database.
- Plan the deployment of the VMM management server.
- Plan and configure a VMM library.

Determining the Topology of a VMM Deployment

A VMM deployment consists of the following components.



Component	Description
VMM management server	<p>The computer that hosts the VMM service. This computer processes commands and controls communications with the following components:</p> <ul style="list-style-type: none"> • VMM database • Library Server • Virtualization hosts <p>Virtualization hosts include Microsoft Hyper-V®, VMware ESX & ESXi, and Citrix XenServer.</p>
VMM Database	A Microsoft SQL Server® database that stores VMM configuration information.
VMM console	An application that connects to the VMM management server. It allows an administrator to centrally view and manage physical and virtual resources.

Component	Description
VMM Library	Collection of resources used to deploy virtual machines and services.
VMM Library Server	A computer that hosts shared folders that store the file-based resources included in a VMM library.
VMM command shell	A Windows PowerShell®-based command shell that allows you to perform administrative functions through the VMM management server.

You can deploy these components on a single computer or distribute them across multiple computers.

When you design a VMM deployment, you must consider the following key factors:

- Number of virtualization hosts that VMM will manage.
- Network topology, including branch sites with virtualization hosts.
- Administrative boundaries.
- Self-service options.
- Availability and recovery goals.
- Availability and recovery time that each component requires.

The number of virtualization hosts determines the physical or virtual resources that each component server in the VMM deployment requires. For example, you will have to decide whether to have a single VMM deployment with multiple VMM library servers or individual VMM deployments at each branch. You will base this decision on the number of branch sites with hosts and the capacity of the wide area network (WAN) links between branch sites and the VMM management server.

You can use System Center 2012 R2 Service Manager to provide self-service virtual machine deployment. When you determine the best type of VMM deployment for your environment, you can plan a self-service deployment that is appropriate for the design. For example, you can use System Center 2012 R2 App Controller to provide self-service virtual machine deployment across five VMM deployments.

The availability and recovery time for components is also important when determining the topology of VMM deployment. VMM is a cluster-aware application. You can install a library server on a clustered file server but not on the same failover cluster that hosts a clustered VMM instance. You can use System Center 2012 R2 Data Protection Manager to back up and restore VMM components.

Additional VMM deployment considerations include:

- Deploy a server hosting the Windows Deployment Services (Windows DS) role if you want to support bare-metal deployment of Hyper-V hosts. A *bare-metal deployment* refers to deploying a host on a computer that does not have an operating system.
- Your deployment requires at least one library server. Typically, you deploy at least one library server to each site that hosts virtualization hosts managed by VMM.
- Use Windows Server Update Services (WSUS) or System Center 2012 R2 Configuration Manager to manage software updates for VMM, virtualization hosts, and virtual machines.
- Use System Center 2012 R2 Operations Manager to enable VMM reporting and to leverage Performance and Resource Optimization (PRO) tips
- Managing ESX and ESXi hosts requires that you integrate VMware vSphere.

In multisite deployment, it may be necessary to configure firewall ports to support your VMM deployment.

The following table lists some default ports that VMM uses.

Port	Description
8100	Provides communication with the VMM console
5985	Provides communication with agents on hosts and library servers
443	Enables file transfers to agents on hosts and library servers
8102	Provides communication with Windows DS
8101	Provides communication with Windows Preinstallation Environment (Windows PE) agents
8103	Provides communication with the Windows PE agent for time synchronization

You can deploy System Center 2012 R2 App Controller to manage multiple VMM servers. You can also use App Controller as a way to manage your on premise virtualization infrastructure as well as any Windows Azure™ subscriptions within your organization.

Planning the VMM Database

The VMM database stores VMM configuration information. When planning the VMM database, you need to consider the version of SQL Server that you will use and the resources that you will allocate to the host. You can deploy the VMM database on a physical or virtual server.

Minimum and recommended configurations for managing up to 150 hosts are as follows.

Up to 150 Hosts		More than 150 Hosts	
Resource	Recommended	Resource	Recommended
CPU	Pentium 4, 2.8 GHz	CPU	Dual-Core, 2GHz
RAM	4 GB	RAM	8 GB
Database Disk Space	50 GB	Database disk space	75 GB

Supported SQL Server configurations:

SQL Server version	Service pack	Editions
SQL Server 2008 R2	Service Pack 2 or newer	Standard, Enterprise, & Datacenter
SQL Server 2012	Service Pack 1 or newer	Standard, Enterprise

- For high availability deploy VMM databases on clustered SQL Server instances or by using AlwaysOn Availability Groups

Hardware component	Minimum	Recommended
Processor	Pentium 4, 2.8 gigahertz (GHz)	Dual-Core 64-bit, 2 GHz
RAM	2 gigabytes (GB)	4 GB
Database disk space	20 GB	50 GB

Minimum and recommended configurations for managing more than 150 hosts are as follows.

Hardware component	Minimum	Recommended
Processor	Dual-Core 64-bit, 2 GHz	Dual-Core 64-bit, 2.8 GHz
RAM	4 GB	8 GB
Database disk space	50 GB	75 GB

When deploying on a virtual server, you will need to allocate the same resources to the virtual machine that you would allocate to the physical server.

You can use the following versions and editions of SQL Server to host the System Center 2012 R2 VMM database.

SQL Server version	Service pack	Editions
SQL Server 2008 R2	Service Pack 2 (SP2) or newer	Standard, Enterprise, and Datacenter
SQL Server 2012	Service Pack 1 (SP1) or newer	Enterprise, Standard

The System Center 2012 and System Center 2012 R2 suites include license rights for the Standard edition of SQL Server.

The VMM database must be in the same Active Directory® Domain Services forest as the VMM management server. Alternatively, there must be a two-way trust relationship between the domain that hosts the VMM management server and the domain that hosts the database. The SQL database server name may not be longer than 15 characters and is not case-sensitive.

You can make the VMM database highly available by:

- Deploying it on a clustered SQL Server instance.
- Placing it on a highly available virtual machine.
- Deploying it on a clustered SQL Server instance deployed on highly available virtual machines.

If you are planning to use a highly available SQL Server instance, you should also plan to have a highly available VMM management server. The VMM database supports the AlwaysOn Availability Groups feature of SQL Server 2012.

Planning a VMM Management Server

The VMM management server runs the VMM service, which processes all commands and handles all communication between the VMM database, the library servers, and the virtual machine hosts. It is necessary that the VMM management server be located on the same server as the server that hosts the VMM database, or have a fast network connection to that server. When you use the VMM console or the VMM command shell, a connection is made to the VMM management server.

In organizations with good connectivity between sites with virtualization hosts, a single VMM management server is probably appropriate. Deploy multiple VMM management servers and VMM instances if you have:

- Large numbers of virtual machine hosts at branch offices or regional sites.
- Administration teams at each site with large numbers of virtual machine hosts.
- Poor inter-site bandwidth.

The VMM management server is cluster-aware and you can deploy it as highly available on a failover cluster. You can also deploy the VMM management server on a highly available virtual machine, which is an easier configuration to deploy. It does not require configuring a failover cluster, either on a physical server, which is more expensive, or where multiple virtual machines participate in a guest failover cluster, which requires configuration of shared storage.

The VMM management server hosts the VMM service, processes commands, and controls communications with the VMM database, library server, and virtualization hosts

- Deploy the VMM management server centrally in organizations with good bandwidth
- Deploy multiple VMM management servers where multiple teams manage large number of virtualization hosts at each site

Hardware requirements for a VMM management server:

Up to 150 hosts		Over 150 hosts	
Resource	Recommended	Resource	Recommended
CPU	Pentium 4, 2.8 GHz or greater	CPU	Dual-Core, 3.2 GHz or greater
RAM	4 GB	RAM	5 GB
Disk space	40 GB	Disk space	50 GB



Note: When you are naming the management server, the computer name cannot contain the character string SCVMM. For example, you cannot name the server ADATUM-SCVMM-01, but you can name it ADATUMSCVMMM01.

System Requirements for a VMM Management Server

The following table describes the hardware requirements for managing up to 150 hosts.

Hardware component	Minimum	Recommended
Processor	Pentium 4, 2 GHz (x64)	Dual-processor, 2.8 GHz (x64) or greater
Random access memory (RAM)	2 GB	4 GB
Hard disk space (without a local VMM database)	2 GB	40 GB
Hard disk space (with a local, full version of SQL Server)	80 GB	150 GB

The following table describes the hardware requirements for managing more than 150 hosts.

Hardware component	Minimum	Recommended
Processor	Pentium 4, 2 GHz (x64)	Dual-processor, 3.6 GHz (x64) or greater
RAM	4 GB	8 GB
Hard disk space	10 GB	50 GB

If you are managing more than 150 hosts, you can enhance performance by separating the VMM components. For example, you can deploy a separate library server instead of using the default library share on the same server as the VMM management server. Conversely, you can use a VMM database on a dedicated computer running SQL Server.

The following table describes the software requirements for installing the VMM management server.

Software requirement	Notes
A supported operating system	Windows Server® 2012 or Windows Server 2012 R2 (full installation), Standard or Datacenter
Windows Remote Management (WinRM) Service	Windows Server 2012 and Windows Server 2012 R2 includes WinRM. By default, the Windows Remote Management (WS-Management) service is set to start automatically. If the Windows Remote Management (WS-Management) service has not started, setup will display an error during the prerequisites check. You must start the service before setup can continue.
At least Microsoft .NET Framework 4	System Center 2012 SP1 and newer require .NET framework 4, which Windows Server 2012 and Windows Server 2012 R2 include.

Software requirement	Notes
Windows Assessment and Deployment Kit (ADK) for Windows® 8	Windows ADK is available at the Microsoft Download Center at Windows Assessment and Deployment Kit (ADK) for Windows 8 at: http://www.microsoft.com/en-us/download/details.aspx?id=30652 . When you install Windows ADK, select the Deployment Tools and the Windows Preinstallation Environment features.

System Requirements for VMM Consoles

The following table describes the hardware requirements for managing up to 150 hosts.

Hardware component	Minimum	Recommended
Processor	Pentium 4, 550 megahertz (MHz)	Pentium 4, 1 GHz or greater
RAM	512 MB	1 GB
Hard disk space	512 MB	2 GB

The following table describes the hardware requirements for managing more than 150 hosts.

Hardware component	Minimum	Recommended
Processor	Pentium 4, 1 GHz	Pentium 4, 2 GHz or greater
RAM	1 GB	2 GB
Hard disk space	512 MB	4 GB

The following table describes the software requirements for installing the Virtual Machine Manager console.

Software requirement	Notes
A supported operating system	Listed in the next table.
Windows PowerShell 2.0 or Windows PowerShell 3.0	Windows Server 2008 R2 and Windows 7 include Windows PowerShell 2.0. Windows Server 2012 includes Windows PowerShell 3.0. Windows Server 2012 R2 includes Windows PowerShell 4.0.
At least Microsoft .NET framework 4	On a computer that is running Windows 7, .NET framework version 3.5.1 is installed by default. On a computer that is running Windows Server 2008 R2, the .NET Framework 3.5.1 feature is not installed by default. However, you can use the VMM Setup Wizard to install the feature. On a computer that is running Windows 8, Windows 8.1, Windows Server 2012, or Windows Server 2012 R2, .NET framework 4 is included. .NET Framework 4.5 is available at the Microsoft Visual Studio® 2012 download page at http://go.microsoft.com/fwlink/p/?linkId=285269 .

The following table lists the supported operating systems on which you can install the Virtual Machine Manager console.

Operating system	Edition	System architecture
Windows Server 2008 R2 SP1 (full installation)	Standard, Enterprise, and Datacenter	x64
Windows 7 SP1	Professional, Enterprise, and Ultimate	x86 and x64
Windows Server 2012/Windows Server 2012 R2	Standard and Datacenter	x64
Windows 8/Windows 8.1 Client	Standard, Pro, and Enterprise	x86 and x64

You can deploy the Virtual Machine Manager console on the same server as the VMM management server, or on another server or workstation that is running a supported operating system.

System Center R2 Virtual Machine Manager does not include the optional VMM Self-Service portal. App Controller is now the web component for self-service. A later module discusses App Controller planning and deployment.

Planning a VMM Library

A VMM library server hosts a VMM library. A VMM library includes virtual hard disks, virtual machine templates, and profiles. VMM library servers are deployed on networks that have fast connectivity to virtual machine hosts that VMM manages. This means that you should deploy VMM library servers at branch offices that host virtual machine hosts. Another reason to deploy VMM library servers at branch offices is that proximity to a library server speeds virtual machine deployment.

VMM libraries can contain:

- Virtual machine templates
 - ISO images
 - Hardware profiles
 - Stored virtual machines
 - Cloud libraries
 - Virtual hard disks
 - Scripts
 - Guest operating system profiles
 - Update baselines
 - Equivalent objects
- VMM library servers should be deployed with good network bandwidth to host machines that use the library objects
 - Consider deploying library servers on deduplicated volumes

When planning the hardware for a library server, you need to consider the number of items that require storage. In particular, you should plan for large items such as virtual machine templates, virtual hard disks, and International Organization for Standardization (ISO) images, which use more space than scripts and profiles.

For example, profiles require little space. However, virtual machine templates and ISO images will consume a considerable amount of space, so you will need to plan for potential growth.

You may want to have an image for Windows Server 2008 R2, an image for Windows 2012, and an image for Windows Server 2012 R2. The ISO files can be approximately 4 GB each, while the virtual machine templates may be larger than 15 GB each.

Consider the performance required to deliver these files efficiently. The larger your deployment size, the more likely it is that you will transfer multiple large files to and from the library simultaneously. Therefore, you may need to monitor library performance to ensure that the library is not overwhelmed by requests or have multiple libraries to balance the load and address demand.

When you plan the VMM library servers, consider deploying them on Windows Server 2012 or Windows Server 2012 R2 volumes configured with deduplication. Often, this dramatically reduces the amount of storage space required to host the library, as many libraries contain duplicate content.

System Requirements for VMM Library Servers

The following table lists the minimum and recommended hardware requirements for installing a VMM library server.

Hardware component	Minimum	Recommended
Processor	Pentium 4, 2.8 GHz	Dual-core, 3.2 GHz or greater
RAM	2 GB	2 GB
Hard disk space	Varies based on the size and number of files stored	Varies based on the number and size of the files stored

Software Requirements for VMM Library Servers

The following table lists the software requirements for installing and running the VMM library server.

Software requirement	Notes
WinRM	Windows Server 2008 R2, includes WinRM 2.0. By default, the WinRM service is set to start automatically (delayed start). Windows Server 2012 and Windows Server 2012 R2 include WinRM. By default, the Windows Remote Management service is set to start automatically.

The following table lists the supported operating systems on which you can install the VMM library.

Operating system	Edition
Windows Server 2008 R2 SP1	Standard, Enterprise, and Datacenter
Windows Server 2012 (full installation or Server Core installation)	Standard and Datacenter
Windows Server 2012 R2 (full installation or Server Core installation)	Standard and Datacenter

Lesson 2

Planning and Implementing a Server Virtualization Host Environment

Once you deploy VMM, you need to understand how you can use it to manage your organization's existing virtualization host infrastructure and deploy new virtualization hosts. In this lesson, you will learn how to add and deploy virtualization hosts, configure host groups, and manage VMM libraries.

Lesson Objectives

After completing this lesson you will be able to:

- Use the VMM management console.
- Add virtualization hosts to VMM.
- Deploy Hyper-V hosts.
- Manage and configure host groups.
- Manage VMM libraries.

Adding Virtualization Hosts to VMM

Before you can use VMM to manage a Microsoft Hyper-V virtualization host, you must deploy the VMM agent software to that host. You can do this by using the Add hosts function in the VMM console. When you add a host, you provide the credentials of an account that has local administrator privileges on the virtualization host that you add.

For a host in a perimeter network, you will need to deploy the agent software manually on the target computer, and then add the host in the VMM console.


To deploy a Hyper-V host in a trusted domain, follow this procedure:

1. Open the VMM console, click the **VMs and Services** workspace, from the ribbon, click **Add Resources**, and then click **Hyper-V Hosts and Clusters**.
2. On the **Resource location** page, click **Windows Server computers in a trusted Active Directory domain**, and then click **Next**.
3. On the **Credentials** page, choose to either use a RunAs account (an account already configured with domain privileges) or manually enter the credentials of an account with privileges to install the agent on the host server, and then click **Next**.
4. On the **Discovery Scope** page, you can either specify computer names by entering them on separate lines in the **Computer name** field or click **Specify an Active Directory query to search for Windows Server computers**, type a query, and then click **Next**.

Add a virtualization host in a trusted domain from the VMM console
Ensure that you have an account with local administrator rights on the virtualization host
Manually install agent on virtualization hosts on perimeter network
Heterogeneous host support requirements:
VMWare vCenter Server 4.1 required for ESX
Citrix Xen Server requires the Citrix System Center integration pack


5. On the **Target resources** page, you can click each host or click **Select all**, and then click **Next**. A dialog box will inform you that you are about to enable the Hyper-V role on all servers as part of the process. If you choose to enable the role, the servers will reboot during the process. You can click **OK** to close the dialog box.
6. On the **Host settings** page, you can assign the host or hosts to a host group. A later section of this module details host groups. Additionally, if you have multiple VMM management servers, and another VMM environment is managing your host currently, you can reassociate the host with this environment by clicking **Reassociate**. You can also assign default placement paths, which are where Windows will store new or migrated Hyper-V virtual machine files. You can do this now or after you add the host. Click **Next**.
7. On the **Summary page**, confirm the settings, and then click **OK**. You can review the progress of the agent deployments in the Jobs window.

When you add a host in a perimeter network by deploying the VMM agent software manually, you must generate an encryption key file and assign a password. Later, this is used to connect the VMM management server to the agent on the host on the perimeter network.

 **Note:** By default, the VMM management server uses port 5986 for agent communication with hosts in a perimeter network and port 443 for file transfers.

You can use VMM to manage VMware ESX and ESXi hosts and Citrix XenServer hosts. Adding other vendor hosts is similar to adding Hyper-V hosts, as long as your environment meets the prerequisites for adding each type. Before you can add a VMware host to VMM, you must configure a vCenter Server and configure VMM to connect to it. Before you can add a Citrix XenServer host, you must add the Citrix System Center integration pack to the host.

Best Practice: You should consider security requirements before adding other vendor hosts to your network. For example, you must decide how to implement certificates for virtualization hosts and you may want to determine how to use a Run As account.

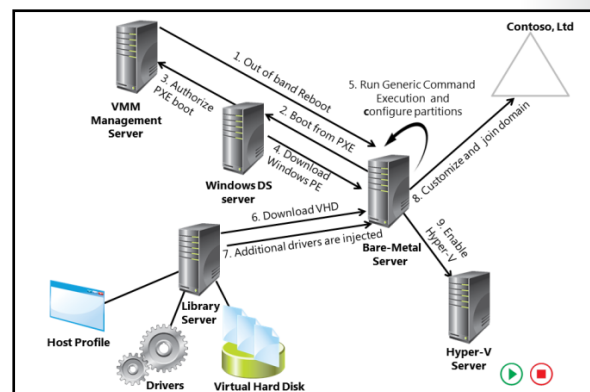
 **Reference Links:** For more information, go to:

- System Requirements: VMware ESX Hosts <http://go.microsoft.com/fwlink/p/?linkId=285337>
- System Requirements: Citrix XenServer Hosts <http://go.microsoft.com/fwlink/p/?linkId=285261>

Deploying Hyper-V Hosts

In a bare-metal deployment, you can use VMM and a Windows Deployment Services server to deploy Windows Server 2008 R2, Windows Server 2012, or Windows Server 2012 R2 host machines to servers that are not running any operating systems.

This process uses Windows DS to deploy a virtual hard disk (.vhd or .vhdx) image hosting the virtualization host operating system to the target machine. When you deploy the operating system, VMM will enable the Hyper-V role and add the server to the chosen host group.



Virtual Machine Hosts

The following table lists the prerequisites for bare-metal deployment.

Software	Notes
A computer that is running Windows Server 2008 R2, Windows Server 2012, or Windows Server 2012 R2, with the Windows Deployment Service role installed	This will serve as the Pre-Boot EXecution Environment (PXE) server, which initiates the operating system installation on the physical computer. Only the Microsoft PXE server is supported.
A physical computer with a baseboard management controller that supports out of band management protocol	Intelligent Platform Management Interface versions 1.5 or 2.0. Data Center Management Interface version 1.0. System Management Architecture for Server Hardware version 1.0 over WS-Management.
A Windows Server 2008 R2, Windows Server 2012, or Windows Server 2012 R2 system image	The operating system image must be Windows Server 2008 R2 or a newer version, because these support booting from virtual hard disks. You can create the .vhd or .vhdx image by using Sysprep.exe.

Follow these high-level steps to prepare for a bare-metal deployment:

1. Ensure that you enable the basic input/output system (BIOS) settings to support virtualization.
2. Set the BIOS boot order to boot from a PXE-enabled network adapter.
3. Add driver files to the library.
4. Make a suitable .vpx or .vhdx image available in the library.
5. Create a host profile.
6. Ensure that the VMM management server can communicate with the baseboard controller, and then run discovery and deploy.

A bare-metal deployment of a new Hyper-V virtualization host follows this general process:

1. VMM issues out of band reboot to the host server.
2. The host server boots from Windows Deployment Services by using boot.wim.
3. The Windows Deployment Services server confirms whether the host is authorized to perform a PXE boot.
4. The host downloads the VMM Windows PE image.
5. The host runs Generic Command Execution and then configures partitions as set in the host profile.
6. The host downloads the .vhd or .vhdx file from the VMM Library server.
7. Any additional drivers are injected.
8. The host reboots. Then you can customize it and join it to a domain.
9. VMM deploys the agent.

Best Practice: Some considerations for bare-metal deployment:

1. You should ensure that your baseboard firmware is current. Consider updating all firmware before deployment.
 - Make sure that you have enough space for the .vhd file on the physical server partition, because VMM will cache drivers during deployment.
 - When creating the .vhd or .vhdx file, consider the size of the host page file. The host's RAM will determine this. After you deploy the server, remote administration will not be enabled. Consider creating Group Policy settings that enable remote administration.

Demonstration: Using the VMM Management Console

In this demonstration, you will see how to:

- Add a Hyper-V host.
- Create a host group.
- Create a cloud.
- Review the VMM management console workspaces.

Demonstration Steps

1. On LON-VMM1, open the Virtual Machine Manager console.
2. Use the Add Resource Wizard in the Fabric workspace to add a new Hyper-V host by using the following information:
3. Resource location: **Windows server computers in a trusted Active Directory domain**
 - Credentials: **Adatum\Administrator** with the password **Pa\$\$w0rd**
 - Computer name: **LON-HOST1**
 - Host group: **All Hosts**
 - Review the status, and then close the Jobs window when the task finishes. A warning message may appear, related to multipath I/O not being enabled. You can ignore this message for this task.
4. In the **VMs and Services** workspace, create a host group named **Sydney Host Group**.
5. Create a new Cloud called **Sydney Cloud** that includes the Sydney Host Group and the Hyper-V capability profile.
6. Close the Jobs window.
7. Review the **Library**, **Jobs**, and **Settings** workspaces, taking note of the items on the ribbon and any available information, such as list of recent jobs and Run As accounts.

Managing Host Groups

Host groups enable you to organize and manage your virtualization hosts. Virtualization hosts are allocated to the default host group (All Hosts) unless you specify another location. You can nest host groups when you want settings that apply to a parent host group to apply to members of child host groups. You create host groups by in the VMs and Services workspace of the VMM console.

Host groups allow you to configure the following:

- General settings
- Placement rules
- Host reserves
- Dynamic optimization
- Network
- Storage
- Custom properties

Host groups allow the grouping of the host servers and assignment of settings at a group level

With host groups, you can create and configure the following:

- Placement rules
- Host reserves
- Dynamic optimization
- Power optimization prerequisites
- Network settings
- Storage settings
- Custom properties

General Settings

You can use a host group's general settings to configure the following:

- Group name
- Group location
- Group description
- Whether to allow unencrypted file transfers to the group

Placement Rules

Placement rules determine which virtualization host will host a new virtual machine. By default, a host group will use the placement setting from the parent host group. If you wish to configure custom placement rules at the individual group level, you can block inheritance by modifying the parent host group setting.

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

- Virtual machine must match the host
- Virtual machine should match the host
- Virtual machine must not match the host
- Virtual machine should not match host

Host Reserves

Host reserves are placement settings that allow the host system to retain resources for its own use. This is useful when a Hyper-V host has additional services running, such as in a branch office where you have configured a VMM library.

The following table details how you can set or override host reserves at the individual host level.

Resource	Notes
Central Processing Unit (CPU)	You can set this as a percentage. The default is 10. However, 10% of one dual-core processor that is running at 2 GHz is not the same as 10% of four six-core processors that are running at 2.8 GHz.
Memory	The default is 256 MB, but you can change this or set this as a percentage.
Disk I/O	The default is zero, but you can set this as a percentage. You may wish to ensure a minimal amount is reserved if you are using a host as a VMM library.
Disk Space	You can set this as a numeric value or percentage.
Network I/O	The default is zero, but you can set this as a percentage. You may wish to ensure that a minimal amount is reserved if you are using a host as a VMM library.

Dynamic Optimization

Dynamic optimization allows VMM to balance the virtual machine loads automatically within a host cluster. By defining minimum resource thresholds for hosts, VMM will migrate virtual machines to alternative hosts if available resources fall below those assigned thresholds.

The following table shows the thresholds that can be set. These settings will affect all hosts within the host group.

Resource	Notes
CPU	Default is 30%.
Memory	Default is 512 MB.
Disk I/O	Default is zero.
Disk Space	Set as a numeric value or percentage.
Network I/O	Default is zero.

In addition to workload balancing, VMM can also invoke power optimization. You can enable this by selecting Settings under the Power Optimization section of the Dynamic Optimization page.

Power Optimization Prerequisites

To enable power management, you must have a baseboard management controller that supports one of these out of band management protocols:

- Intelligent Platform Management Interface versions 1.5 or 2.0
- Data Center Management Interface version 1.0
- System Management Architecture for Server Hardware version 1.0 over WS-Management

The right corner of the VMM overview page displays the number of computer hours saved in the last 30 days. If you integrate System Center 2012 R2 Operations Manager with VMM, then power optimization reports are available.

Demonstration: Configuring Host Groups

In this demonstration, you will see how to create host groups. Then you will discuss the options that you can configure and assign to them.

Demonstration Steps

- On LON-VMM1, on the desktop, double-click **Virtual Machine Manager Console**.
- In the **Connect to Server** dialog box, ensure that the **Use current Microsoft Windows session identity** check box is selected, and then click **Connect**. The Virtual Machine Manager console opens.
- Click **VMs and Services**, and then in the navigation pane, click **All Hosts**.
- On the ribbon, click **Create Host Group**.
- Type **Classroom** for the host group name.
- Right-click the host group that you created, and then click **Properties**.
- Review the options and settings on each of the pages.

Considerations for Implementing Host Groups

Host groups allow you to collect virtual machine hosts, so that you can organize those hosts and the virtual machines that run on them in ways that are suitable for your organization. Each organization may use host groups in different ways. The most common uses for host groups include:

1. Organizing large numbers of hosts and virtual machines. The simplest way to use host groups is to organize hosts and virtual machines, by creating host groups for each organizational location, for example.
2. Reserving resources. You can use host reserves to specify the CPU, memory, network capacity, disk space, and disk I/O capacity reserved for the host operating system. By creating separate host groups, you can reserve different amounts of each resource on collections of hosts.
3. Specifying self-service hosts. You can specify the virtual machine hosts where users and groups with self-service privileges can create and manage virtual machines. Using host groups allows you to restrict self-service virtual machines to specific virtual machine hosts.
4. Implementing Performance and Resource Optimization (PRO). You implement PRO on host groups and host clusters. PRO uses System Center 2012 R2 Operations Manager to optimize virtual machine placement based on resource utilization through PRO tips.

You can use host groups to:

- Organize and manage large numbers of hosts and virtual machines
- Reserve resources for hosts
- Specify which users can create and operate their own virtual machines
- Implement PRO

Host groups are hierarchical, meaning that, in most cases, a child host group will inherit settings configured in a parent host group. The following table describes host group inheritance.

Host group action	Host reserves settings	PRO settings
Create a new child host group	Child host group will inherit host reserve settings from parent group.	Child host group will inherit PRO settings from parent group.

Host group action	Host reserves settings	PRO settings
Move child host group to a parent host group	Child host group will not inherit reserve settings from parent host group.	Child host group will inherit PRO settings if the Inherit PRO settings from parent Host Group setting is enabled.
Modify parent host group settings	You have the option to apply changes only to parent group or to cascade changes to all child host groups.	Child host group will inherit PRO settings if the Inherit PRO settings from parent Host Group setting is enabled.

Working with VMM Libraries

Your VMM deployment can include one or more library servers, which can be any Windows server that meets the prerequisites. When you add a library server to VMM, VMM will deploy an agent to the server. Then you must configure a file share to store the content, and add a library share in VMM.

When you add files to the VMM library share, VMM indexes the files. This speeds up the process of searching for files. By default, the index refreshes every hour. You can disable the refresh or change the interval period (in hours). Refresh settings are global to VMM and apply to all libraries.

The following table lists the files that VMM will index.

When configuring VMM libraries, you can:

- Add a library server or library share
- Associate a Library server with a host group
- Add file-based resources
- Create or modify equivalent objects
- View and remove orphaned resources

Resource	File name extension
VHDs	.vhd .vhdx .vmdk
ISO images	.iso
Windows PowerShell Scripts	.ps1
SQL Server scripts	.sql
Microsoft Web Deployment Tool packages	.zip
SQL Server data-tier applications (DACs)	.dacpac
Microsoft Application Virtualization (App-V) packages	.osd
Driver files	.inf
Answer files	.inf .xml
Custom resources	Folders with .CR extension
Virtual floppy disks	.vfd .flp

MOCT USE ONLY STUDENT USE PROHIBITED



Additional Reading: For more information, go to Configuring the Library Overview at <http://go.microsoft.com/fwlink/p/?linkId=285262>

You can mark file-based library resources in the VMM library as equivalent objects. For example, you may have a Windows Server 2012 R2 .vhdx file stored in the London, Sydney, and Toronto sites. You can mark this .vhdx file as an equivalent object. When you create a template for a new virtual machine, you can specify the .vhdx file marked as an equivalent object. Then, when a deployment occurs at one of these sites, you can use the equivalent object rather than dragging the original .vhdx file from the location in which the template was created. Equivalent objects allow a single template to be used across multiple sites.

Servers that host library servers and library shares servers must meet the operating system requirements and must be in the same domain as the VMM management server or in a domain with a two-way trust. You must enable file and print sharing on the library servers. All hosts that are using the library server must be able to use Server Message Block (SMB) to access the library server.



Note: When you install VMM, a default library share is created. By default, the library share is named MSSCVMMLibrary and is placed on the VMM management server. However, if you are installing a highly available or clustered VMM management server, you cannot place the library on the cluster on which you are installing VMM.

Demonstration: Managing VMM Libraries

In this demonstration, you will see how to:

- Add a library server or library share.
- Associate a library server with a host group.

Demonstration Steps

Configure a library and library shares

- On TOR-SVR1, on the Server Manager Dashboard, click **File and Storage Services**, and then click **Shares**.
- In the Shares workspace, click **Tasks**, click **New Share**, click **SMB Share – Quick**, click **Next** twice, type **TORVMMLibrary**, click **Next** three times, and then click **Create**.
- Click **Close**.
- On LON-VMM1, open the Virtual Machine Manager console, and then click **Library** in the bottom left of the screen. Add **TOR-SVR1\TORVMMLibrary** as a new library server by using the following credentials:
 - User name: **Adatum\Administrator**
 - Password: **Pa\$\$w0rd**
- Review the job status, and then close the Jobs window.

Question: Why would you want to create an equivalent object?

Question: What are five types of objects that a library might store?

Question: How could you prevent two web servers from running on the same host server?

Lab: Planning and Implementing a Server Virtualization Strategy

Scenario

The information technology (IT) infrastructure at A. Datum Corporation is expanding rapidly, but the company would like to minimize expansion costs. The London data center has almost reached its maximum capacity for space and power, and the company wants to avoid the cost of expanding or building a new data center.

Due to these constraints, one of the key goals in the Windows Server 2012 deployment project is to virtualize as many servers, services, and applications as possible. A critical component of this virtualization strategy is to optimize the virtual environment's management.

The first step in deploying the virtual environment at A. Datum is to plan and implement the host and management layers for virtualization. You are responsible for planning the host configuration for the Hyper-V server deployment and for planning host components, such as Virtual Machine Manager libraries. After completing the design, you will configure the required host layer components in VMM.

Objectives

In this lab, you will learn how to:

- Plan the Hyper-V host deployment.
- Configure Hyper-V hosts and host groups in VMM.
- Configure the VMM library.

Lab Setup

Estimated Time: 60 minutes

Virtual machines: 20414C-LON-HOST1, 20414C-LON-HOST2, 20414C-LON-DC1, 20414C-LON-VMM1, 20414C-TOR-SVR1

User name: **Adatum\Administrator**

Password: **Pa\$\$w0rd**

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

1. Sign in to LON-HOST1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. On the host computer, start **Hyper-V Manager**.
3. In Hyper-V Manager, click **20414C-LON-DC1**, and then in the Actions pane, click **Start**. Wait until the machine is fully started.
4. Repeat step 3 for **20414C-LON-VMM1** and **20414C-TOR-SVR1**.
5. Click **20414C-TOR-SVR1**, and then in the Actions pane, click **Connect**. Wait until the virtual machine starts.
6. Sign in by using the following credentials:
7. User name: **Administrator**
8. Password: **Pa\$\$w0rd**
9. Domain: **Adatum**
10. Repeat steps 5 and 6 for **20414C-LON-VMM1**.
11. Sign in to LON-HOST2 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.

Exercise 1: Planning the Hyper-V Host Deployment

Scenario

A. Datum is planning to purchase approximately 30 servers that it will deploy as Hyper-V servers. A. Datum plans to deploy most of the servers in the London data center, but is considering deploying some of the servers in the Toronto and Sydney data centers.

A. Datum will purchase the servers from suppliers located in each country, so the hardware specifications for the Hyper-V servers are not likely to be identical. Your design must optimize the virtual machine placement on the Hyper-V servers and allow for variations in server hardware configurations.

A. Datum has deployed a single VMM management server in the London data center and plans to centralize the management of the Hyper-V servers by using a VMM management server. A. Datum expects to create several virtual machine templates. The organization expects to store installation ISO files in the VMM environment. There is some concern about the amount of bandwidth between offices that this configuration will consume. Your design for the Hyper-V and VMM deployment should minimize the network bandwidth utilization between London, Toronto, and Sydney.

A. Datum wants to take advantage of the options available within VMM to manage host machines differently based on the host server location. You must ensure that different groups of administrators can manage the host servers in London, Toronto, and Sydney, and that the rules for distributing virtual machines automatically across the hosts are different for each location.

Supplemental Documentation

Email from Cristina Potra at A. Datum

From: Cristina Potra
Sent: 24 Aug 2014 09:05
To: Thomas Andersen
Subject: Re: VMM Deployment considerations
Thomas,

Regarding the deployment of VMM, we have discussed the deployment with the Network team and feel that there are some bandwidth constraints. We hope your deployment topology can avoid deployment of VHD files and ISO files across the WAN. There are no concerns between the primary and secondary data center in London should you wish to deploy across both.

I can tell you that all the hardware is now here. In Sydney and Toronto, the teams have received their servers. As you know, they do not have blade chassis in the branch offices. The hardware that they have in both locations is the following:

4 servers with 96GB RAM and 2 x 6 core processors and 1 server with 64GB RAM and 2 quad core processors. Our blades for the London chassis are 128GB RAM and 2 x 6 core processors.

I know from our last meeting that the Toronto representatives have requested that servers not be powered down during the working day. Is that feasible?

The development manager in Sydney has suggested that the web farm of four servers can have downtime on only two servers at a time. Sometimes, they have one of four offline. He is concerned that, once they are converted to virtual, problems may arise if two were on the same host and that host failed while a third was already offline.

I am aware that the Operations team has upgraded System Center Operations Manager and System Center Data Protection Manager to the latest version. They are used to monitor and protect SQL Server and a number of physical servers and applications at both sites.

I think that covers everything. Good luck with your design.

Regards,

Christina Potra

Project Manager

----- Original Message -----

From: Thomas Andersen
 Sent: 24 Aug 2014 08:45
 To: Cristina Potra
 Subject: VMM Deployment considerations

Cristina,

Do you have any information about the VMM deployment, specifically the new branches? I am working on the final design proposal and need all the requirements. Additionally, is there any other information that you have of which I should be aware?

Thank you,

Thomas

A. Datum Server Virtualization Strategy	
Document Reference Number: AD070405	
Document Author	Thomas Andersen
Date	25 th Aug
<ul style="list-style-type: none"> • Requirements Overview <ol style="list-style-type: none"> 1. Centralized management of host servers and virtual machines <ul style="list-style-type: none"> • Delegation of administration to branch offices • Accommodation for variations in server hardware • Minimized bandwidth utilization • Prevention of virtual machine migrations during the day (Toronto) 	
<ul style="list-style-type: none"> • Additional Information <ul style="list-style-type: none"> ○ Configure bare-metal deployment ○ Have a maintenance host in each location ○ Keep bandwidth utilization low ○ Provide management reporting on power savings 	
<p>Proposals</p> <ol style="list-style-type: none"> 1. Is it possible to have a single VMM management server that can accommodate the requirements, or should we implement three VMM management servers? 2. How many library servers should we deploy? 3. Can we prevent virtual machines from powering down during the day? 4. How can we keep the number of templates for server deployment low? 5. What measures will we take to avoid bandwidth utilization? 6. Which hosts servers may require different host reserve settings? Where do we choose these settings? 7. What is the benefit of a maintenance host? It seems like a waste of resources. Can it do anything else? 8. How can we get reports on power savings? 	

The main tasks for this exercise are as follows:

- Read the supporting documentation.
- Update the proposal document.
- Examine the suggested proposals in the Lab Answer Key.

► **Task 1: Read the supporting documentation**

Read the documentation that the student handbook provides.

► **Task 2: Update the proposal document**

Answer the questions in the proposals section of the A. Datum Server Virtualization Strategy document.

Question: Is it possible to have a single Virtual Machine Manager (VMM) server to accommodate the requirements or should you implement three VMM management servers?

Question: How many library servers should you deploy?

Question: Can you prevent virtual machines from powering down during the day?

Question: How can you keep the number of templates for server deployment low?

Question: What measures will you take to avoid bandwidth utilization?

Question: Which hosts servers may require different host reserve settings?
Where do you choose these settings?

Question: What is the benefit of a maintenance host? It seems like a waste of resources.
Can it do anything else?

Question: How can you get reports on power savings?

► **Task 3: Examine the suggested proposals in the Lab Answer Key**

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

Exercise 2: Configuring Hyper-V Host Groups

Scenario

Now that the design for the Hyper-V server and VMM deployment is complete, you need to implement the host group configuration in VMM. You will add the Hyper-V host to VMM and then configure the host groups based on the design.

The main tasks for this exercise are as follows:

- Add Hyper-V host to VMM.
- Create the host groups.
- Configure the host groups.

► **Task 1: Add Hyper-V host to VMM**

Add Hyper-V hosts

1. On LON-VMM1, on the desktop, double-click **Virtual Machine Manager Console**.
2. In the **Connect to Server** dialog box, ensure that the **Use current Microsoft Windows session identity** check box is selected, and then click **Connect**. The Virtual Machine Manager console opens.

3. Use the Add Resource Wizard to add a new Hyper-V host by using the following information:
4. Computer location: **Windows Server computers in a trusted Active Directory domain**
 - Credentials: **Adatum\Administrator** with the password **Pa\$\$w0rd**
 - Computer name: **LON-HOST1**
 - Host settings: **All Hosts**
5. Review the status, and then close the Jobs window when the task finishes. A warning message may appear, related to multipath I/O not being enabled. You can ignore this message for this task.
6. Repeat step 3 and 4 to add LON-HOST2 to the **All Hosts** group.

► Task 2: Create the host groups

Create the host groups

1. On LON-VMM1, from the Virtual Machine Manager console, click **VMs and Services**.
2. Click the **All Hosts** node, and then create the following host groups:
 - **London Hosts**
 - **Toronto Hosts**
 - **Sydney Hosts**

► Task 3: Configure the host groups

Configure host groups

1. On LON-VMM1, from the Virtual Machine Manager console, click **VMs and Services**.
2. Open the **Properties** page for **London Hosts**.
3. Change the Host Memory reserve to **1024 MB**.
4. Remove dynamic optimization from being inherited from the parent host group.
5. Configure automatic virtual machine migration, and change the frequency of migrations to every 60 minutes.
6. Enable Power Optimization with the default schedule.
7. Move **LON-HOST1** and **LON-HOST2** to the **London Hosts** host group.
8. Open the properties for **Toronto Hosts**.
9. Remove dynamic optimization from being inherited from the parent host group.
10. Configure automatic virtual machine migration, and change the frequency of migrations to every 60 minutes.
11. Enable the power optimization settings to prevent power optimization from running between 7 a.m. and 7 p.m. Monday to Friday.

Results: After completing this exercise, you will have added Microsoft Hyper-V® hosts to VMM and created and configured host groups.

Exercise 3: Configuring VMM Libraries

Scenario

You have decided to deploy a VMM library on a file server located in the Toronto data center, so that you can minimize the network utilization related to the Hyper-V deployment. You need to configure the VMM library on the server, and then distribute VMM components to each library. You also need to ensure that as local network administrators create virtual machines in the Toronto data center, they have access to all required files in the library.

The main tasks for this exercise are as follows:

- Configure a library server and share.
- Copy the required files to the branch office's data center.
- Assign a library to a host group.
- To prepare for the next module.

► Task 1: Configure a library server and share

Configure a library server and share

1. On TOR-SVR1, on the Server Manager Dashboard, click **File and Storage Services**, and then click **Shares**.
2. In the Shares workspace, click **Tasks**, and then create a new **SMB Share – Quick** share named **TORVMMLibrary**.
3. On LON-VMM1, in the Virtual Machine Manager console, click the **Library** workspace.
4. Add **TOR-SVR1\TORVMMLibrary** as a new library server with default resources. Use the following credentials:
 - User name: **Adatum\Administrator**
 - Password: **Pa\$\$w0rd**
5. Review the job status, and then close the Jobs window when the job is complete.

► Task 2: Copy the required files to the branch office's data center

Manually copy files to library shares

1. On LON-VMM1, in the Virtual Machine Manager console, click the **Library** workspace.
2. Expand **Library Servers**, and then click **LON-VMM1.Adatum.com**.
3. In the details pane, right-click **Blank Disk – Small.vhdx**, and then click **Open File Location**.
4. In the VHDs window, right-click **Blank Disk – Small.vhdx**, and then click **Copy**. Close the VHDs window, in the navigation pane, expand **Library Servers**, and then expand **tor-svr1.adatum.com**.
5. Right-click **TORVMMLibrary**, and then click **Explore**. In the TORVMMLibrary window, right-click an empty space, and then click **Paste**.
6. Close the TORVMMLibrary window. Right-click **TOR-SVR1.adatum.com**, and then click **Refresh**. Confirm that the **Blank Disk - small.vhdx** file appears in the objects list.

► **Task 3: Assign a library to a host group**

Assign a library to a host group

1. On LON-VMM1, in the Virtual Machine Manager console, click the **Library** workspace.
2. Expand **Library Servers**, and then right-click **tor-svr1.adatum.com**. Click **Properties**.
3. Add **tor-svr1.adatum.com** to the **Toronto Hosts** host group.
4. Close the Virtual Machine Manager console.

► **Task 4: To prepare for the next module**

Do not revert the virtual machines, as you will need them during the next module.

Results: After completing this lab, you will have configured a VMM library server, copied files to the branch office's data center, and assigned the library to a host group.

Question: What is the primary factor in design decisions for the VMM deployment?

Module Review and Takeaways

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
Unable to add VMware ESX or Citrix XenServer hosts	

Review Questions

Question: Why would you use an equivalent object? What steps do you need to take after you create it to ensure that VMM placement algorithms use it?

Question: What are some of the reasons for deploying more than one VMM management server?

Real-world Issues and Scenarios

When planning a virtualization project, there are a number of physical, data center, and host considerations. Server density can change, which will affect power consumption and the type of cooling required. For example, a fully populated blade chassis may change the cooling requirements for a data center if you wanted a farm of host servers with physical graphics processing units (GPUs) to provide Microsoft RemoteFX®. This could result in greater heat output, thereby increasing your cooling requirements.

Tools

- Microsoft Virtual Machine Converter Plug-in for VMware vSphere Client
- Microsoft Virtual Machine Servicing Tool 2012 at <http://go.microsoft.com/fwlink/p/?linkId=285267>
- Microsoft Assessment and Planning Toolkit at <http://go.microsoft.com/fwlink/p/?linkId=285265>
- System Center 2012 Virtual Machine Manager Component Add-ons and Extensions at <http://go.microsoft.com/fwlink/p/?linkId=285266>

Module 3

Planning and Implementing Networks and Storage for Virtualization

Contents:

Module Overview	3-1
Lesson 1: Planning a Storage Infrastructure for Virtualization	3-2
Lesson 2: Implementing a Storage Infrastructure for Virtualization	3-10
Lesson 3: Planning and Implementing a Network Infrastructure for Virtualization	3-16
Lesson 4: Planning and Implementing Network Virtualization	3-30
Lab: Planning and Implementing Virtualization Networks and Storage	3-38
Module Review and Takeaways	3-47

Module Overview

After host design and deployment, the next step in implementing a server virtualization infrastructure is to plan the storage and network components. Planning the storage infrastructure is particularly critical because you will be running many virtual machines that all share the same storage.

This module describes the factors that you must consider when you are planning the storage and network infrastructure for your virtual environment. You will learn how to deploy these components in Windows Server® 2012 Hyper-V®, Windows Server 2012 R2 Hyper-V, and Microsoft® System Center 2012 R2 Virtual Machine Manager (VMM).

Objectives

After completing this module, you will be able to:

- Plan a storage infrastructure for server virtualization.
- Implement a storage infrastructure for server virtualization.
- Plan and implement a network infrastructure for server virtualization.
- Plan and implement network virtualization.

Lesson 1

Planning a Storage Infrastructure for Virtualization

Before you can deploy virtual machines, you must plan and deploy the storage infrastructure that the virtual machines will use. In this lesson, you will explore the storage options available for virtual machines.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe storage options for server virtualization.
- Describe the virtual hard disk configuration options.
- Explain the considerations for implementing Fibre Channel storage.
- Explain the considerations for implementing Internet small computer system interface (iSCSI) storage.
- Explain the considerations for implementing network file system (NFS) and Server Message Block (SMB) 3.0 storage.
- Describe the options available for making storage highly available in server virtualization scenarios.

Considerations for Implementing Storage

When provisioning virtual machines, you must ensure that the underlying storage infrastructure is reliable and can provide adequate performance. Storage must be able to cope with peak utilization times, such as backups, antivirus sweeps, and multiple, concurrent virtual machine boots. Storage is one of the more complicated and costly resources to manage in virtualization projects. Organizations should design storage solutions that have the flexibility to scale up and meet future growth, but not overprovision capacity.

Windows Server 2012 R2 builds upon existing storage options and introduces new storage options for virtualization. These options provide small to midsize companies with highly available storage solutions. Formerly, these solutions were available only by investing in storage area network (SAN) technologies or third-party software. A later section of this module details the new storage features that are important to virtualization deployments.

In a virtualization environment, the virtual hard disk performance can affect the virtual machine's performance. Servers with adequate random access memory (RAM) and processor capacity can still experience unsatisfactory performance if you misconfigure the storage system or it becomes overwhelmed with traffic. It is important to ensure that the storage design provides adequate performance and includes a plan for monitoring storage for availability and performance.

When planning storage for virtualization hosts, you should:

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

Consider the following factors when planning the storage:

- High performance connection to the storage. You can locate virtual hard disk (.vhd or .vhdx) files on local or remote storage. When you locate these files on remote storage, you must ensure that there is adequate bandwidth and minimal latency between the host and the remote storage. Slow network connections to storage, or connections where there is latency, result in poor virtual machine performance.
- Redundant storage. The volume on which the virtual hard disk files are stored should be fault tolerant, whether the virtual hard disk is stored on a local disk or a remote SAN device. Hard disks do fail. Therefore, the virtual machine and the Windows Server 2012 R2 Hyper-V host should remain in operation after a disk failure. Replacement of failed disks should not affect the operation of the Hyper-V host or virtual machines.
- High performance storage. The storage device where virtual hard disk files are stored should have excellent I/O characteristics. Many enterprises use solid-state drive (SSD) hybrid drives in Redundant Array of Independent Disks (RAID) 1+0 arrays. This achieves maximum performance and redundancy, particularly when multiple virtual machines are running simultaneously on the same storage. This can place a tremendous I/O burden on a disk subsystem, so you must choose high performance storage, or your virtual machine performance may suffer.
- Adequate growth space. If you configure virtual hard disks to grow automatically, it is important that there is adequate space in which these files can grow. Additionally, you need to monitor growth carefully so that you experience no service disruptions when a virtual hard disk consumes all available space.
- Data protection. It is important to consider the performance of your backup solution and its impact on your storage design, as well as the amount of data virtual machines will host. Review existing data and ensure that you will be able to back up required virtual machines and their storage within an acceptable timeframe.

Hyper-V offers flexible storage options including most of the options that Windows Server supports, such as locally attached storage including Serial Advanced Technology Attachment (SATA), small computer system interface (SCSI), and SSD. In addition, Hyper-V supports remotely connected Fibre Channel and iSCSI SANs. Hyper-V supports running virtual machine in file shares using the new SMB 3.0 protocol. The following lessons cover these storage types in more detail.

Options for Configuring Virtual Disks

Hyper-V in Windows Server 2012 introduces a new virtual hard disk, *.vhdx*, which supports capacities up to 64 terabytes. *.vhdx* provides better protection against data corruption during power failures and has improved alignment to work on large sector physical disks.

Hyper-V supports a number of virtual hard disk types that utilize either *.vhd* or *.vhdx*, including:

- Dynamically expanding virtual hard disks
- Fixed-size virtual hard disks
- Differencing disks
- Physical disks or volumes

Hyper-V supports a number of different virtual hard disk types including:

- Dynamically expanding disks
- Fixed disks
- Differencing disks
- Physical disks

Dynamically Expanding Virtual Hard Disks

These virtual disks start very small and then grow as you write data to them. Typically, organizations have used dynamically expanding disks in test and development environments. However, performance improvements mean that you can use dynamic disks in production, because they often offer performance similar to that of fixed disks. A dynamically expanding disk grows only to the space that you allocate to it when you create it. You can expand the disk, if necessary, to the size of the corresponding disk format, which is 2 terabytes for .vhd and 64 terabytes for .vhdx.

One potential problem with using dynamically expanding virtual hard disks is that you must manage storage utilization after deployment. If you have multiple dynamically expanding virtual hard disks in a single storage location that is less than the virtual hard disks' total maximum size, you must monitor the storage location to ensure that the virtual hard disks do not expand to utilize all available space.

Another potential problem with dynamically expanding virtual hard disks is that fragmentation of the virtual hard disk file may occur on the host computer's physical hard disk. This could affect the virtual disk's performance.

Fixed-Size Virtual Hard Disks

These disks use as much physical disk space as you specify when you create the disk. For example, if you create a 100-gigabyte (GB) fixed-size virtual hard disk, it will use 100 GB of physical disk space. The primary benefit of using fixed-size disks is that all storage that the disks require is committed when you create the disks. This reduces the likelihood that you will overcommit your storage resources.

One reason for choosing fixed-size virtual hard disks is that dynamically expanding virtual hard disks may not support some applications. For example, Microsoft does not support Microsoft Exchange Server 2010 or Exchange Server 2007 deployed on dynamically expanding virtual hard disks.

One disadvantage of fixed-size disks is that the disks may take longer to move from one server to another.

Differencing Disks

A differencing virtual hard disk is associated with another virtual hard disk in a parent-child relationship. The differencing disk is the child, and the associated virtual disk is the parent. The parent disk can be any type of virtual hard disk. The differencing disk stores a record of all changes made to the parent disk, and then provides a way to save changes without altering the parent disk. In other words, by using differencing disks, you ensure that changes occur on the differencing disks and not on the original virtual hard disk. You can merge changes from the differencing disk to the associated virtual hard disk when appropriate.

The differencing hard disk expands dynamically as data intended for the parent disk writes to the differencing disk. You should write-protect or lock the parent disk. If another process modifies the parent disk without recognizing the differencing disk's parent-child relationship, then all differencing disks related to the parent disk become invalid. As a result, any data written to them is lost. However, by locking the parent disk, you can mount the disk on more than one virtual machine, similar to a read-only disk or CD.

You cannot specify a size for a differencing disk. Differencing disks can grow as large as the parent disks to which you associate them. However, unlike dynamically expanding disks, you cannot compact differencing disks directly. You can compact differencing disks only after merging the disk with a dynamically expanding parent disk.

If you are using differencing disks, it is important to have a standardized naming convention for your virtual hard drives. When you examine a virtual hard disk in Hyper-V Manager, it is not readily apparent whether it is a differencing drive or a parent disk.

Differencing disks have poor performance and are used to save space. Windows Server 2012 R2's deduplication functionality minimizes the space-saving advantage.

Physical Disks or Volumes

Hyper-V also allows for the use of a physical disk type (previously known as pass-through disk). When you configure a virtual machine to use a physical disk, the virtual machine will use an entire physical disk or volume on the host computer. If you intend to use physical disks to support an operating system installation, you must store the guest configuration file in an alternate location. This is because the operating system installation will consume the entire physical disk. For example, you could locate the configuration file on another internal drive in the Hyper-V server itself. Conversely, if the host computer is part of a failover cluster, you can host the configuration file on a separate cluster that provides highly available file services. Be aware that you cannot expand physical disks dynamically. Additionally, here are some other factors to consider:

- Physical disks improve performance
- Physical disks provide storage by enabling you to associate an external data source with a virtual machine. Then the virtual machine writes directly to the data source without encapsulation in a virtual hard disk. Compatible storage drives that you can use for physical disks for data sources include physical disks, partitions, logical unit numbers (LUNs), SANs, and iSCSI.
- Physical disks have no virtual hard disk. They use a physical drive instead.
- Physical disks forward all read and write requests directly to the physical volume.
- Physical disks do not support dynamically expanding virtual hard disks, differencing disks, or virtual machine snapshots.
- You must configure the physical disk as offline if you want the Hyper-V server to implement physical disks.
- If you are going to install an operating system on the physical disk, you must first bring the disk online on the host and initialize it, then take it offline. As a boot drive, you must use an IDE channel.
- If you plan to use the physical disk as a storage drive, you must prepare it in the operating system before you can place data on it.

Planning Fibre Channel Storage

You can add virtual Fibre Channel adapters to a virtual machine to enable it to access Fibre Channel storage on SANs. To deploy a virtual Fibre Channel:

- You must configure the Hyper-V host with a Fibre Channel host bus adaptor (HBA).
- The Fibre Channel HBA must have a driver that supports Virtual Fibre Channel.
- The virtual machine must support virtual machine extensions.

To support Fibre Channel, you must:

- Configure the Hyper-V host with a Fibre Channel HBA
- Ensure that the Fibre Channel HBA has a driver that supports Virtual Fibre Channel
- Ensure that the virtual machine supports virtual machine extensions

Virtual Fibre Channel adapters support port virtualization by exposing HBA ports in the guest operating system. This allows the virtual machine to access the SAN by using a standard World Wide Name (WWN) associated with the virtual machine.

Hyper-V supports Multipath I/O (MPIO) to provide highly available access to the LUNs that have been exposed to the host. With Virtual Fibre Channel adapters, you can provide access to LUNs directly from virtual machines that can also support MPIO. You can use a combination of both in your virtualization environment.

You will probably want to use Fibre Channel storage with your organization's virtualization infrastructure if your organization has an existing Fibre Channel deployment. As both iSCSI and Fibre Channel storage address the same needs, the nature of the existing infrastructure will influence the choice of storage infrastructure technology that you use with your virtualization infrastructure.

Planning iSCSI Storage

An inexpensive and simple way to configure a connection to remote disks is to use iSCSI storage. Many application requirements require that remote storage connections be redundant for fault tolerance or high availability. Many companies already have fault-tolerant networks that are less expensive than SANs for retaining redundancy.

When designing your iSCSI storage solution, you should consider the following:

- Deploy the iSCSI solution on at least a 1 gigabyte per second network. Review specific features that you intend to use, such as jumbo frames, and ensure that your hardware can support these features.
- Design a highly available network infrastructure, because network devices and components conduct the transfer of data from servers to iSCSI storage. This is a crucial consideration.
- Design an appropriate security strategy for your iSCSI storage solution.
- Involve all relevant teams in the storage implementation, and confirm whether automation of storage deployment will be possible for virtualization and application administrators.

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
- Involve all relevant teams

Planning SMB 3.0 and NFS Storage

SMB 3.0

SMB file share provides an alternative to storing virtual machine files on iSCSI or Fibre Channel SAN devices. Hyper-V supports the storage of virtual machine data, such as virtual machine configuration files, snapshots, and virtual hard disk files, on SMB 3.0 file shares. If you decide to implement SMB file shares, consider the following:

- The file share must support SMB 3.0, which limits placement of virtual hard disks on file shares that file servers that are running on a Windows Server 2012 or Windows Server 2012 R2 host. Earlier versions of Windows Server do not support SMB 3.0.

- SMB 3.0:
 - Enables virtual machine storage on SMB 3.0 file shares
 - Requires Windows Server 2012 or Windows Server 2012 R2 file servers
 - Requires fast network connectivity
 - Provides redundancy and performance benefits
- NFS:
 - Enables you to use NFS shares to deploy VMware to virtual machines

- You must ensure that network connectivity to the file share is 1 GB or more.
- When creating a virtual machine in Hyper-V on Windows Server 2012 or Windows Server 2012 R2, you can specify a network share when you select the virtual machine location and virtual hard disk location. In addition, you can attach disks stored on SMB 3.0 file shares. You can use both .vhd and .vhdx disks with SMB file shares.



Note: Hyper-V over SMB assigns the computer account permissions on the share, so you can configure it only in an Active Directory® Domain Services (AD DS) environment.

The SMB protocol in Windows Server 2012 and Windows Server 2012 R2 includes the following features:

- **SMB Transparent Failover.** The SMB protocol has the built-in ability to handle failure so that the client and server can coordinate a transparent move that allows continued access to resources with only a minor I/O delay. There is no failure for applications.
- **SMB Scale-Out.** The SMB Scale-Out feature allows you to access shares through multiple cluster nodes by using Cluster Shared Volumes (CSVs). This allows you to balance loads across a cluster.
- **SMB Direct (SMB or RDMA).** Formerly seen in high performance computing scenarios only, SMB Direct is now available in Windows Server 2012. SMB Direct allows a Remote Direct Memory Access–enabled (RDMA-enabled) network interface to perform file transfers by using technology onboard the network interface, without operating system intervention.
- **SMB Multichannel.** SMB Multichannel is enabled automatically. It allows SMB to detect a network's configuration. For example, if it detects that two network interfaces are configured and teamed on the client and server, SMB can utilize all available bandwidth.
- **SMB Encryption.** SMB Encryption allows for encryption without the need for Internet Protocol security (IPsec). You can configure it per share or at the server level. Older SMB clients cannot connect to encrypted shares or servers.
- **VSS for SMB File Shares.** Volume Shadow Copy Service (VSS) is enhanced to allow snapshots at the share level. Remote file shares act as a provider and integrate with a backup infrastructure.
- **SQL Server over SMB.** You can store Microsoft SQL Server® databases on SMB 3.0 shares, which could allow infrastructure consolidation.

NFS

NFS is a file sharing solution that uses the NFS protocol and enables you to transfer files between computers that are running Windows Server 2012, Windows Server 2012 R2, and operating systems other than Windows®.

Windows Server 2012 and Windows Server 2012 R2 include an updated NFS stack, which provides transparent failover to NFS clients by using continuously available NFS shares. You can use NFS as storage for VMware-based virtual machines. However, it is not an option for Hyper-V storage.

Planning High-Availability Options for Hyper-V Storage

Highly available storage options are necessary for ensuring that virtual machines run uninterrupted when a hardware component fails. Before you consider Hyper-V–specific storage options, you should consider the typical components that comprise storage hardware and how you can use these in storage solutions.

You can use the following basic hardware components when building a highly available storage solution:

- **Power.** Where possible, use redundant power supplies in storage solutions and, optionally, keep spares on site. The devices that support secondary power supplies are disk shelves, servers, and switches.
- **Disks.** A highly available storage solution should be able to tolerate a disk failure and continue serving storage without interruption. Typical vendor storage controllers offer a variety of RAID levels that allow for the loss of one or more disks.
- **Storage network hardware.** Usually, creating a highly available hardware solution involves removing all single points of failure, including the storage network components, such as HBAs or network adapters. This may require that you use MPIO and have two storage switches and Storage Controllers. Network Teaming may be necessary when you are using multiple storage paths.

Hyper-V highly available storage options include:

- Storage spaces that provide tiering, mirroring, and parity resilience across multiple heterogeneous disk types
- Multiple HBAs with storage path redundancy via MPIO or adapter teaming
- Failover clustering that provides highly available iSCSI targets or continuously available SMB and NFS shares

There are several highly available storage options in Windows Server 2012 and Windows Server 2012 R2, including:

- **Storage spaces.** Storage spaces allow the use of universal serial bus (USB), Fibre Channel, iSCSI, Serial Attached SCSI, and SATA attached disks that create virtual disks that span all of these technologies. Windows provides a level of disk protection by creating these virtual disks by using mirror or parity protection. These allow disk failure without loss of data. You will learn more about this in a later module. However, you should understand that Windows Server can create pools of storage across heterogeneous disk types, and offers several solutions that can all be highly available. These solutions include NFS and SMB file shares and iSCSI disks. In Windows Server 2012 R2, storage spaces also support storage tiering, allowing commonly accessed data blocks to be stored on faster disks, increasing performance.
- **Failover clustering.** iSCSI targets and NFS and SMB shares are compatible with Windows failover cluster resources. Therefore, you can eliminate a single Windows server as a point of failure in a storage solution. However, to present storage spaces as a clustered resource, you can use Serial Attached SCSI– attached storage only.
- **MPIO.** MPIO is a Windows feature that allows Fibre Channel or iSCSI storage to be accessible over multiple paths. This is necessary on a client, such as Hyper-V host, so that it can connect to the storage solution. When using two paths, MPIO sees both storage locations, but determines that they actually are a single device. MPIO handles failover of paths and uses load balancing algorithms to provide agreed access to storage through the least saturated path. A later module in this course provides more detail about this topic.

- Network Teaming and fault tolerance. If you want to connect Hyper-V hosts to storage, we recommend having multiple network cards and paths. You should provide redundancy by using SMB or NFS network teaming on both the file server and the Hyper-V hosts. If you use SMB, SMB multichannel will provide the best available bandwidth automatically.



Reference Links: For more information, go to The Microsoft Storage Team Blog at <http://go.microsoft.com/fwlink/?LinkID=285270>.

Lesson 2

Implementing a Storage Infrastructure for Virtualization

Now that you understand the storage options available with Hyper-V, the next step is to implement and configure storage. One of the options available for implementing Hyper-V storage is iSCSI.

In this lesson, you will learn how to configure iSCSI storage and how to manage storage from VMM.

Lesson Objectives

After completing this lesson, you will be able to:

- Explain how to configure iSCSI.
- Configure an iSCSI target.
- Explain how to manage storage using VMM.
- Configure virtual disks and storage in VMM.

iSCSI Configuration Overview

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. You have the option of utilizing specialized hardware, such as HBA or network switches. iSCSI uses TCP/IP (typically, TCP port 3260). This means that iSCSI enables two hosts to negotiate tasks, such as session establishment, flow control, or packet size, and then exchange SCSI commands by using an existing Ethernet network. iSCSI uses a high performance, local storage bus subsystem architecture, which it emulates over LANs and WANs, thereby creating a SAN.

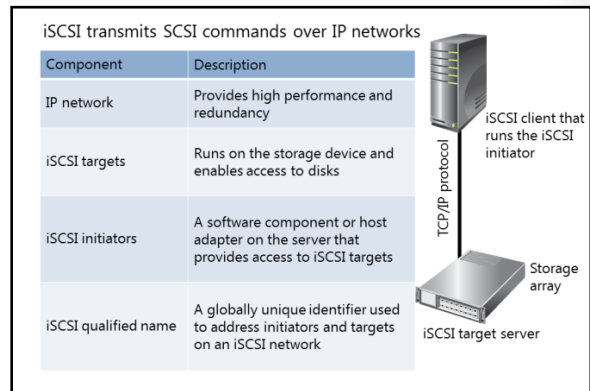
Unlike some SAN technologies, iSCSI requires no specialized cabling. You can run it over the existing switching and IP infrastructure. However, as a best practice, you can increase the performance of an iSCSI SAN deployment by operating it on a dedicated network or subnet.



Note: You can use either a standard Ethernet network adapter or dedicated iSCSI HBAs to connect the server to the iSCSI storage device.

An iSCSI SAN deployment includes the following:

- TCP/IP network. You can use standard network interface adapters and 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 Gbps and multiple paths to the iSCSI target. As a best practice, use a dedicated physical and logical network to achieve fast, reliable throughput.



- iSCSI targets. iSCSI targets present or advertise storage similarly to controllers for hard disk drives of locally attached storage. However, you access this storage 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 and at least one Ethernet adapter. Windows Server 2012 provides the iSCSI target server, a role service that acts a driver for the iSCSI protocol.
- iSCSI initiators. The iSCSI target displays storage to the iSCSI initiator, or *client*, which acts as a local disk controller for the remote disks. Windows Server 2008 and newer versions include the iSCSI initiator and can connect to iSCSI targets.
- iSCSI Qualified Name. iSCSI Qualified Names are unique identifiers for initiators and targets on an iSCSI network. When you configure an iSCSI target, you must configure the iSCSI Qualified Name for the iSCSI initiators that will connect to it. iSCSI initiators also use iSCSI Qualified Names to connect to the iSCSI targets. However, if name resolution on the iSCSI network is a possible problem, you can identify iSCSI endpoints (both target and initiator) by their IP addresses.

Demonstration: Configuring iSCSI Storage for Virtualization

In this demonstration, you will:

- Add the iSCSI target server role service.
- Create a storage pool.
- Create an iSCSI virtual disk.
- Create an iSCSI target.
- Connect to an iSCSI target.

Demonstration Steps

Add virtual disks to LON-SVR1

1. On LON-HOST1, create the folder **C:\StoragePool**.
2. In Hyper-V Manager, edit the settings of 20414C-LON-SVR1, and add and attach three 50 GB dynamically expanding .vhdx format virtual hard disks named **iSCSI1.vhdx**, **iSCSI2.vhdx** and **iSCSI3.vhdx** to the SCSI controller. Configure these disks to be stored in the **C:\StoragePool** folder.
3. Start LON-SVR1, and then sign in as **Adatum\Administrator** with the password **Pa\$\$w0rd**.

Add the iSCSI Target Server Role Service

- On LON-SVR1, in Server Manager, add the **iSCSI Target Server** role (located under File and iSCSI Services).

Create a storage pool

1. In Server Manager, in **File and Storage Services**, click **Storage Pools**.
2. In the Storage Pool section, create a new storage pool named **VMPool** that uses the three new disks.
3. Complete the New Storage Pool Wizard, accepting the default settings. On the **Results** page, click **Create a virtual disk when this wizard closes**, and then click **Close**. The New Virtual Disk Wizard launches.

4. Follow the wizard, and then on the **Virtual Disk Name** page, in the **Name** field, type **VMStorage**. Choose **Parity** and **Thin** provisioning, and then make the size **100 GB**.
5. Close the **Results** page.
6. When the New Volume Wizard launches, follow the wizard, and on the **Server and Disk** page, in the Disk area, click the **VMStorage Virtual Disk**. On the **Size** page, leave **(99.9) GB**.
7. On the **Drive Letter or Folder** page, leave the drive letter **F**. On the **File System Settings** page, type **VMStorage** for the **Volume Label**.
8. In the File and Storage Services pane, click **iSCSI**.
9. In the iSCSI Virtual Disks pane, create a new iSCSI virtual disk.
10. Under **Storage location**, click **F**, and then use **LONHOST1-iSCSIDisk1** for the **iSCSI virtual disk name**. Make the size **90 GB and dynamically expanding**, and then create a new **iSCSI target** named **LON-HOST1**.
11. On the **Specify access servers** page, browse to **LON-HOST1**.
12. Complete the wizard, leaving the defaults, and then click **Close**.
13. In the iSCSI Virtual Disks pane, click **Tasks**, and then in the **Tasks** drop-down list box, click **New iSCSI Virtual Disk**.
14. Create another iSCSI virtual disk on **C:** with the name **iSCSIDisk2**.
15. Make the size **5 GB**.
16. Assign **LON-HOST1** for the target.

Configure iSCSI initiators

1. On LON-HOST1, in Server Manager, click **Tools**, click **iSCSI Initiator**, and then when prompted to start the Microsoft iSCSI service, click **Yes**.
2. On the **Targets** page, in the **Target** field, type **172.16.0.12**. Click **Quick connect**, click **Done**, and then click **OK** to close the page.
3. On LON-HOST1, press and hold the Windows key, press **X**, and then click **Disk Management**.
4. Find the new 90 GB disk, bring the disk online, initialize the disk, and create a new simple volume. Leave the default size, assign the letter **V**, and then add a label **VMStorage**.
5. Close Disk Management.

Managing Storage in VMM

Managing Host Storage

VMM can access the storage attached to Hyper-V hosts. When VMM makes virtual machine placement calculations, it can calculate the amount of free storage on the associated host disks. When you add storage to a virtualization host, it will not be visible in VMM until you refresh the host. To access newly added storage immediately, you can perform a manual refresh to make it visible.

VMM can discover, classify, and provision remote storage on supported storage arrays through the VMM console. To implement this storage management, perform the following steps:

1. Install an SMI-S storage provider
2. Connect to the SMI-S storage provider to discover the storage
3. Classify the storage
4. Create logical units from the storage pool
5. Allocate logical units or storage pools to hosts, host groups, or host clusters

To see and work with host storage, you can right-click a host from the Properties page, and then click Storage. In the Storage area, you will see areas for each disk subsystem type, which you can expand, review, or configure.

Managing Storage Arrays

Through the VMM console, you can discover, classify, and provision remote storage on supported storage arrays. VMM fully automates storage assignment to a Hyper-V host or Hyper-V host cluster, and then tracks any storage that it manages. To enable new storage features, VMM uses the new Storage Management Service to communicate with external arrays through a Storage Management Initiative Specification (SMI-S) provider.

By default, the Storage Management Service installs during the VMM installation. Then, in order to manage storage, you must install a supported Storage Management Initiative Specification provider on an available server and add the provider to VMM.

Storage Overview Display

In the Fabric workspace, select Storage, and then click Overview on the ribbon to display the Storage Overview. As your virtual data center grows, the overview displays what the data center is provisioning and which resources continue to remain available.

You must complete the following steps to discover, classify, and assign storage through VMM:

1. For a supported storage array, obtain a Storage Management Initiative Specification storage provider from your storage array vendor. Then install the Storage Management Initiative Specification storage provider on an available server as instructed by your storage vendor.
2. From the VMM console, in the storage node, connect to the Storage Management Initiative Specification storage provider to discover and classify the storage. Connect to the provider by using either the Internet Protocol version 4 (IPv4) address or the fully qualified domain name (FQDN).
3. Classifying storage entails assigning a meaningful classification to storage pools. For example, you may assign a classification of Gold to a storage pool that resides on the fastest, most redundant storage array. This enables you to assign and use storage-based classification without remembering the specifics of its hardware characteristics. You can use these classifications when providing virtual machine self-service solutions. Service consumers pay more for Gold tier storage than they do for Bronze tier storage.
4. In the storage node, you have the option to create logical units from a managed storage pool.
5. From either the VMM console storage node or the target host group's Properties dialog box, you must allocate precreated logical units or storage pools to specific host groups. If you allocate storage pools, you can create and assign logical units directly from managed hosts in the host group that can access the storage array. In addition, if you use rapid provisioning to provision virtual machines by using SAN snapshots or cloning, VMM can create logical units from the storage pool automatically.
6. In the VMM console, from either the host or host cluster Properties dialog box, assign logical units from the host group to specific Hyper-V hosts or to Hyper-V host clusters, as shared CSV or available storage. If you allocated a storage pool to a host group, you can create and optionally assign logical units directly from a host or host cluster's Properties dialog box. If the storage array supports iSCSI host connectivity, you can create iSCSI sessions to the storage array from a host's Properties dialog box.



Reference Links: For a list of supported storage arrays, see Supported Storage Arrays for System Center 2012 VMM at:
<http://social.technet.microsoft.com/wiki/contents/articles/16100.supported-storage-arrays-for-system-center-2012-vmm.aspx>.

Assigning Storage to Host Groups

You can assign storage pools to host groups. By using host groups, you can define logical groups of physical computing resources. These groups can help VMM and administrators determine placement for deploying new virtual workloads. For example, you could create host groups and assign storage pools with classifications as shown in the table.

Host group name	Storage classification	Storage type	Host server central processing unit (CPU)
Tier 1	Platinum	Solid-state drives (SSD)	3.6 GHz
Tier 2	Gold	Serial Attached SCSI 15,000 Revolutions Per Minute (RPM)	3.0 GHz
Tier 3	Silver	Serial Attached SCSI 10,000 RPM	2.4 GHz
Tier 4	Bronze	Serial ATA 7200 RPM	2.0 GHz

To add storage to a host group, you must perform the following steps:

1. Right-click the host group, and then click **Properties**.
2. Click **Storage**, and then click **Allocate Storage Pools** or **Allocate Logical Units**.
3. Add storage as required.



Note: The iSCSI Target Storage Management Initiative Specification Provider for Windows Server can be found on the System Center Virtual Machine Manager 2012 Service Pack 1 (SP1) Installation CD in path: \amd64\Setup\msi\iSCSITargetSMISProvider.msi, or on the VMM Server under \Program Files\Microsoft System Center 2012\Virtual Machine Manager\setup\msi\iSCSITargetProv\iSCSITargetSMISProvider.msi.

Demonstration: Configuring Storage in VMM

In this demonstration, you will see how to use VMM to review and configure storage for virtualization hosts.

You will also review how to add classifications to storage and how to add storage providers.

Demonstration Steps

Update the LON-HOST1 virtual machine placement path in VMM

1. Open the Virtual Machine Manager console.
2. Click the **Fabric** workspace, expand **Servers**, click **LON-HOST1.adatum.com** on the Managed Computers pane, and then click **Refresh** on the ribbon.
3. On the ribbon, click **Home**, click **Jobs**, wait for **Refresh host job** to complete, and then close the Jobs window.
4. Click **LON-HOST1.adatum.com**, and then click **Properties** on the ribbon.
5. Click **Placement Paths**, click **Add** next to Specify default virtual machine paths to be used during virtual machine placement, click **VMStorage**, and then click **OK**. Click **OK** again to close the **Properties** page.

Create storage classifications

1. In the Virtual Machine Manager console, click the **Fabric** workspace, and then click **Storage**.
2. On the ribbon, click **Create Storage Classification**, in the **Name** field, type **Gold**, in the **Description** field, type **15K SAS Drives**, and then click **Add**.
3. On the ribbon, click **Create Storage Classification**, in the **Name** field, type **Silver**, in the **description** field, type **7K SATA Drives**, and then click **Add**.
4. Expand Storage, Click the **Classification and Pools** node, and then note that there is no capacity available.

Add storage providers

1. On LON-VMM1, in the Virtual Machine Manager console, click **Fabric**, right-click **Storage**, and then click **Add Storage Devices**.
2. On the **Select Provider Type** page, click **Windows-based file server**, and then click **Next**.
3. Click the **Provider IP address or FQDN** field, type **lon-svr1.adatum.com**, and then click **Browse**.
4. On the **Select a Run As account** page, click **Administrator**, and then click **OK**.
5. On the **Specify Discovery Scope** page, click **Next**.
6. On the **Gather Information** page, review the discovery result, and then click **Next**.
7. On the **Select Storage Devices** page, click **Next**.
8. On the **Summary** page, click **Finish**, and then close the Jobs window.

Create file shares from VMM

1. On LON-VMM1, click **Fabric**, and then on the ribbon, click **Create File Share**.
2. On the **Create File Share** page, in the **Name** field, type **GoldDisks**. In the **Local path** field, type **F:**, and then click **Create**.

Lesson 3

Planning and Implementing a Network Infrastructure for Virtualization

Your network infrastructure is an important component of implementing a server virtualization solution. In a server virtualization deployment, many virtual machines need to share the network adapters that the host machine provides.

In this lesson, you will learn about how to plan your network infrastructure for a virtualization environment, including the facets of host and guest planning, storage networking, clustering, and multitenant hosting.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the requirements for Hyper-V networking.
- Explain the options for virtual network configuration.
- Plan virtual switch extensibility.
- Plan high performance options for networking.
- Configure virtual networks.
- Explain the network planning considerations for a Hyper-V deployment.
- Manage virtual networks in VMM.
- Configure virtual network components in VMM.

Networking Requirements for a Hyper-V Deployment

Many network devices now have virtual equivalents, or *virtual appliances*, and Hyper-V provides support for many demanding applications that were not virtualization candidates previously. This means that you can build very large Hyper-V clusters that are capable of hosting scores of virtual machines, each with its own networking requirements. Furthermore, this means that your virtualization network planning is dependent fully on infrastructure or solution requirements.

You should identify the network requirements for a Hyper-V deployment during the host and guest environment planning process. You must define the features that you plan to use so that you can determine what requirements you need to meet.

In older Hyper-V versions, optimal designs for the host clusters required you to dedicate multiple network adapters to different services. Windows Server 2012 and newer implementations support vendor-independent network teaming, which allows you to aggregate host network adapters. This provides

Network options:

- Host type: standalone, networked, clustered
- Storage type: iSCSI, SMB 3.0, or Fibre Channel
- Live and storage migration traffic
- Replica and backup traffic

Virtual machine network options:

- Private, external, isolated, VLAN

Network adapters and switches:

- Bandwidth
- Teamed, multiple paths support VMQ and SR-IOV

hardware redundancy and higher aggregated bandwidth. You can use Quality of Service (QoS) to guarantee bandwidth to multiple services that are running on the same aggregated network adapters. This has the potential to reduce the complexity and cost of the physical infrastructure.

Some considerations for planning requirements include:

- Stand-alone Hyper-V hosts. It is possible to run a Hyper-V server in isolation, with no network adapters. You can use Hyper-V Manager to create internal private networks that will allow multiple virtual machines to communicate. This configuration is useful only in test and development environments, in which the virtual machines do not need connectivity to any other computers on the network.
- Network-enabled Hyper-V hosts. If you want to manage a Hyper-V host server remotely, or have its guest virtual machines communicate with other machines or devices on the physical network, you must have at least one physical network adapter in the host machine. You can configure Hyper-V so that the host and its virtual machines share a single network adapter. This ensures that all network connectivity occurs over one network adapter. We recommend that, for resiliency and performance, you configure more than one network adapter in a host.
- Hyper-V clusters. You can build Hyper-V host clusters by using a number of solutions. Each has its own specific networking requirements, such as SMB 3.0-enabled file shares or utilization of iSCSI and Fibre Channel SANs. When using SMB 3.0 or iSCSI, you will need at least one physical network adapter. We recommend strongly that you deploy at least two physical network adapters. One is for the host and virtual machine connectivity to the rest of the network, and one is for accessing the storage network. Modern network adapters support various features that assist specifically with virtualization and iSCSI storage, such as Virtual Machine Queue (VMQ), single root I/O virtualization (SR-IOV), and IPsec task offloading. When you are considering which features you want to use, you should remember that you must configure the network adapters for your hosts to support these features.
- Storage migration. Storage migrations provide the ability to migrate virtual machine storage from one location to another. If you are using iSCSI or SMB shares for your storage, you need to consider the impact that this migration will have on other network traffic. You can set the number of simultaneous migrations. Additionally, if you will be using this feature in production, you should test to see both the network and disk I/O impact to determine optimal settings.
- Live migration. Hyper-V hosts can send and receive live migrations of virtual machines, but live migration traffic can affect other traffic. Therefore, you should configure it to use any available network or a dedicated network. You can set the number of simultaneous live migrations to best suit the available bandwidth. However, you need to test this extensively when designing the network for Hyper-V host clusters, because this will determine the time that is necessary to evacuate a host. Additionally, you can determine the impact of virtual machine movement when you enable dynamic optimization and when clusters move between virtual machines to balance resource utilization.
- Hyper-V Replica. You can use the Hyper-V Replica feature in Windows Server 2012 to replicate virtual machines between Hyper-V host servers. You should consider the bandwidth required between hosts and what impact this may have on other network traffic:
 - You can schedule initial Hyper-V replication of virtual machines and, optionally, send the initial replication to a portable disk drive. Then you can import this to the target server. We recommend that you test replication in a nonproduction environment and calculate the amount of synchronization traffic between hosts. You can configure Hyper-V Replica to replicate to specific named Hyper-V hosts and clusters or any Hyper-V hosts.
 - You can assign the HTTP or HTTPS port that you want to use for replication. Additionally, you need to ensure that you configure all firewalls between Hyper-V hosts to allow the replication traffic.

- Network hardware. You should determine what your network hardware requirements are, and ensure that the end-to-end hardware will support your chosen feature set. If you want to isolate virtual machines by using virtual LANS (VLANs), you must ensure that your network adapters and switches between each cluster node support this. You should consider the newer Hyper-V network features and determine whether the network adapters that you plan to purchase will support them. Some network adapter hardware features can improve virtual machine network performance, including:
 - VMQ. VMQ is a feature that uses hardware packet filtering to deliver data directly to virtual machines from an external network. This reduces the overhead of routing packets from the management operating system to the virtual machine.
 - SR-IOV. The SR-IOV interface is an extension of the peripheral component interconnect (PCI) Express specification. SR-IOV allows a device, such as a network adapter, to distribute access to its resources among PCI Express hardware functions, including:
 - PCI Express physical function.
 - The device's primary function and the advertised SR-IOV capabilities.
 - The primary function associated with the Hyper-V parent partition in a virtualized environment.
 - IPsec task offloading. Hyper-V 3.0 enables IPsec task offloading at the machine level, which moves some demands on the virtual machine's CPU to a dedicated processor on the physical network adaptor.

The follow planning questions can help you identify networking requirements for a Hyper-V deployment:

- What storage technology is in place or planned? If it is iSCSI or SMB 3.0 file shares, you can use a separate physical network infrastructure, or take advantage of QoS and use a converged network.
- How many guests per host will there be? What sort of bandwidth is necessary for guests and hosts? What monitoring statistics are available for migration of existing physical servers?
- Will you isolate management traffic to hosts?
- Will you dedicate networks to backup and replication traffic?
- When deploying clustering, will you use a dedicated or shared network adapter for live migration traffic?
- How will you secure network traffic?
- Will you be operating a multitenant virtualization environment? Will you use VLANs?
- Will your network team manage all network settings and IP address assignment? Will they provide address ranges to assign to VMM pools?
- How many hosts and how many different switch configurations will you require? Creating them in VMM may save considerable time.

Configuration Options for Virtual Networks

You can create three types of virtual switches on your Hyper-V host servers, including:

- **External.** You can create a virtual network switch that you bind to a physical network adapter in the host server. Once you have created this virtual network switch, you can connect one or more virtual machine network adapters, permitting virtual machines' access to a physical network. You can create only one external virtual switch for each physical network adapter or set of teamed adapters on the virtualization host. Optionally, you can allow the host to share the network adapter with the virtual switch.

Virtual switch types:

- External
- Internal
- Private

Network adapter settings:

- MAC Address
- DHCP Guard
- Router Guard
- Port Mirroring
- NIC Teaming
- VLAN

- **Internal.** Creating an internal switch allows virtual machines to communicate with each other and with the Hyper-V host. However, internal switches do not allow any communication with the physical network.
- **Private.** The private virtual switch allows virtual machines to communicate with each other. You can create multiple private virtual switches on a single Hyper-V host, to isolate different groups of virtual machines.

VLAN settings and external virtual switches enable you to share the network adapter with the virtual guest machines. If you do share the network adapter, you can set VLAN IDs for the host server. However, this does not control virtual machine VLAN configuration.

Network Adapter Settings

You can configure the following network adapter settings:

- **MAC Address.** At the Hyper-V switch level, you can define the range of addresses that Hyper-V can assign dynamically to virtual network adapters. At the virtual network adapter, the default is dynamic, although you can set a range manually, if necessary. You can also enable MAC address spoofing. By default, this is disabled, and the server can use only the assigned MAC address to send and receive packets.
- **DHCP Guard.** The Dynamic Host Configuration Protocol (DHCP) guard drops DHCP server messages from unauthorized virtual machines that are pretending to be DHCP servers. You can set this at the network adapter level.
- **Router Guard.** A router guard, which you can set at the virtual machine level, drops router advertisement and redirection messages from unauthorized virtual machines that are pretending to be routers.
- **Port Mirroring.** Port mirroring allows network traffic to travel to and from a virtual machine for monitoring. It does this by copying incoming and outgoing packets and then forwarding the copies to another virtual machine that you configure for monitoring.
- **NIC Teaming.** By enabling NIC teaming in the virtual network adapter settings, you can use NIC teaming within a virtual machine. This is useful when the Hyper-V server does not use NIC teaming and you want to create a redundant connection across multiple virtual switches.
- **VLAN.** You can assign each guest to a VLAN, if necessary.

Planning for Virtual Switch Extensibility

Prior to Windows Server 2012, Hyper-V included a simple virtual switch that was built on a closed architecture. It provided only basic networking functionality and was not extensible in any way. Windows Server 2012 and later versions of Hyper-V use a completely redesigned virtual switch, which is built on an open framework. This virtual switch is extensible and allows developers to extend existing features and add new features. For example, other companies can add their own monitoring, filtering, and forwarding features without having to replace all of the Hyper-V

Extension	Purpose	Extensibility component
Network packet inspection	Inspects network packets that are exchanged between virtual machines and passed through a virtual switch	NDIS filter driver
Network packet filter	Creates, filters, and modifies network packets that are entering or leaving the virtual switch	NDIS filter driver
Network forwarding	Provides network packets with a forwarding logic extension	NDIS filter driver
Intrusion detection or firewall	Filters and modifies network packets, monitors or authorizes connections, and filters traffic based on different criteria	WFP callout driver

virtual switch functionality. In addition, you can implement extensions by using network device interface specification (NDIS) filter drivers or Windows Filtering Platform (WFP) callout drivers, which are two public Windows platforms used for extending the Windows networking functionality. You can use each platform to extend the virtual switch in different ways. NDIS filter drivers or WFP callout drivers have the following characteristics:

- **NDIS filter driver.** The NDIS filter driver is a filtering service to monitor and modify network packets in Windows. For example, you can use the NDIS filter driver to perform packet inspection, modify packets when transiting virtual switch, or perform packet forwarding based on packet content. NDIS filters were introduced with the NDIS 6.0 specification, which Windows Server 2008 and Windows Vista® first implemented.
- **WFP callout drivers.** Windows Server 2008 and Windows Vista first implemented WFP callout drivers. Developers can use these drivers to filter and modify TCP/IP packets, monitor or authorize connections, filter IPsec-protected traffic, and filter remote procedure calls (RPCs). Filtering and modifying TCP/IP packets provides unlimited access to the TCP/IP traffic that passes through the virtual switch. WFP callout drivers can examine and modify outgoing and incoming packets before additional processing occurs. By using WFP callout drivers, developers can create firewalls, antivirus software, diagnostic software, intrusion detection software, and other applications and services.

Non-Microsoft Extension Support

Non-Microsoft extensions can extend three aspects of the switching process: inbound filtering, destination look-up and forwarding, and outbound filtering. In addition, monitoring extensions can gather statistical data by monitoring traffic at different layers of the virtual switch. You can add multiple monitoring and filtering extensions to a virtual switch. However, you can use only one instance of the forwarding extension per switch instance. If you use a non-Microsoft forwarding extension, it will override the default forwarding of the virtual switch.

After you install virtual switch extensions, you can control them on the Extensions settings for the virtual switch, or by using Windows PowerShell®. By default, Hyper-V includes two virtual switch extensions. These form the Microsoft NDIS Capture monitoring extension, which is disabled by default, and the Microsoft Windows Filtering Platform filtering extension, which is enabled by default.

The following table lists some of the virtual switch extensions, the functionalities they provide, and which platform you can use to provide such functionality.

Extension	Purpose	Extensibility component
Network packet inspection	Inspects network packets that are exchanged between virtual machines and passed through a virtual switch. You cannot modify network packets.	NDIS filter driver

Extension	Purpose	Extensibility component
Network packet filter	Creates, filters, and modifies network packets that are entering or leaving the virtual switch.	NDIS filter driver
Network forwarding	Provides network packets with a forwarding logic extension. This extension replaces the default forwarding extension, because the virtual switch can have only one forwarding extension.	NDIS filter driver
Intrusion detection or firewall	Filters and modifies network packets, monitors or authorizes connections, and filters traffic based on different criteria (for example, if IPsec protects the network packets).	WFP callout driver

Getting Started Writing a Hyper-V Extensible Switch Extension

[http://msdn.microsoft.com/en-us/library/windows/hardware/jj673961\(v=vs.85\).aspx](http://msdn.microsoft.com/en-us/library/windows/hardware/jj673961(v=vs.85).aspx)

Planning High Performance Options for Networking

Windows Server 2012 supports several high performance options for networking. These options allow you to virtualize workloads that you would previously have deployed physically because of unusual network performance requirements. These options are SR-IOV, Virtual Machine Queue, Dynamic Virtual Machine Queue, NIC Teaming, and Virtual Receive Side Scaling.

SR-IOV

SR-IOV is a standard that specifies how a hardware device can make its functionality available for direct use by virtual machines. SR-IOV virtual functions are associated with physical hardware functions.

- SR-IOV allows multiple virtual machines to share the same PCI Express physical hardware resources
- VMQ delivers network traffic directly to the guest
- Dynamic VMQ allows the VMQ to be associated with a processor dynamically
- NIC Teaming
 - Provides redundancy and aggregates bandwidth
 - Is supported at the host and virtual machine level
- Virtual Receive Side Scaling balances network processing load across multiple virtual processor cores

SR-IOV in Hyper-V uses remapping of interrupts and direct memory access (DMA), and allows you to assign SR-IOV-capable devices directly to a virtual machine. Hyper-V enables support for SR-IOV-capable network devices. It also allows you to assign an SR-IOV virtual function of a physical network adapter directly to a virtual machine. When you do this, the network adapter bypasses the virtual switch, increasing network throughput and reducing network latency and CPU overhead on the Hyper-V host.

If you want to use SR-IOV, the Hyper-V host hardware, the network device, and its device driver must support it. Because SR-IOV requires compliant hardware, it can be associated with only an external virtual switch that maps to an SR-IOV-capable network adapter in the Hyper-V host. You can configure SR-IOV only at the time that you create the virtual switch. You cannot convert an external virtual switch with SR-IOV enabled to an internal or private switch. You can enable SR-IOV on virtual machine network adapters.

In Windows Server 2012 and newer Windows Server operating systems, you can use live migration to move running virtual machines without noticeable downtime, even when virtual machines are configured to use SR-IOV. During live migration, Hyper-V can check whether the destination server has SR-IOV capabilities and if so, move the virtual machine to the destination server. You can configure live migration to refuse migrations of SR-IOV-dependent virtual machines to a Hyper-V host that does not have SR-IOV capabilities.

You can use live migration to move virtual machines that are configured to use SR-IOV between Hyper-V hosts, even if those Hyper-V hosts have different SR-IOV-enabled network adapters. When you move a virtual machine, you will notice that it uses a different network adapter, but the configuration and network connectivity is preserved.

If you want to enable and use SR-IOV, the Hyper-V host must meet the following requirements:

- Server hardware must support SR-IOV. This means it must include chipset support for interrupt and DMA remapping, in addition to firmware support to make the hardware system SR-IOV capabilities available to the Windows Server operating system.
- An SR-IOV-capable network adapter and network adapter device driver must be present on the Hyper-V host (in the parent partition). The network adapter device driver must be present in each virtual machine where an SR-IOV-capable network adapter (its virtual function) is assigned.



Note: When you use SR-IOV, virtual machine traffic bypasses the Hyper-V virtual switch. If you have set any switch port policies, SR-IOV functionality is disabled for that virtual machine.



Everything you wanted to know about SR-IOV in Hyper-V. Part 1

<http://blogs.technet.com/b/jhoward/archive/2012/03/12/everything-you-wanted-to-know-about-sr-iov-in-hyper-v-part-1.aspx>

Virtual Machine Queue

Hyper-V in Windows Server 2008 R2 first supported Virtual Machine Queue (VMQ), a hardware virtualization technology that improves the network performance of virtual machines.

VMQ uses network adapter queues to:


- Classify received packets.
- Group received packets.
- Apply VLAN filtering.
- Provide concurrent processing on the network traffic for multiple virtual machines.
- Distribute interrupts to multiple cores for multiple virtual machines.
- Avoid copying receive buffers to virtual machine address spaces.

VMQ allows the efficient transfer of the incoming network traffic to a virtual machine. A VMQ-capable network adapter can use DMA to transfer incoming packets to the appropriate virtual machine. This reduces CPU overhead when transferring packets to the virtual machines. This is most beneficial when virtual machines are receiving large amounts of traffic due to tasks such as file backup, database replication, or data mirroring.

Dynamic Virtual Machine Queue

Hyper-V in Windows Server 2008 R2 associated the VMQ queue with each virtual machine statically. In Windows Server 2012 and newer versions, Hyper-V provides automatic configuration and tuning for VMQ queues. You can utilize this by associating VMQ with a processor dynamically, based on processor networking and CPU load. The number of processors that network processing uses can increase or decrease automatically, based on the network load. This allows the Hyper-V host to process more networking traffic and support higher network bandwidth. The ability to dynamically adjust the number of processor cores that process VMQ queues is called *Dynamic Virtual Machine Queue*.

Dynamic Virtual Machine Queue is enabled automatically in the virtual switch whenever an administrator enables VMQ on the virtual network adapter that is connected to the switch. There are only two ways to disable the VMQ feature. You can disable VMQ in the virtual network adapter Hardware Acceleration settings, or you can use the Windows PowerShell cmdlet **Set-VMNetworkAdapter**.

 **Note:** VMQ requires a physical network adapter that supports this feature. If you enable the VMQ feature on a virtual network adapter, but the Hyper-V host does not have a physical adapter that supports VMQ, you cannot use this feature.


NIC Teaming

NIC Teaming is a feature in Windows Server 2012 R2 that you can use to consolidate up to 32 physical network adapters. Then you can use those adapters as a single interface. This strategy provides both higher network throughput and redundancy. NIC Teaming is not a Hyper-V-specific feature. Because of this, all applications that are running at the system level on Windows Server 2012 R2 can benefit from NIC Teaming, including Hyper-V.

NIC Teaming is available to guest operating systems that are running inside virtual machines, regardless of whether NIC Teaming is used at the system level or not. This enables virtual machines with multiple virtual network adapters to team the adapters and still have connectivity, even when an adapter is disconnected or a virtual switch fails. This is especially important when using SR-IOV, because SR-IOV traffic bypasses the virtual switch and cannot benefit from NIC Teaming at the system level, whereas the Hyper-V virtual switch can.

To benefit from virtual machine NIC Teaming, you should create at least two external virtual switches and then connect virtual machine network adapters to them. You can configure physical network adapters that are connected to virtual switches to use SR-IOV, although this is not mandatory. If virtual machine network adapters are connected to SR-IOV-enabled virtual switches, the virtual machine will install virtual functions for them and will be able to use them in an NIC team. If one of the physical network adapters fails or is disconnected, the virtual machine continues to use the virtual functions of the remaining SR-IOV-enabled network adapters and has network connectivity. You will see the same result if virtual switches are connected to physical network adapters that are not SR-IOV-enabled. Another option is to use a combination of SR-IOV-enabled and nonenabled network adapters in the same virtual machine NIC team.

You can enable virtual machine NIC Teaming from the Advanced Properties settings page of the virtual network adapter. Alternatively, you can use the Windows PowerShell cmdlet **Set-VmNetworkAdapter**. Virtual machine NIC Teaming is not enabled by default. If you do not enable it, and if one of the physical network adapters stops working, the NIC team that is created in the guest operating system in the virtual machine loses connectivity.

 **Note:** Failover between network adapters in a virtual machine results in traffic being sent with the MAC address of the other network adapter. Because of this, each virtual network adapter that uses NIC Teaming must be set to allow MAC address spoofing, or must have the AllowTeaming=On parameter set by using the Windows PowerShell cmdlet **Set-VmNetworkAdapter**.

At the Hyper-V host level, NIC Teaming is not supported when physical network adapters are using SR-IOV or Remote Direct Memory Access (RDMA), because network traffic is delivered directly to the adapter, thereby bypassing the network stack and not allowing path redirection. When you configure NIC Teaming at the virtual machine level, physical network adapters that are connected to virtual switches can use SR-IOV.

Virtual Receive Side Scaling

Virtual Receive Side Scaling (vRSS) enables network adapters to balance the network processing load across the processor cores assigned to a virtual machine. Virtual Receive Side Scaling enables a virtual machine to process higher amounts of network traffic than it could process if only a single CPU core was responsible for processing traffic. You can implement Virtual Receive Side Scaling if you allocate virtual machine multiple cores through the advanced network settings. To use vRSS, the host's processor must support Receive Side Scaling (RSS) and the host's network adapters must support VMQ.

Demonstration: Configuring Virtual Networks

Demonstration Steps

Create a virtual network

1. Launch the Hyper-V Virtual Switch Manager and create a new virtual network switch named Classroom demo that allows the management operating system to share the network adapter and for which single-root I/O virtualization is enabled.
2. Set the virtual switch type to Internal network

Create a virtual network adapter

- Create a new virtual machine called **demo1**, and then connect the network adapter to the new **Classroom demo** network switch.

Review virtual network adapter settings

1. Review the network options for the newly created virtual machine demo1.
2. Delete the virtual machine **demo1**, and then remove the new **Classroom demo** switch.

Planning Networks for a Hyper-V Deployment

When you are configuring virtual networks, you must ensure that you provision virtual machines with adequate bandwidth. A bandwidth-intensive operation, such as a large file copy or website traffic spike, will affect performance on all virtual machines.

Best practices for configuring virtual networks include:

- **NIC Teaming.** You should deploy multiple network adapters to the Hyper-V host and then configure those adapters as part of a team. This ensures that network connectivity continues if the individual network adapters fail. Configure multiple teams that connect to different switches to ensure that connectivity continues if a hardware switch fails.
- **Bandwidth management.** You can use bandwidth management to allocate a minimum and a maximum bandwidth allocation for each virtual network adapter. You should configure bandwidth allocation to guarantee that each virtual machine has a minimum bandwidth allocation. This ensures that if another virtual machine hosted on the same Hyper-V server experiences a traffic spike, other virtual machines are able to communicate with the network normally.

When configuring virtual networks:

- Use NIC Teaming on the Hyper-V host to ensure connectivity to virtual machines if an adapter fails
- Enable bandwidth management to ensure that no single virtual machine is able to monopolize the network interface
- Use network adapters that support VMQ
- Use network virtualization to ensure that virtual machines keep their original IP addresses after migrating to a new host

- VMQ. You should provision the Hyper-V host with an adapter that supports VMQ, which uses hardware-packet filtering to deliver network traffic directly to the virtual machine. This improves performance because the host does not need to copy the packet from the host operating system to the virtual machine. When you do not configure virtual machines to support VMQ, the host operating system can become a bottleneck when it processes large amounts of network traffic.
- Network virtualization. Network virtualization is complicated to configure, but offers a distinct advantage compared to VLAN, because it is not necessary to configure VLANs on all of the switches that connect to the Hyper-V host. You can perform all necessary configurations when you need to isolate servers on the Hyper-V host, without having to involve the network team. However, if you are hosting large numbers of virtual machines and need to isolate them, use network virtualization rather than VLANs.
- Availability, protection, and recovery. When planning clusters or site-to-site replication and backups, consider where bottlenecks may occur. You can use 10 GB Ethernet cards in infrastructure sections that service many hosts at once, such as on scale-out file servers or iSCSI target hosts and backup servers. When possible, run a pilot to capture sample data and then calculate the throughput. Plan to monitor traffic flow between different sections of the virtual network, so that you can identify potential problems.

Managing Virtual Networks in VMM

Networking in VMM includes several enhancements that enable administrators to provision network resources efficiently for a virtualized environment. You can take advantage of the following networking capabilities:

- Create and define logical networks.
- Assign static IP addresses and static MAC addresses.
- Integrate load balancers.

Networking components include:

- Logical networks
- Network sites
- Static IP address pools
- MAC address pools
- Virtual IP templates
- Load balancer integration

Logical Networks

A logical network that you combine with one or more associated network sites is a user-defined, named grouping of IP subnets, VLANs, or IP subnet/VLAN pairs. These organize and simplify network assignments. Some possible logical network examples include BACKEND, FRONTEND, LAB, MANAGEMENT, and BACKUP. Logical networks represent an abstraction of the underlying physical network infrastructure, which enables you to model the network based on business needs and connectivity properties. After you create a logical network, you can use it to specify the network on which to deploy a host or a virtual machine (stand-alone or part of a service). Users can assign logical networks when they create a virtual machine and service, without understanding network details.

You can use logical networks to describe networks with different purposes, such as traffic isolation or provisioning networks for different types of service level agreements (SLAs). For example, for a tiered application, you may group IP subnets and VLANs that you use for the front-end web tier into a logical network named FRONTEND. You may choose to group backend servers into a logical network named BACKEND for the IP subnets and VLANs that you use. When a self-service user models the application as a service, he or she can choose the logical network for virtual machines in each service tier to which he or she wants to connect.

At least one logical network must exist if you want to deploy virtual machines and services. By default, when you add a Hyper-V host to VMM, VMM creates logical networks automatically. These logical

networks match the first Domain Name System (DNS) suffix label of the connection-specific DNS suffix on each host network adapter. To make a logical network available to a host, you must associate the logical network with a physical network adapter on the host, and make it available through an external virtual network or external virtual switch. You perform this association for each network adapter individually.

By default, when you add a Hyper-V host to VMM, if a physical network adapter on the host does not have an associated logical network, VMM automatically creates and associates a logical network that matches the first DNS suffix label of the connection-specific DNS suffix. For example, if the DNS suffix for the host network adapter is corp.contoso.com, VMM creates a logical network named corp. If you do not associate a virtual network with the network adapter, then if a job connects a virtual machine to a logical network that is associated with the physical network adapter, VMM creates an external virtual network automatically. Additionally, it associates it with the logical network. VMM does not create any network sites automatically, however. These default settings provide a solution to help you create and deploy virtual machines on your existing network.

Network Sites

When you create a logical network, you can create one or more associated network sites. A network site associates one or more subnets, VLANs, and subnet/VLAN pairs with a logical network, and enables you to define the host groups to which the network site is available. For example, if you have a Seattle host group and a New York host group, and you want to make the BACKEND logical network available to each, you can create two network sites for the BACKEND logical network. You can scope one network site to the Seattle host group (and any desired child host groups), and the other network site to the New York host group (and any desired child host groups). Additionally, you then would add the appropriate subnets and VLANs for each location.

IP Address Pools

If you associate one or more IP subnets with a network site, you can create an IP address pool. A *static* IP address pool enables VMM to assign static IP addresses to hosts, such as when you use VMM to convert a bare metal computer to a Hyper-V host or for Windows-based virtual machines that are running on any supported hypervisor platform. Static IP address pools enable your VMM administrator to manage IP addresses for the virtual environment. However, configuring static IP address pools is optional. You also can assign addresses automatically through DHCP if it is available on the network. If you use DHCP, you do not have to create IP address pools.

MAC Address Pools

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

Virtual IP Templates

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

Load Balancer Integration

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

Logical Switches

Logical switches allow you to apply a single configuration to multiple hosts, and you configure them to use native port profiles, port classification, and virtual-switch extensions. The type of switch extensions supported are:

- **Monitoring.** Monitoring extensions monitor and report network traffic but cannot modify packets.
- **Capturing.** You can use capturing extensions to inspect and sample traffic but not to change packets.
- **Filtering.** You can use filtering extensions to block, modify, or defragment packets. You can also use them to block ports.
- **Forwarding.** You can use forwarding extensions to direct traffic by defining destinations. They can also capture and filter traffic. To avoid conflicts online, one forwarding extension can be active on a logical switch.
- **Virtual switch extension manager.** A virtual switch extension manager makes it possible to use a vendor network-management console and VMM together. To do this, you need to install the vendor's provider software on the VMM server.

Native Port Profiles

You can use native port profiles, also called Hyper-V port profiles, to configure uplink adapters that must be available on the physical network adapters to which a switch connects. You can assign these to host groups, and then enable them to support Windows network-virtualization. You also can use native port profiles to configure virtual adapters for enabling offload settings, such as VMQ, IPsec task offloading, and SR-IOV. Virtual network adapter port profiles allow you to reuse the same settings across multiple switches, which simplify the deployment of your virtual environments.

Additionally, you can specify minimum and maximum bandwidth settings and relative bandwidth weights to define how much bandwidth the virtual network adapter can use in relation to other virtual network adapters that connect to the same switch.

The following default profiles have already been created in VMM:

- SR-IOV Profile
- NLB NIC Profile
- Low, medium, and high bandwidth adapters
- Host management
- Live Migration
- Cluster
- iSCSI
- Default

Each of these has been configured with varying offload, security, and bandwidth settings.

Port Classifications

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

The following default port classifications have already been created in VMM:

- SR-IOV
- Network load balancing
- Live migration workload
- Host Cluster Workload
- Low, medium and high bandwidth
- iSCSI workload

Gateways

In VMM, a gateway allows network traffic in and out of a virtual machine network that is using network virtualization. You can configure this for local network routing, which routes traffic between the virtual machine network and the physical network, or you can configure it for remote network routing, which first creates a virtual private network (VPN) connection with another endpoint of a site-to-site VPN. It then routes in and out of the virtual machine network through the VPN tunnel. The remote option is most relevant for hosting providers.

Demonstration: Configuring Virtual Network Components in VMM

In this demonstration, you will see how you can use VMM to create and configure the following network components:

- Logical networks and logical network IP pools
- Native port profiles, uplinks, and virtual adapters
- Logical switches

In addition, you will see where you assign logical switches to Hyper-V hosts.

Demonstration Steps

Create a logical network

1. In the Fabric workspace of the Virtual Machine Manager console, create a logical network named **Classroom1**, enabling the **Allow new VM networks created on this logical network to use network virtualization** option.
2. Configure this logical network so that the All Hosts group can use this network site, and so that VLAN ID 3 and IP subnet 192.168.3.0/24 are associated.

Create a logical network IP Pool

- In the **Fabric** workspace, create an IP Pool named **Classroom1 IP Pool** that uses the Classroom1 logical network, the Classroom1_0 network site, and the 192.168.3.0/24 subnet.

Create an uplink native port profile

1. Create a Hyper-V port profile named
2. **Classroom1 Uplink**. Set it to use the Classroom 1_0 network site, and then enable Hyper-V network virtualization.

Create a logical switch

1. Create a new logical switch with the name **Classroom switch1** that uses the default extensions and uses Classroom1 Uplink as the port profile.
2. On the **Virtual Port**, add a virtual port that uses the **High Bandwidth** port profile classification set to the native virtual network adapter port profile that uses the **High Bandwidth Adapter** port profile.

Add a logical switch to a host server

3. From the Fabric workspace, create a new virtual switch of the logical switch type on **LON-HOST1**. Click **OK** if presented with a warning.
4. On the **Properties** page, select **Hardware**, and expand **Network adapters**. Select your physical network adapter, and note that you can select or clear the adapter for virtual machine placement and management use.
5. Select the logical network, and then on the right under **Logical network connectivity**, verify that you can assign the logical networks and IP subnets.

Lesson 4

Planning and Implementing Network Virtualization

Network virtualization allows you to completely isolate virtual machines running on the same host. Network virtualization simplifies support of multitenant networking by removing the requirement to implement independent network infrastructure to ensure virtual machine separation. In this lesson, you will learn about multitenant networking scenarios. In addition, you will learn how you can use network virtualization through Hyper-V and VMM to ensure virtual machine isolation.

Lesson Objectives

In this lesson, you will learn how to:

- Understand multitenant networking scenarios.
- Understand options for implementing multitenant networking.
- Describe how network virtualization works.
- Describe VMM network virtualization components.
- Understand Windows Server Gateway.

Multitenant Virtual Networking Scenarios

Virtualization provides many benefits, including consolidation, better hardware utilization, and virtual machine separation from the physical server hardware. As a result, many companies are virtualizing most of their server load.

Organizations now have the ability to host virtual machines from different departments or even different companies in the same data center. Therefore, it is important to be able to separate and isolate each organization's virtual machines. Multitenant hosting is the process of hosting workloads owned by disparate groups of people.

When you isolate virtual machine workloads running on the same virtual machine host, you can:

- Remove virtual machine ownership from the considerations around virtual machine placement
- Group virtual machines on the same host so that they can use the same IP address space without conflict
- Allow more efficient use of virtualization host resources

One of the basic requirements of multitenant hosting is to be able to isolate virtual machines that run on the same physical hardware. Isolation ensures that one virtual machine is unable to communicate directly with another. This makes it impossible for the virtual machine to determine what other workloads are co-located on the same virtualization host. Until recently, there was no easy, inexpensive, and scalable solution for separating or isolating the network traffic that different tenants generate on the same network infrastructure. You can implement physical network separation by deploying separate network hardware, but this option is neither scalable nor inexpensive.

When you can isolate virtual machines so that they run on the same physical hardware without any ability to communicate directly with each other, you have more flexibility in managing your overall virtual machine hosting capacity. By isolating virtual machines, you can move those virtual machines across hosts on the data center according to resource availability. You do not have to restrict movement depending on ownership rules, such as: a virtual machine owned by customer A can never be placed on the same virtual host as a virtual machine owned by customer B. Isolating workloads reduces the complexity of multitenant hosting by separating virtual machine ownership and virtual machine placement considerations.

Separation is not just for Internet service providers (ISPs) or hosting providers that host workloads for external customers. You can use it for large organizations with internal customers. For example, the IT department at your organization might charge each department based on the metered resource utilization of their virtual machines. This is similar to how public cloud providers, such as Windows Azure, charge based on the utilization of computer, storage, and bandwidth resources. When you segment and isolate those workloads, you can charge each department based on utilization.

Network isolation is useful during mergers and acquisitions. Many organizations use the same internal private IP address space. Typically, when one organization merges with or acquires another that uses the same internal IP address space, a networking team performs the complex task of address reassignment. With the ability to isolate virtual machines, groups of virtual machines on the same virtualization host are able to use the same IP address space without requiring address reassignment.

When you isolate virtual machines that run on the same physical hardware, you can provide elastic capacity to meet the needs of all tenants from one hardware pool, rather than requiring separate pools of hardware to meet the elastic requirements of different tenants. For example, your organization may host virtual machines for 10 separate companies or departments, which have 10 different sets of elastic requirements. You could approach this by using two strategies:

- You have 10 separate sets of virtual machine hosts, one set for each tenant workload. You add and remove virtual machine hosts from each set as capacity requirements increase and decrease.
- You have one set of virtual machine hosts with each tenant workload segmented from the other through isolation. You add and remove virtual machine hosts from the total set only as overall capacity requirements increase and decrease.

With the second strategy, decreased capacity requirements in one group of tenants might provide you with the resources to address the increased capacity requirements of another.

Options for Implementing Multitenant Networking

You can use different solutions to provide network isolation in a multitenant environment, such as:

- VLANs. Most organizations use this solution to support address space reuse and multitenant isolation. A VLAN uses an additional header that contains a VLAN ID. It relies on switches to enforce isolation of network traffic between computers that are connected on the same network but use different VLAN IDs. One of the drawbacks of VLAN is that it provides limited scalability. Theoretically, because VLAN ID uses only 12 bits, you can have a maximum of 4,094 different VLANs on the same infrastructure. However, many switches can support much fewer than 4,094 VLANs. A second drawback is that VLANs cannot span multiple logical subnets. This limits the number of computers in a single VLAN and restricts the placement of virtual machines based on physical location.

Multiple isolated networks on the same infrastructure
The Hyper-V virtual switch supports three solutions:

VLANs
Private VLANs
Port ACLs

- When you use VLANs, consider the following:
 - Limited scalability (maximum of 4094 VLANs)
 - VLANs cannot span multiple subnets
 - Challenge to reconfigure when adding or moving virtual machine

Although you can enhance or stretch VLANs across physical locations, a stretched VLAN must be on the same subnet. You should configure switches and routers to support VLANs. You need to reconfigure them whenever virtual machines or isolation boundaries move in the dynamic data center. You can automate this to a certain extent, but it increases the risk of an inadvertent network outage due to incorrectly performed reconfiguration.

- Private VLANs. You can use private VLANs to avoid some of the VLAN scalability limitations. You implement private VLANs in a similar way to VLANs, but you can use private VLANs to divide a VLAN into a number of separate and isolated subnetworks, which you can then assign to tenants. Private VLANs consist of a primary and secondary VLAN pair, which share the parent VLAN's IP subnet. Although computers that are connected to different private VLANs still belong to the same IP subnet, they require a router to communicate with each other, and with resources on any other network.

When you use private VLANs, you can assign a large number of tenants to the same primary VLAN and have isolated secondary VLAN IDs. For example, if you have 4,000 tenants and you cannot use private VLANs, you need 4,000 VLANs to provide isolation. However, if you use private VLANs, you can use one primary VLAN only, and assign each tenant a different secondary VLAN. When using such a configuration, you need a single VLAN ID only, instead of 4,000.

- Port access control lists (ACLs). You can use port ACLs to configure network traffic filtering based on MAC addresses, IP addresses, or IP ranges. By using port ACLs, you can configure virtual network isolation by creating two lists: one list contains addresses of computers with which a virtual switch port can communicate, and the second list contains addresses of computers with which a virtual switch port cannot communicate or share data.

When you add a new virtual machine or move an existing virtual machine, you must manage and update these two lists need. This can be a challenging and error-prone process. Technically, it is possible to provide multitenancy isolation by using only port ACLs. However, you do not typically use the port ACLs feature for this purpose, but for ensuring that virtual machines are not pretending to have different IP or MAC addresses than those you have assigned to them.

The Hyper-V virtual switch supports all three solutions—VLANs, private VLANs, and port ACLs. However, the virtual switch also supports network virtualization, which is the best solution for providing multitenant networking.

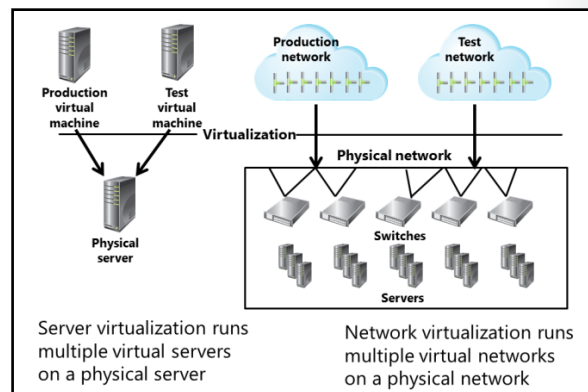
How Network Virtualization Works

You can use network virtualization to isolate virtual machines from different organizations, even if they share the same Hyper-V host. For example, you might be providing an infrastructure as a service (IaaS) to different businesses, or you might want deploy a copy of certain virtual machines without creating conflicting IP addresses. You can use network virtualization to assign these virtual machines to separate virtual networks that are running on the same physical or logical network but are isolated from each other. Network virtualization offers you the advantage of configuring all network isolation on the Hyper-V host without VLANs. Furthermore, by using gateways or VPN extensions, you can extend virtualized networks for isolated communication between hosts.

You can continue to run isolated machines with VLANs. However, you must configure the physical switches with the appropriate VLAN IDs.

When you configure network virtualization, each guest virtual machine has two IP addresses:

- Customer IP address. The customer assigns this IP address to the virtual machine. You can configure this IP address so that communication with the customer's internal network can occur even if the virtual machine is hosted on a Hyper-V server that connects to a separate public IP network. You can see the customer IP address by using the **ipconfig** command on the virtual machine.



- Provider IP address. The hosting provider assigns this IP address, which is visible to the hosting provider and to other hosts on the physical network. This IP address is not visible from within the virtual machine.

You can use network virtualization to host multiple machines that use the same customer address ranges, such as 10.x.x.x, on the same Hyper-V host. When you do this, the hosting provider assigns the virtual machines different IP addresses, though these addresses will not be visible from the virtual machines.

You can manage network virtualization by using Windows PowerShell cmdlets. All network virtualization cmdlets are in the NetWNV Windows PowerShell module. Tenants gain access to virtual machines that take advantage of network virtualization through routing and remote access. They make a tunneled connection from their network to the virtualized network on the Hyper-V server.

Hyper-V Network Virtualization supports using Network Virtualization for Generic Route Encapsulation as the method of virtualizing IP addresses. This method encapsulates the virtual machine's packets inside packets that have the externally used source and destination information. The outer packet's header also stores the Virtual Subnet ID used with the virtualized network, which allows hosts to identify the virtual machine associated with a specific packet.



Reference Links: For a complete overview of network virtualization, go to <http://go.microsoft.com/fwlink/?LinkID=285279>.

For detail about network virtualization in System Center 2012 R2 Virtual Machine Manager, review the following series of product team blog posts:

<http://blogs.technet.com/b/scvmm/archive/2013/11/25/adopting-network-virtualization-part-i.aspx>

Implementing Network Virtualization

In larger VMM environments, you will need to administer a greater number of logical networks, virtual machine networks, and virtual network components. If you have multiple administrators, the potential for complexity and error or increases further.



Best Practice: In most sections of the VMM console, you can filter the view by entering text in the search field. Keep this feature in mind and apply a good naming convention to all your virtual network components. This will help you and other administrators when you are working with and or troubleshooting virtual networking. This applies to everything you can label in VMM.

- Before deleting a virtual machine network, confirm that there are no dependent resources

When you use VMM to manage virtual machine networks, you can:

- Quickly discover which virtual machines connect to which networks by using the built-in VMM network diagrams
- Delegate access to virtual machine networks by setting an owner for a virtual machine network

You should be aware of a few considerations before you start working with virtual machine networks in VMM. As a first step, you should plan your network and document the proposed configurations. You will need to determine if you should implement isolation. Then you must create the underlying logical network components.

After you have created your prerequisite logical network, perform the following steps to create a virtual machine network in the VMM console:

1. Open the VMM console, click the **VMs and Services** workspace, and then on the ribbon, click **Create VM Network**.
2. On the **Name** page, type the name and description of your virtual machine network, click the drop-down list box, select the logical network, and then click **Next**.
3. On the **Isolation** page, click to select either **Isolate using Hyper-V network-virtualization** or **No isolation**, choose between **IPv4** and **IPv6** for your virtual machine network and logical network, and then click **Next**.
4. On the **VM Subnets** page, click **Add**, and in the **Name** field, type the name for your virtual machine subnet. In the **Subnet** field, type the IP address and mask for your subnet. If necessary, add and remove further subnets, and then click **Next**.
5. On the **Connectivity** page, choose the setting for connecting directly to an additional logical network, and specify whether that connection will use network address translation (NAT). If you have not added a gateway, no option will be available. Review the message, and then click **Next**.
6. On the **Summary** page, review the summary, and then click **Finish**.
7. Close the Jobs window.

In a large host or environment, you may need to discover quickly which virtual machines connect to which networks. Rather than investigate each virtual machine individually, you can investigate in Windows by using the built-in VMM network diagrams.

You can review hosts and virtual machine network topology by performing the following steps:

1. Open the VMM console, and then click the **Fabric** workspace.
2. In the Fabric navigation pane, click to expand the host group containing your hosts. In the main section of the console, right-click the host that you want to review, and then click **View Networking**.
3. On the left, you can choose the hosts, host groups, and clouds that you want to include in the diagram. On the ribbon, you can choose to view the following diagrams:
 - VM Networks
 - Host Networks
 - Host/VM Networks
 - Network Topology

To delegate access to virtual machine networks, you set an owner for a virtual machine network and delegate access to other administrators and self-service users. You can configure access by performing the following steps:

1. Open the VMM console, click the **VMs and Services** workspace, and then on the ribbon, click **Properties**.
2. On the left, click **Access**. You can select an owner and delegate access to the virtual machine network.

If you want to delete a virtual machine network, first you must confirm that there are no dependent resources. You can review dependent resources by following these steps:

1. Open the VMM console, and then click the **VMs and Services** workspace.
2. In the VMs and Services navigation pane, click **VM Networks** on the right, click to highlight a virtual machine network, and then on the ribbon, click **View Dependent Resources**.

3. Review the Names and Type of resources, and then click **OK**.
4. To delete other VMM resources that may have dependent resources, you can right-click them. If they have dependent resources, the dependent resource option will display. You can click this option to display those dependencies.

Demonstration: Configuring Network Virtualization in VMM

In this demonstration, you will learn how to enable and configure network virtualization in Windows Server 2012 R2 and VMM. For PCs with only one network adapter, you can follow many of the steps, but you cannot assign the logical switch to the host or deploy a virtual machine that will use network virtualization.

Demonstration Steps

Create the Test VM network

- On LON-VMM1, open the Virtual Machine Manager console, and click the **VMs and Services** workspace. Create a VM network named **Classroom1_Test**, select network isolation, create a subnet named **Test Network**, and then assign the subnet **192.168.3.0/24**.

Create the Production VM network

- In the **VMs and Service** workspace, create another VM Network named **Classroom1_Production**, select network isolation, create a subnet named **Production Network**, and then assign the subnet **192.168.3.0/24**.

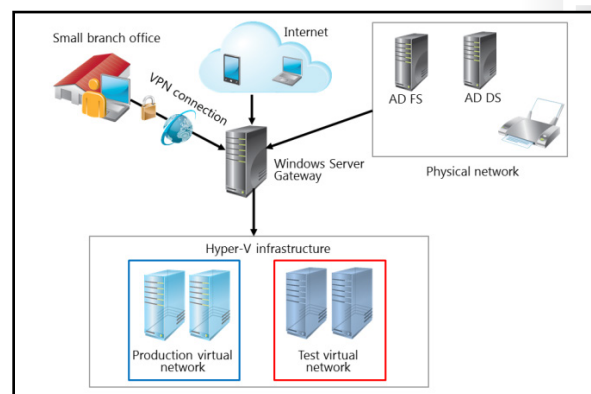
Create the VM IP pools

1. Create an IP pool for the Production VM network named **Production VM Network IP Pool**, and then accept the defaults. Note that the first address in the pool is reserved.
2. Create an IP pool for the Test VM network, named **Test VM Network IP Pool**.

Windows Server Gateway

When you use the Hyper-V virtual switch to implement network virtualization, the switch operates as a router between different Hyper-V hosts in the same infrastructure. The network virtualization policies define how packets are routed from one host to another.


However, a virtual switch cannot route to networks outside the Hyper-V server infrastructure when you use network virtualization. If you do not use network virtualization, you connect the virtual machine to an external switch, and the virtual machine connects to the same networks as the host machine. However, in a network virtualization scenario, you may have multiple virtual machines that share the same IP address running on a Hyper-V host. You want to be able to move the virtual machine to any host in the network without disrupting network connectivity. You also need to be able to connect the virtualized networks to the Internet by using a mechanism that is multitenant-aware, so that traffic to external networks is routed correctly to the internal addresses that the virtual machines use.



Windows Server 2012 R2 provides Windows Server Gateway to address this scenario. *Windows Server Gateway* is a virtual machine–based software router that allows you to route network traffic between the virtual networks on the Hyper-V hosts and physical networks. This enables the virtual machines to connect to other resources on the internal network and to external networks such as the Internet.

You can implement Windows Server Gateway in three different configurations:

- Multitenant-aware virtual private network (VPN) gateway. In this configuration, you configure Windows Server Gateway as a VPN gateway that is aware of the virtual networks deployed on the Hyper-V hosts. By deploying the Windows Server Gateway with this configuration, you can connect to the Windows Server Gateway by using a site-to-site VPN from a remote location. Alternatively, you can configure individual users with VPN access to the Windows Server Gateway. The Windows Server Gateway operates similarly to any other VPN gateway; it allows the remote users to connect directly to the virtual networks on the Hyper-V servers. The main difference is that the Windows Server Gateway is multitenant-aware, so you can have multiple virtual networks with overlapping address spaces located on the same virtual infrastructure. This configuration is useful for organizations that have multiple locations, or multiple business groups that share the same address spaces and must be able to route traffic to the virtual networks. Hosting providers can also use this configuration to provide remote clients with direct network access between their on-premise network and the hosted networks.
- Multitenant-aware network address translation (NAT) gateway for Internet access. In this configuration, Windows Server Gateway provides access to the Internet for virtual machines on virtual networks. The Windows Server Gateway is configured as an NAT device; it translates addresses that can connect to the Internet to addresses used on the virtual networks. In this configuration, Windows Server Gateway is multitenant-aware, so that all virtual networks behind the Windows Server Gateway can connect to the Internet, even if they use overlapping address spaces.
- Forwarding gateway for internal physical network access. In this configuration, Windows Server Gateway provides access to internal network resources that are located on physical networks. For example, an organization may have some servers that are still deployed on physical hosts. When you configure it as a forwarding gateway, Windows Server Gateway enables computers on the virtual networks to connect to those physical hosts.

 Windows Server Gateway is a Microsoft implementation of a multitenant-aware gateway. Third-party vendors have developed similar gateways. For more details about Windows Server Gateway, go to the following link: <http://technet.microsoft.com/en-us/library/dn313101.aspx>

You can configure Windows Server Gateway by deploying appropriate Windows Server 2012 R2 roles and by configuring the network settings by using Windows PowerShell. To implement Windows Server Gateway, follow these high-level steps:

1. Verify that your Hyper-V deployment meets the requirements for the Windows Server Gateway deployment. Although you can deploy a Windows Server Gateway on a host with single network adapter, we recommend that you configure multiple network adapters on the host. You must configure multiple virtual network adapters on the Windows Server Gateway virtual machine. As a best practice, configure the physical and virtual network adapter names to match the intended use for each network.
2. Install the Remote Access role on the Windows Server Gateway virtual machine, including the Direct Access and VPN (RAS) and Routing role service. Install the required management tools.
3. On the Hyper-V host running the Windows Server Gateway virtual machine, do the following:
4. Enable the multitenancy mode on the virtual machine network adapter by using the **Set-VMNetworkAdapterIsolation** cmdlet with the **-IsolationMode** parameter.

5. Map the tenant's routing domains and virtual subnets by using the **Add-VmNetworkAdapterRoutingDomainMapping** parameter.
6. Configure the network virtualization settings by using the **New-NetVirtualizationProviderAddress**, **New-NetVirtualizationLookupRecord**, and **New-NetVirtualizationCustomerRoute** cmdlets.
7. On Windows Server Gateway, configure the IP addresses and network routes for each tenant network.
8. On the Hyper-V hosts running the tenant virtual machines, configure the network virtualization settings by using the **New-NetVirtualizationProviderAddress**, **New-NetVirtualizationLookupRecord**, and **New-NetVirtualizationCustomerRoute** cmdlets.

When deploying a gateway in VMM, configure a gateway in the Network Service section of the Fabric workspace. The gateway connects to remote networks using a VPN tunnel. To add a gateway, you must first install its provider software. You can review the list of installed providers by using the following procedure:

1. Open the VMM console.
2. Click the **Settings** workspace, and then in the Settings pane, click **Configuration Providers**. The lists of providers displays, along with information such as Type, Version, Publisher, Manufacturer, and Model. The default providers in VMM are:
 - Microsoft IP Address Management Provider
 - Microsoft Network Load Balancing (NLB)
 - Microsoft Standards-Based Network Switch Provider
 - Microsoft Windows Server Gateway Provider

The default installation directory for providers is:

- C:\Programs Files\Microsoft System Center 2012 R2\Virtual Machine Manager\Bin\Configuration Providers

Confirm that the necessary provider software for the gateway device has been installed and is listed.

For more information about gateway prerequisites, and to review the setup steps, refer to:



Configuring VM Networks and Gateways in VMM

<http://technet.microsoft.com/en-us/library/jj721575.aspx>



How to Add a Gateway in VMM in System Center 2012 R2

<http://technet.microsoft.com/en-us/library/dn249416.aspx>

Lab: Planning and Implementing Virtualization Networks and Storage

Scenario

After designing and implementing the host server and the virtualization infrastructure's management layers, the next step is to plan and implement the storage and network layers.

Because many virtual machines will be sharing the same storage and network infrastructure, it is critical that these layers be highly available. At the same time, A. Datum Corporation does not have the budget to overprovision its virtual environment, so your design must provide sufficient, not excessive, capacity.

The storage administrators in London are designing the storage infrastructure using the Fibre Channel SAN. The network administrators in London are responsible for designing the network configuration. However, you are responsible for creating the storage design for the Hyper-V deployment in Toronto. You are also responsible for creating the network design in Toronto.

Objectives

- Plan a storage infrastructure.
- Plan a network infrastructure.
- Implement iSCSI storage for the virtual machines.
- Configure network components for the virtual machine deployment.
- Configure network virtualization.

Lab Setup

Estimated Time: 60 minutes

Virtual machines: 20414C-LON-HOST1, 20414C-LON-HOST2,
20414C-LON-DC1, 20414C-LON-VMM1, 20414C-LON-SVR1

User Name: **Adatum\Administrator**

Password: **Pa\$\$w0rd**

Lab Setup

To perform this lab, you must ensure you have completed the lab in Module 2. You will use the virtual machine environment that is available after completing that lab. Before you begin this lab, you must complete the following steps:

1. You should have started the following computers and virtual machines after completing the lab in Module 2:
 - 20414C-LON-HOST1
 - 20414C-LON-HOST2
 - 20414C-LON-DC1
 - 20414C-LON-VMM1
2. Do not start **20414C-LON-SVR1** until directed to do so.
3. If necessary due to host resources, you can shut down 20414C-TOR-SVR1.

Exercise 1: Planning a Storage Infrastructure for Virtualization

Scenario

In Toronto, A. Datum is planning to use an iSCSI-based storage infrastructure for the virtual machines. The server operations team has provided the following information about the server deployments that it is planning for the Toronto branch office. They will deploy two domain controllers, which are also DNS servers, as virtual machines in the Toronto data center. The servers require only a standard operating system partition. Specifically, the server operations team will deploy:

- Four web servers as virtual machines in the Toronto data center. The web servers require a standard operating system partition and approximately 30 GB of storage each.
- Two file servers as virtual machines in the Toronto data center. The file servers require a standard operating system partition and approximately 1 terabyte of storage each.
- Three database servers as virtual machines in the Toronto data center. The database servers require a standard operating system partition and approximately 900 GB of storage each.
- Four Exchange Server 2010 servers as virtual machines in the Toronto data center. The Exchange servers require a standard operating system partition and about 3.5 terabytes of storage each.

You must plan the storage infrastructure for the iSCSI storage deployment at Toronto. A. Datum is planning to deploy two Windows Server 2012 servers with iSCSI targets. You need to create a storage design that provides sufficient storage. In addition, your design must provide availability based on the company requirements, which the following table details.

Server role	Number deployed	Operating system storage	Data storage (maximum)	Suggested performance requirements
Domain Controllers/DNS servers	2	40 GB	30 GB	Medium RAID 1
Web servers	4	40 GB	30 GB	Medium RAID 1 or 5
File servers	2	40 GB	1000 GB	Medium RAID 5
Database servers	3	40 GB	900 GB	High Performance RAID 1 or 10
Exchange servers	4	40 GB	3500 GB	Medium RAID 5

A. Datum Toronto Storage Strategy	
Document Reference Number: BS0905/1	
Document Author	Brad Sutton
Date	5 th September
Requirements Overview	
<ul style="list-style-type: none"> • To plan a storage strategy to support the following objectives: • Provide sufficient storage to virtualize the servers listed. • Provide high availability where possible. 	

A. Datum Toronto Storage Strategy

- Identify risks and bottlenecks.

Additional Information

- Provide an alternative lower cost solution that offers the same availability.

Tasks & Questions

1. The servers provided have a single 10 GB onboard network adapter with all the latest iSCSI features. For redundancy, you want to purchase more. There are three PCI Express slots available. How many should you buy?
2. The network team advises you that there will be a significant cost and delay should you wish to implement a separate iSCSI network. The current network is relatively new and has some 10 GB capability, though its number of connections is limited. You must advise the team on the number of 1 or 10 GB connections that will be required.
3. You have considered using another host server as a maintenance host, but your budget does not permit this. What is an alternative solution?
4. You identified data protection as both a risk and potential bottleneck. What could the cause be and how can you resolve this?
5. What Windows Server role or feature can you enable to help keep your virtual machines running when a network storage component, (such as a network adapter or network switch) fails?

The main tasks for this exercise are as follows:

1. Read the supporting documentation.
2. Update the proposal document with your planned course of action.
3. Examine the suggested proposals in the Lab Answer Key.
4. Discuss your proposed solution with the class, as guided by your instructor.

► **Task 1: Read the supporting documentation**

Read the documentation that the student handbook provides.

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

Answer the questions in the proposals section of the A. Datum Toronto storage strategy document.

Question: The servers provided have a single 10 GB onboard network adapter with all the latest Internet Small Computer System Interface (iSCSI) features. For redundancy, you want to purchase more. There are three PCI Express slots available. How many should you buy?

Question: The Network team advises you that there will be a significant cost and delay should you wish to implement a separate iSCSI network. The current network is relatively new and has some 10 GB capability, though its number of connections is limited. You must advise the team on the number of 10 GB connections required. How many are required?

Question: You have considered using another host server as a maintenance host, but your budget does not permit this. What is an alternative solution?

Question: You identified data protection as both a risk and potential bottleneck. What do you think is the cause, and how can you resolve the issue?

Question: What Windows Server feature will help keep your virtual machines running when a network storage component (such as a network adapter or network switch) fails?

► **Task 3: Examine the suggested proposals in the Lab Answer Key**

Compare your proposals with those in the Lab Answer Key.

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

Be prepared to discuss your proposals with the class.

Exercise 2: Planning a Network Infrastructure for Virtualization

Scenario

You must plan the virtual network configuration for the VMM configuration at the Toronto data center. Clients at Toronto area branch offices interact with the servers hosted virtually at this data center. The network team at A. Datum has provided the following information for the network configuration at the Toronto data center.

VLAN name/identifier	Purpose	IP addresses
VLAN3/3	Virtualization servers	172.16.1.0/24
VLAN4/4	Infrastructure servers (DHCP/DNS/Domain Controllers)	172.16.2.0/24
VLAN5/5	Application servers	172.16.3.0/24
VLAN6/6	Backup network	172.16.4.0/24
VLAN7/7	iSCSI network	172.16.5.0/24

Recently, A. Datum has acquired a smaller company named Wingtip Toys. Wingtip Toys uses a highly virtualized server infrastructure with the following configuration.

VLAN name/identifier	Purpose	IP addresses
VLAN3/3	iSCSI network	172.16.1.0/24
VLAN4/4	Backup network	172.16.2.0/24
VLAN5/5	Application servers	172.16.3.0/24
VLAN6/6	Infrastructure servers (DHCP/DNS/Domain Controllers)	172.16.4.0/24
VLAN7/7	Virtualization servers	172.16.5.0/24

As with A. Datum, workers in Wingtip Toys branch offices access resources in their data center from client computers.

After the acquisition, workers at the A. Datum and Wingtip Toys Toronto locations will remain in their separate branch offices. However, you will centralize all server infrastructure by hosting both the A. Datum and Wingtip Toys virtual machines on the Hyper-V servers in the A. Datum data center. Once you have completed the migration from the Wingtip Toys data center and decommissioned that data center, clients on the A. Datum and Wingtip Toys branch office networks should be able to communicate with the A. Datum and Wingtip Toys servers in the A. Datum data center.

A. Datum and Wingtip Toys Toronto Network Virtualization Strategy	
Document Reference Number: BS0906/1	
Document Author	Brad Sutton
Date	8 th September
<p>Requirements Overview</p> <p>Plan a network virtualization strategy to meet the following objectives:</p> <ul style="list-style-type: none"> • Allow multiple Hyper-V virtual machines with the same IP address on the same host. • Allow traffic to pass between isolated virtual machines located on different Hyper-V hosts. 	
<p>Proposals</p> <ol style="list-style-type: none"> 1. Which of the listed networks might you need to virtualize to support the objectives? 2. Where and how can you can you configure virtualized networks? 3. What are the steps to configure the virtualized networks? 4. Are there other ways to host these overlapping virtual machines? 5. How can you ensure that clients in the A. Datum and Wingtip Toys branch offices are able to access the correct servers? 	

The main tasks for this exercise are as follows:

1. Read the supporting documentation.
2. Update the proposal document with your planned course of action.
3. Examine the suggested proposals in the Lab Answer Key.
4. Discuss your proposed solution with the class, as guided by your instructor.

► **Task 1: Read the supporting documentation**

Read the documentation and scenario provided.

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

Answer the questions in the proposals section of the A. Datum Toronto Network Virtualization Strategy document.

Question: Which of the listed networks might you need to virtualize to support the objectives?

Question: Where and how can you configure virtualized networks?

Question: What are the steps to configure the virtualized networks?

Question: Are there other ways to host these overlapping virtual machines?

Question: How can you ensure that clients in the A. Datum and Wingtip Toys branch offices are able to access the correct servers?

► **Task 3: Examine the suggested proposals in the Lab Answer Key**

Compare your proposals with those in the Lab Answer Key.

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

Be prepared to discuss your proposals with the class.

Exercise 3: Implementing a Storage Infrastructure for Virtualization

Scenario

A. Datum has decided to implement virtual machine storage in the Toronto data center by using an iSCSI storage deployment. You must configure the iSCSI targets and iSCSI initiators to enable this deployment.

The main tasks for this exercise are as follows:

1. Configure iSCSI targets for the virtual machine deployment.
2. Configure iSCSI initiators.

► **Task 1: Configure iSCSI targets for the virtual machine deployment**

Add virtual disks to LON-SVR1

1. On LON-HOST1, create the folder **C:\StoragePool**.
2. In Hyper-V Manager, edit the settings of 20414C-LON-SVR1, and add and attach three 50 GB dynamically expanding VHDX format virtual hard disks named **iSCSI1.vhdx**, **iSCSI2.vhdx**, and **iSCSI3.vhdx** to the SCSI controller. Configure these disks to be stored in the **C:\StoragePool** folder.
3. Start and connect to LON-SVR1 and sign in as **Adatum\Administrator** with the password **Pa\$\$wOrd**.

Add the iSCSI Target Server Role Service

- On LON-SVR1, in Server Manager, add the **iSCSI Target Server** role (located under File and iSCSI Services).

Create a storage pool

1. In the Storage Pool section of Server Manager, create a new storage pool named **VMPool** that uses the three new disks.
2. Run the New Virtual Disk Wizard and create a new virtual disk named **VMStorage** by using **Parity** and **Thin** provisioning that is **100 GB** in size.
3. Run the New Volume Wizard and set the Size of the new volume using the VMStorage Virtual Disk at **(99.9) GB**, set the drive letter as **F:**, and the Volume Label as **VMStorage**.
4. In the iSCSI Virtual Disks pane, create a new iSCSI virtual disk. Set the storage location to **F:** and use **LONHOST1-iSCSIDisk1** for the **iSCSI virtual disk name**. Make the size **90 GB and dynamically expanding**, and then create a new named **LON-HOST1**.
5. On the **Specify access servers** page, browse to LON-HOST1.
6. Complete the wizard, leaving the defaults.
7. Create a second iSCSI virtual disk named **iSCSIDisk2** on **C:** that is **5 GB** and is assigned **LON-HOST1** as the target.

► Task 2: Configure iSCSI initiators

Configure iSCSI initiators

1. On LON-HOST1, start the Microsoft iSCSI service.
2. Use the iSCSI initiator's Quick Connect function to connect to **172.16.0.12**.
3. On LON-HOST1, open **Disk Management**, bring the 90 GB disk online, initialize the disk, and create a new simple volume. Leave the default size, assign the letter **V**, and then add a label **VMStorage**.
4. Close Disk Management.

Results: After completing this exercise, you will have configured iSCSI targets and iSCSI initiators and implemented iSCSI storage.

Exercise 4: Implementing a Network Infrastructure for Virtualization

Scenario

Now you must configure the network infrastructure for the virtualization deployment.

The main tasks for this exercise are as follows:

1. Configure logical networks that your design requires.
2. Configure network virtualization.
3. Assign virtual machines to VM networks.
4. To prepare for the next module.

► Task 1: Configure logical networks that your design requires

Create the logical network

1. On LON-VMM1 open the **Virtual Machine Manager Console** and in the Fabric workspace, create a logical network named **Toronto Production Network**.
2. Enter a description of **Adatum Toronto – Production Logical Network**, and then enable **Allow new VM networks created on this logical network to use network virtualization**.
3. Add **All Hosts** in the Host groups section, and then add the following row VLAN: 0, IP subnet: 172.16.3.0/24.

Create an IP pool

1. In the Fabric workspace, create and assign an IP Pool named **Toronto Apps** to the **Toronto Production Network** network site. On the **Network site** page, click **Use an existing network site**, select **Toronto Production Network_0**, click the **IP subnet** drop-down list box, select **172.16.3.0/24**, and then click **Next**. Reserve the 172.16.3.100-172.16.3.120 range for load balancer virtual IPs (VIPs), and then assign **172.16.3.1** for the gateway. In the description field, type **Toronto Production Application IP Pool**.
2. **Create a native port profile**
3. Create a Hyper-V Port Profile named **Toronto Default Network Adapter**. This should be an Uplink port profile that you assign to the Toronto Production network site, and for which you enable Network Virtualization.

Create a logical switch

1. Create a logical switch named **Toronto Logical Switch**, and then associate this with the **Toronto Default Network Adapter**.
2. Add the port classifications and virtual network adapter port profiles for the following:
 - **Host management/Host management**
 - **Live migration workload/Live migration**
 - **Host Cluster Workload/Cluster**
 - **High bandwidth/High Bandwidth Adapter**
3. Close the Jobs window.

Add the Microsoft Loopback adapter

1. On LON-HOST1, open the Control Panel, open Device Manager, right-click **LON-HOST1**, click **Add legacy hardware**, choose to **Install the hardware that I manually select from a list**, add a Network adapter with the Manufacturer **Microsoft** and the type **Microsoft KM-TEST Loopback Adapter**. Complete the Wizard.

Refresh LON-HOST1

2. On LON-VMM1, in the **Fabric** workspace, expand **Servers**, expand **All Hosts**, expand **London Hosts** and refresh **LON-HOST1**.

Update Hyper-V hosts to use logical networks

1. From the Fabric workspace, edit the properties of **LON-HOST1**, and create a new Virtual Switch of the Logical Switch type. Use the **Microsoft KM-TEST Loopback Adapter**.
2. Select the Logical network that is listed under the Microsoft KM-TEST Loopback Adapter, and then on the right of the **Hardware** page, under **Logical network connectivity**, select **Toronto Production Network**.

► Task 2: Configure network virtualization

Configure a virtual machine network for the application servers by using network virtualization

1. In the Virtual Machine Manager console, create a VM network named **Toronto Applications VM Network**, with the description **Toronto Application Servers**. Set the Logical network to **Toronto Production Network** and enable the **Isolate using Hyper-V network-virtualization** option.
2. Add the following VM subnet:
 - **Toronto Application Servers: 172.16.3.0/24.**
3. Set the gateway to **No connectivity** selected.
4. Repeat steps 1 through 3 to create a VM network named **Toronto Partner_Applications VM Network**, and then use the same subnet.

Create the VM Network IP pools

- Select the **Toronto Applications VM Network** and create an IP Pool. Set the IP Pool name to **Toronto Applications VM IP Pool**, set the network to **Toronto Applications VM Network**, and set the VM subnet to **Toronto Application Servers (172.16.3.0/24)**. Use the defaults for all other settings.

► **Task 3: Assign virtual machines to VM networks**

Assign virtual machines to VM Networks

1. Create a virtual machine by using the **Create the new virtual machine with a blank virtual hard disk** option named **TOR-CRM1**. Set this virtual machine so that Network Adapter 1 is connected to the **Toronto_Applications VMs Network**.
2. Assign the High bandwidth port profile, and when completing the rest of the virtual machine creation, use default settings.

Review the virtual machine

- From the VMs and Service workspace, review the network properties of TOR-CRM1, and then confirm that it has been assigned an IP address and is part of the correct VM Network.

► **Task 4: To prepare for the next module**

Do not revert the virtual machines, as you will need them for the next module.

Question: What type of business would benefit from network virtualization?

Question: Which two workloads could you consolidate into a single cluster?

Question: What are the new SAN types available to Hyper-V and its virtual machines?

Module Review and Takeaways

Best Practice: Storage and networking bandwidth is crucial to the planning process. Always look for bottlenecks, and calculate the amount of data that will transfer point-to-point. For example, if you host 500 servers on a SAN, and schedule an antivirus sweep, what is the impact?

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
Virtual machines are all paused	
Random iSCSI connectivity issues	
Live migrations fail, poor connectivity to virtual machines	

Review Question

Question: What is a benefit of logical switches?

Real-world Issues and Scenarios

After creating various logical networks and virtual machine networks, you are unable to remove a VMM object due to an error with a dependency on a temporary template. If this happens, you can remove the template by using Windows PowerShell®. Follow this procedure:

In the VMM console, on the ribbon, click **Windows PowerShell**, at the command prompt, type the following command, and then press Enter:

Get-SCVMTemplate | where {\$_.Name -like "Temporary*"}

Review the output, and then confirm that the only listed item is the suspicious temporary template, and that you do not have valid templates with the name "Temporary" in them.

Remove the problematic template by typing the following command at the command prompt, and then pressing Enter:

Get-SCVMTemplate | where {\$_.Name -like "Temporary*" } | Remove-SCVMTemplate

This should clear the dependent template, which will allow you to delete objects, such as a virtual machine network.

Tools

Microsoft Assessment and Planning Toolkit (MAP)

<http://go.microsoft.com/fwlink/?LinkID=285277>

Module 4

Planning and Deploying Virtual Machines

Contents:

Module Overview	4-1
Lesson 1: Planning a Virtual Machine Configuration	4-2
Lesson 2: Preparing for Virtual Machine Deployments with VMM	4-10
Lesson 3: Deploying Virtual Machines	4-21
Lesson 4: Planning and Implementing Hyper-V Replica	4-25
Lab: Planning and Implementing a Virtual Machine Deployment and Management Strategy	4-31
Module Review and Takeaways	4-41

Module Overview

To plan and deploy virtual machines, you must analyze existing workloads, identify application resources and requirements, and then configure suitable virtual machines for deployment to the best available hosts. In this module, you will learn how to configure virtual machines, reusable profiles, and templates to aid in deployment. You will also review application specific workloads and learn about the options for physical and virtual machine migrations.

Objectives

After completing this module, you will be able to:

- Plan virtual machine configurations.
- Plan and configure the Microsoft® System Center 2012 R2 Virtual Machine Manager (VMM) profiles and templates that you can use to implement a VMM deployment.
- Plan and implement a virtual machine deployment in VMM.
- Plan and implement a Microsoft Hyper-V® Replica.

Lesson 1

Planning a Virtual Machine Configuration

The first step in implementing server virtualization is to evaluate your organization's current server environment and determine which components you should virtualize. In most organizations, you can use virtualization to address many issues. However, getting the maximum benefit out of virtualization requires careful planning. This lesson provides an overview of the process and tools that you can use to evaluate and plan virtualization in an organization.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe considerations for virtualizing server roles.
- Describe the considerations for planning virtual machine configuration.
- Plan for Microsoft SQL Server® virtualization.
- Plan for Microsoft Exchange Server virtualization.
- Plan for Microsoft SharePoint® Server virtualization.
- Plan for Microsoft Active Directory® Domain Services (AD DS) virtualization.

Considerations for Virtualizing Server Roles

Virtualization is an established technology that many organizations have adopted. Enterprises often develop a server implementation policy that seeks to virtualize all new and replaced systems. These enterprises will consider deploying physical hardware as an alternative only when there is a valid reason not to virtualize. Server computing has evolved. With improved hypervisor capabilities, virtualizing broader and more intensive computer workloads has become an accepted practice.

You now can use Windows Server® 2012 and Windows Server 2012 R2 to deploy servers with up to 320 logical processors and 4 terabytes (TBs) of system memory. This is a significant improvement over older hypervisors and provides new capabilities for virtual workloads.

The following table provides an overview of the evolution of Windows Server 2012 Hyper-V virtual machine capabilities. This will help you to understand virtualization of more intensive workloads.

When identifying server workloads to virtualize, consider the following:

- Hardware requirements
- Compatibility
- Applications and services
- Supportability
- Licensing
- Availability requirements

	Release date	Random access memory (RAM)	Number of processors (virtual)	Storage (per drive)
Microsoft Virtual Server 2005 R2 Service Pack 1 (SP1)	2007	3.6 gigabyte (GB)	1	127 GB
Hyper-V in Windows® Server 2008	2008	64 GB	4	2,000 GB
Hyper-V in Windows Server 2008 R2 SP1	2009	64 GB	4	2,000 GB
Hyper-V in Windows Server 2012 and Windows Server 2012 R2	2012	1 TB	64	64 TB

Resource and performance are not the only considerations when implementing a server virtualization solution, so it is important to review all aspects of an application's requirements before deciding whether you can host it virtually.

Choosing Server Roles to Virtualize

You should consider these factors when choosing whether to virtualize server workloads:

- **Hardware requirements.** Typically, virtual machines require approximately the same resources as a physical server. For example, if a physical server is currently utilizing 1 GB of RAM, you should expect the virtual machine to use the same amount of RAM, assuming that it runs the same operating system and applications as the physical server. If a single virtual machine consumes more than half of your host's workload, you should consider whether virtualization is appropriate or if the host's sizing is adequate.
- **Compatibility.** You must determine whether the application can run in a virtualization environment. Business applications range from simple programs to complex, distributed multiple tier applications. You must consider requirements for specific components of distributed applications, such as requirements for communication with other infrastructure components or direct access to the system hardware. While you can virtualize some servers easily, some components may need to continue running on dedicated hardware.
- **Applications and services.** Applications and services that have specific hardware or driver requirements are not well suited for virtualization. An application may not be a good candidate for application virtualization if it contains low-level drivers that require direct access to the system hardware. This access may not be possible through a virtualization interface, or it may affect performance negatively.
- **Supportability.** You need to evaluate if a virtualized environment will support your operating system and requisite applications. Verify vendor support policies for deployment of the operating system and the application using the virtualization technologies.
- **Licensing.** You must evaluate whether you can license the application for use in a virtual environment. Reduced licensing costs for multiple applications or operating systems could add up and make virtualization financially practical.

- Availability requirements. Most organizations have some applications that must be virtually available for users at all times. Some applications provide built-in options for enabling high availability, while other applications are difficult to make highly available outside of a virtual machine environment. When considering whether to virtualize a server, evaluate whether the application has high availability options, whether a virtual machine environment supports those options, and whether you can use failover clustering to make the virtual machine highly available.

The goal in most organizations is to utilize all servers adequately, whether they are physical or virtual. You can fully utilize some server roles, such as SQL Server or Exchange Server Mailbox servers, by deploying additional SQL Server instances or moving more mailboxes to the server. In some cases, you can virtualize server workloads in one scenario, but not in another. For example, in a very large domain, with thousands of users logging on simultaneously, it may not be practical to virtualize a domain controller. However, in a smaller domain or in a branch office deployment, the virtualization of domain controllers may be your best option.

Planning Virtual Machine Configuration

When developing a virtualization strategy, you should aim to simplify and standardize the host computer and virtual machine configuration as much as possible. Consider the following general guidelines that apply to all the virtual machines:

- Develop a small number of standard virtual machine builds. To streamline the deployment and management of virtual machines, develop a set of standard virtual machine builds. For example, consider creating a standard low-end server build, a medium server build, and a high-end server build. Then assign a standard central processing unit (CPU) and memory configuration for each role. You should also consider configuring each of the virtual machines with a standard 50 GB system partition and providing additional disks to store data or install applications. Consider using SCSI controllers for all hard disks other than the disks containing the boot and system partition. With Windows Server 2012, you can add new virtual hard disks that you connect to a SCSI controller without restarting the server. With the new Generation 2 virtual machine in Windows Server 2012 R2, you can start from the SCSI controller, which has the added benefit of greatly improved performance. However, Generation 2 virtual machines are not backward-compatible with older versions of the Windows operating system and cannot be used in all scenarios. For example, in Windows Server 2012 R2 you cannot use these virtual machines for Remote Desktop Services Virtual Desktop Infrastructure (VDI) Virtual Machine Pools.
- Plan virtual machines for specific server roles. Although you should be able to configure most virtual machines with the same basic disk and operating system configuration, the actual physical requirements for each virtual machine will vary. For example, some virtual machines will require significantly more RAM or CPU resources than others will. To design the physical requirements for a virtual machine, consider the following guidelines:
 - Monitor the servers before virtualizing them. Collect performance data on the servers to evaluate how specific applications perform on physical servers. If an application uses a very low percentage of the hardware resources on a physical server, deploy a virtual server with significantly less capacity to run the same application.

To standardize the virtual machine configuration:

- Develop a small number of standard virtual machine builds
- Plan virtual machines for specific server roles by:
 - Monitoring the servers before virtualization
 - Configuring each virtual machine with a hardware configuration that is similar to the hardware required on a physical server
- Deploy Windows Server 2012 R2–based virtual machines whenever possible
- Consider other options for ensuring physical server utilization

- Configure each virtual server with a hardware configuration that is similar to the hardware required for the application on physical servers. The fact that you are virtualizing a server does not change the hardware resources that the server requires.
- Deploy Windows Server 2012 R2–based virtual machines whenever possible. Use Dynamic Memory, and review any support statements pertaining to virtualization and Dynamic Memory for any application you will host on your virtual machine.

When considering virtualization, review the applications and make use of affinity and anti-affinity. *Affinity* is the process of grouping certain virtual machines together on single host. *Anti-affinity* is purposely preventing certain virtual machines from deploying to, or residing on, the same cluster nodes.

Planning Virtual Machines for SQL Server

SQL Server is one of the server workloads that organizations are virtualizing. This is especially practical in development, testing, and training environments that require SQL Server, in which you often perform new installations. If you are planning a new SQL Server deployment or a SQL Server consolidation project, you should consider virtualization, which will provide maximum utilization of your hardware.

Consider the following recommendations for configuring virtual machines that run SQL Server:

Consider the following recommendations for configuring virtual machines that run SQL Server:

- Ensure that the Hyper-V integration components are installed
- Plan virtual machine hardware settings to match physical server hardware
- Use fixed-size virtual hard disks and SCSI controllers for database and log file drives
- To ensure adequate CPU capacity:
 - Remember that virtual machines are limited to 64 virtual CPUs
 - Do not overcommit CPU resources
 - Remember that networking-intensive workloads require more CPU capacity

- Ensure that you install the Hyper-V integration components on the guest virtual machine. Additionally, use standard network adapters rather than legacy network adapters when configuring networking for the virtual machine, if you are using Generation 1 virtual machines. Following these two steps will provide enhanced performance for the virtual machines.
- Plan to configure the hardware settings for the virtual machines to match the hardware settings that you would configure on a physical server for the same workload.
- Plan virtual machine storage. If you want to ensure optimal performance for any SQL Server instance, you must use a storage system that is the correct size and configuration. The storage hardware should provide sufficient I/O throughput, in addition to adequate storage capacity to meet the current and future needs of the planned virtual machines. You can now use a virtual hard disk (.vhdx) to provide storage for up to 64 TB drives. Alternatively, you can use Fibre Channel directly in the virtual machine. Introduced in Windows Server 2012, .vhdx provides a significantly larger storage capacity than the previous .vhd format.
- Follow the recommended best practices for configuring disks for transaction logs and database storage. While pass-through disks provide the best performance for SQL Server, their lack of portability can make the deployment more complicated. Fixed-size virtual hard disks provide almost the same performance. Typically, they are the best disk option for SQL Server.
- Generation 1 Hyper-V virtual machines must use an IDE controller for the boot and system partitions, but you should use synthetic SCSI controllers for the disks containing SQL Server databases and logs. Generation 2 Hyper-V virtual machines start from a synthetic SCSI controller.

- Provide adequate CPU capacity, which is critical to SQL Server performance. When designing a virtualization host that will run multiple SQL Server virtual machines, you should ensure that the host's cumulative physical CPU resources are adequate to meet the needs of all guest virtual machines. Just as when you are deploying multiple SQL Server instances on a physical server, the only way to guarantee adequate performance is to test the deployment thoroughly. When running SQL Server on a virtual machine, you will need to consider the following CPU-based limitations:
 - When using Hyper-V, you can assign up to 64 CPU cores to a virtual machine. If you are working with large virtualization hosts, you should not overcommit CPU resources. This happens when the total number of logical CPU cores that you configure across all guest virtual machines is more than the actual number of physical CPU cores that the server has available. Overcommitting the CPU cores can affect server performance significantly when you are utilizing all the virtual machines heavily.
 - Networking-intensive workloads will result in higher CPU overhead and more performance impact on a virtual machine. During planning, consider using guest network adapter teaming, single-root I/O virtualization (SR-IOV), and Virtual Machine Queue (VMQ), if supported.
- SQL Server versions that support Hot Add Memory also support Dynamic Memory. However, while earlier versions of SQL Server are supported, they can only utilize the assigned startup memory.



Best Practices for Virtualizing and Managing SQL Server:

<http://go.microsoft.com/fwlink/?LinkID=393716>



For more information on the support policy for SQL Server products that run in a hardware virtualization environment, go to:

<http://go.microsoft.com/fwlink/?LinkID=285286>

Planning Virtual Machines for Exchange Server

You can use a virtualization environment to run all Exchange Server 2007, Exchange Server 2010 SP1, and Exchange Server 2013 server roles. Consider the following guidelines when virtualizing Exchange servers:

- Use standard server sizing. Running Exchange on a guest virtual machine does not change the Exchange Server design requirements from an application perspective. The Exchange Server guest virtual machine still must be the appropriate size to handle the workload.
- Configure appropriate storage. Exchange Server virtual machines can use fixed-size virtual hard disks, SCSI physical storage, or Internet SCSI (iSCSI) storage. As with SQL Server-based servers, physical storage provides the best performance. However, it does not support dynamically expanding virtual disks or differencing drives.

When designing virtual machines for Exchange Server:

- Use standard server sizing rules
- Configure appropriate storage
- Do not use virtual machine checkpoints
- Configure adequate CPU resources
- Consider how to use Hyper-V and native Exchange Server high availability
- Consider I/O requirements

You should use separate logical unit numbers (LUNs) on Redundant Array of Independent Disks (RAID) arrays for the host operating system, each guest operating system disk, and all virtual machine storage. As with physical servers, you should create separate LUNs for each database and set of transaction log files.

- Do not use virtual machine checkpoints. Virtual machine checkpoints are not application-aware, and using them can cause unintended and unexpected consequences for a server application that maintains state data, such as Exchange Server. Therefore, Exchange Server virtual machines do not support checkpoints.
- Configure adequate CPU resources. Exchange Server supports a ratio of virtual processors to logical processors of no greater than two to one. For example, a dual-processor system that uses quad-core processors contains eight logical processors in the host system. On a system with this configuration, do not allocate more than 16 virtual processors to all guest virtual machines combined. If you are utilizing the CPUs heavily for all virtual machine instances, overcommitting the CPUs will affect performance significantly. In these scenarios, do not assign more virtual processors to virtual machines than the number of processor cores on the host computer.
- High availability for Exchange servers. Exchange Server provides several options for high availability. For server roles such as Client Access servers, Hub Transport server roles, and Edge Transport server roles, you can deploy multiple servers for each role to ensure that the server role is available if a single-server failure occurs. For Mailbox servers, Exchange Server 2007 provides several Exchange clustering solutions, such as cluster continuous replication (CCR) and single copy clusters (SCCs). Exchange Server 2010 and Exchange Server 2013 provide database availability groups (DAGs). These solutions provide various options for automatic failover if a server failure occurs.

With Hyper-V, you can make virtual machines highly available by deploying them in a failover cluster. You can use failover clustering to make virtual machines running the Client Access server, the Hub Transport server role, and the Edge Transport server role highly available. You can combine the Exchange Server Mailbox server high availability options with failover clustering.

- Mailbox server performance. The most common performance bottlenecks for Mailbox servers are disk I/O and network I/O. Running Mailbox servers in a virtual environment means that the virtual machines have to share this I/O bandwidth with the host machine and with other virtual machine servers that you deploy on the same host. If a single virtual machine is running on the physical server, the disk I/O and network I/O that are available to the virtual machine are almost equivalent to the I/O that is available to a physical server. However, a heavily utilized Mailbox server can consume all available I/O bandwidth, which makes it impractical to host additional virtual machines on the same physical server.



Note: Exchange Server does not support the Dynamic Memory feature in Hyper-V.



Best Practices for Virtualizing & Managing Exchange 2013:

<http://go.microsoft.com/fwlink/?LinkID=393717>

Planning Virtual Machines for SharePoint Server

When using virtual machines, you have the option of deploying Microsoft SharePoint Foundation and Microsoft SharePoint Server.

Consider the following recommendations for deploying Windows SharePoint Services or SharePoint Server on a virtual machine:

- Ensure that you configure each virtual machine with the same hardware capacity that a physical server would require. Additionally, consider the overhead performance on the host computer for each virtual machine.
- Do not take checkpoints of virtual servers that connect to a SharePoint server farm. If you do, the timer services and the search applications might become unsynchronized during the checkpoint process. To take server checkpoints of, first detach the server from the server farm.
- Avoid overcommitting the number of virtual CPUs. Although Hyper-V will allow you to allocate more virtual CPUs than physical CPUs, this causes performance issues because the hypervisor software has to swap out CPU contexts. This is a problem only if the virtual machines are utilized heavily.
- Ensure that you assign adequate memory to each virtual machine. Inadequate memory will have the greatest impact on server performance. The amount of memory required depends on the server workload, so you will need to test and optimize memory configuration for each scenario.
- Choose the right storage implementation. If you run only front-end web servers or query servers on virtual machines, the disk performance is not as important as when the image hosts the index role or a SQL Server database. If the image hosts the index role, you should use a fixed-size VHD or a physical disk.
- Consider monitoring as a design element of a virtualized SharePoint server farm. Microsoft System Center 2012 R2 Operations Manager has management packs, which help monitor the Windows Server operating system, Hyper-V, SQL Server, Internet Information Services (IIS), SharePoint, and your server hardware. By using these management packs, you can create a baseline for the SharePoint server farm, configure error alerting, and monitor ongoing performance closely. You will learn more about these management packs in a later module.

When designing virtual machines for SharePoint Server:

- Configure virtual machine hardware with the same capacity as physical server hardware
- Do not take checkpoints of virtual servers
- Avoid overcommitting the virtual CPUs
- Assign adequate memory
- Choose the right storage implementation
- Configure monitoring

The resource centers for both SharePoint Server 2010 and SharePoint Server 2013 have topology and planning guides that specifically guide administrators seeking to build virtualized SharePoint server farms.



For the SharePoint Server 2010 resource center, go to:

<http://go.microsoft.com/fwlink/?LinkID=285291>



Best Practices for Virtualizing & Managing SharePoint 2013:

<http://go.microsoft.com/fwlink/?LinkID=393718>

Planning Virtual Machines for AD DS

You should be aware of the risks involved in virtualizing AD DS domain controllers. Typically, you can mitigate these risks by following some basic recommendations, including:

- Do not perform online physical-to-virtual (P2V) migrations.
- Do not stop or pause domain controllers.
- Do not restore checkpoints of domain controllers unless the domain controllers and the hypervisor support Virtual Machine Generation ID (VM Generation ID).
- Consider building virtual servers, and then promoting them to domain controllers and demoting the physical servers.

- Windows Server 2012 introduces new, safer virtualization of domain controllers, including domain controller cloning
- Windows Server 2012 and Windows Server 2012 R2 domain controllers support checkpoints through VM Generation ID on supported hypervisors
- Recommendations:
 - Do not perform online physical-to-virtual (P2V) migrations
 - Do not stop or pause domain controllers
 - Do not restore checkpoints of domain controllers unless the domain controllers and the hypervisor support Virtual Machine Generation ID (VM Generation ID)
 - Consider building virtual servers, and then promoting them to domain controllers and demoting the physical servers

Prior to Windows Server 2012, virtualized domain controllers contained no specific technology to accommodate or benefit from virtualization. Windows Server 2012 introduces support for virtualizing AD DS domain controllers safely, specifically for hypervisor platforms that have the VM Generation ID identifier. When Windows Server 2012 detects the VM Generation ID identifier, it provides measures to protect the AD DS environment, as long as you roll back the virtual machine to a previous version.



For more information about the VM Generation ID identifier, go to:

<http://go.microsoft.com/fwlink/?LinkID=260709>

In addition to providing protection, you can now clone Windows Server 2012 and Windows Server 2012 R2 domain controllers. This means that clones are able to identify that they are clones. Until now, you were required to use the System Preparation tool (Sysprep) to prepare a Windows Server, and then deploy the domain controller services after the server was deployed. This change allows you to use simpler disaster recovery.

When virtualizing domain controllers for production systems, evaluate the virtual machine placement and use anti-affinity to avoid placing domain controllers on the same cluster node.



For an introduction to AD DS virtualization, go to:

<http://go.microsoft.com/fwlink/?LinkID=285289>

Lesson 2

Preparing for Virtual Machine Deployments with VMM

VMM provides a solution for creating and storing virtual machines, templates, and profiles. This solution can speed up delivery of a virtual machine and reduce configuration errors. After building templates and profiles, you can delegate permissions to ensure that administrators can deploy new virtual machines only, based upon the validated and assigned template.

Lesson Objectives

After completing this lesson, you will be able to:

- Configure guest operating systems profiles.
- Configure hardware profiles.
- Configure a virtual machine in VMM.
- Configure SQL Server profiles.
- Configure application profiles.
- Configure virtual machine templates.
- Configure service templates.
- Describe considerations for planning VMM profiles and templates.

Configuring Guest Operating System Profiles


A guest operating system profile contains operating system settings that you use during a virtual machine deployment. You can use the guest operating system profile as one of the building blocks for constructing a virtual machine template.

Guest Operating System Profile Overview

A *guest operating system* refers to any operating system that runs on a virtual machine in Windows Server using Hyper-V technology. As you install a guest operating system on multiple virtual

machines, you may soon realize that many virtual machines contain similar system settings, such as domain or workgroup membership, product keys, time zone, and the local administrator password.

You can create and use a guest operating system to support an automated and standardized virtual machine deployment process. Guest operating system profiles contain a collection of operating system settings that the virtual machine deployment process imports into a virtual machine template. The virtual machine template provides a consistent operating system configuration for any virtual machine that you create using the template.



You can use a guest operating system profile to provide predefined configuration settings for the guest operating system. These settings include:

- Identity information
- Local administrator password
- Product key
- Time zone
- Operating system version
- Server roles and features
- Domain/workgroup membership
- Answer file references

The *guest operating system profile* is a database object that you create and access from within the Library workspace in the VMM console.

Creating a Guest Operating System Profile

To create a guest operating system profile, perform the following steps:

1. In the VMM console, click the **Library** workspace.
2. In the navigation pane, expand **Profiles**, and then click **Guest OS Profiles**. Any existing profiles display in the results pane.
3. On the **Home** tab, click **Create**, and then click **Guest OS Profile**. The **New Guest OS Profile** dialog box opens.
4. On the **General** page, provide the profile's name and description.
5. On the **Guest OS Profile** page, configure settings as required, and then click **OK**.

The Guest OS Profile page allows you to configure the following settings:

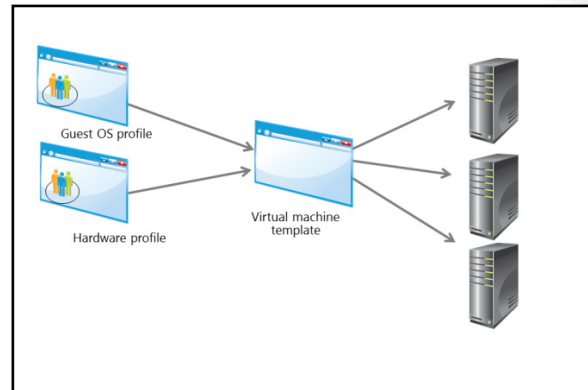
- **General Settings.** This section includes settings such as computer name, local administrator password, product key, time zone, and the type of operating system that you will be deploying to the virtual machine.
- **Roles and Features.** This section allows you to select one or more server roles and features that you want to install on the virtual machine that you are deploying.
- **Networking.** This section allows you to specify the workgroup or domain that the virtual machine should join.
- **Scripts.** This section allows you to include additional settings, which you specify in a Unattend.xml file or a Sysprep.inf file, or through commands that you configure within the [GUIRunOnce] section of the registry key.

Configuring Hardware Profiles

You use a hardware profile to define a standard set of hardware settings that you want to use during a virtual machine deployment. The hardware profile is another building block that you can specify when you construct your virtual machine template.

Hardware Profile Overview

A hardware profile contains specifications for various hardware components such as the number of processors, memory allocation, IDE devices, SCSI adapter configuration, and network adapter configuration. Although you can deploy a virtual machine without a hardware profile, using a hardware profile in conjunction with a virtual machine template ensures that your virtual machine deployment uses a consistent hardware configuration.



Creating a Hardware Profile

To create a hardware profile, perform the following steps:

1. In the VMM console, click the **Library** workspace.
2. In the navigation pane, expand **Profiles**, and then click **Hardware Profiles**. Any existing profiles display in the results pane.
3. On the **Home** tab, click **Create**, and then click **Hardware Profile**. The **New Hardware Profile** dialog box opens.
4. On the **General** page, provide the profile's name, description, and virtual machine type (Generation 1 or Generation 2).
5. On the **Hardware Profile** page, configure settings as required, and then click **OK**.

On the **Hardware Profile** page for a Generation 1 virtual machine, you can configure the following settings:

- **Compatibility.** This setting provides an option to select a preconfigured capability profile, which ensures that the hardware profile meets specific hardware capability requirements.
 - **General.** This section allows you to configure settings related to the processor, memory, floppy drive, communications ports, and video adapter.
 - **Bus configuration.** This section allows you to configure settings for IDE devices and SCSI adapters.
 - **Network adapters.** This section allows you to specify connectivity settings for one or more network adapters.
 - **Fibre Channel adapters.** This section allows you to specify settings for one or more Fibre Channel adapters.
 - **Advanced.** This section provides a number of settings related to availability, basic input/output system (BIOS) configuration, virtual non-uniform memory access (NUMA), and CPU and memory priorities.
6. On the **Hardware Profile** page for a Generation 2 virtual machine, there are fewer hardware settings. Note the changes under the following hardware sections:
 - **General.** This section allows you to configure settings related to the processor and memory.
 - **Bus configuration.** This section allows you to configure settings SCSI adapters and devices.

Demonstration: Configuring a Virtual Machine in VMM

In this demonstration, you will see how to configure a virtual machine by using VMM.

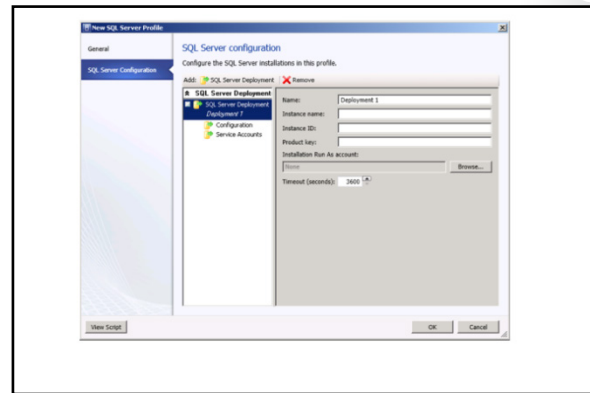
Demonstration Steps

Configuring a virtual machine

1. Start the Virtual Machine Manager console.
2. Click the **VMs and Services** workspace, on the ribbon, click **Create Virtual Machine**, and then in the drop-down list box, click **Create Virtual Machine**.
3. Work through the Create Virtual Machine Wizard to create a virtual machine from a blank virtual hard disk.
4. On the **Configure hardware** page of the Create Virtual Machine Wizard, configure the virtual hardware for your virtual machine, and then click **Save As** to save a hardware profile with a unique name such as **New SharePoint Web Farm**.
5. Navigate through the remaining configuration pages, carefully reviewing all of the options for deploying a virtual machine.
6. Notice that when you reach the **Select Host** page, a placement assessment evaluates the options based on the configuration that you select.

Configuring SQL Server by Using SQL Server Profiles

Many web-based applications and multitier services use SQL Server for database functionality. Often you must deploy database applications to support virtualized services within the private cloud. You can use a SQL Server profile as a building block for deploying instances of SQL Server to virtual machines. A SQL Server profile provides the building blocks for configuring a prepared instance of SQL Server on a virtual machine image. This prepared virtual hard disk must have SQL Server deployed. You must have generalized it by using the Sysprep tool. System Center 2012 R2 Virtual Machine Manager supports SQL Server 2008 R2 and SQL Server 2012. The profile contains configuration settings for each instance that was prepared previously on the virtual machine.



To create a SQL Server profile, complete the following steps:

1. Open the VMM console, and then click the **Library** workspace.
2. In the navigation pane, expand **Profiles**, and then click **SQL Server Profiles**.
3. In the ribbon, click **Create**, and then click **SQL Server Profile**. The **New SQL Server Profile** dialog box opens.
4. On the **General** page, provide a **Name** and **Description** for the profile.

5. On the **SQL Server Configuration** page, for each instance that you need to configure, click **SQL Server Deployment**, and then configure the following:
 - **Name.** Use this required setting to specify the name of the SQL Server deployment in the profile. Each instance will have a unique name for identification.
 - **Instance name.** Use this required setting to specify the SQL Server instance name.
 - **Instance ID.** Use this required setting to enter the instance ID that you documented when you prepared the SQL Server image.
 - **Product key.** Use this optional setting to specify the product key for SQL Server. If you do not configure this setting, the evaluation version installs.
 - **Installation Run As account.** Use this optional setting to specify the Run As account with which you want to run the SQL Server setup. If you do not specify an account, the installation uses the virtual machine service account.
 - **Timeout (seconds).** Use this optional setting to specify a timeout window in which the SQL Server installation has to finish. By default, this value is configured for 3,600 seconds or one hour.
6. Click **Configuration**, and then configure the following:
 - **Media source.** Use this required setting to specify the path to the installation media folder. You can place the media locally on the VHD or you can specify a path to a network share. If you use a network share, you must configure the Installation Run As account with credentials that have permission to access the network share and administrator privileges for the guest virtual machine.
 - **SQL Server administrators.** Use this required setting to specify users or groups that should be members of the system administrator role.
 - **Security mode.** Use this optional setting to choose between Windows Authentication (which is the default) and SQL Server Authentication.
 - **System Administrator (SA) password Run As account.** If, under **Security mode**, you selected **SQL Server Authentication**, use this setting to provide the password for the system administrator account.
 - **TCP/IP.** Use TCP/IP for remote connections. This optional setting enables the TCP/IP protocol for the SQL Server service.
 - **Named pipes.** Use named pipes for remote connections. You can use this optional setting to enable the named pipes protocol for the SQL Server service.
 - **SQL Server configuration file.** Use this optional setting to specify a SQL Server configuration file. The file must reside on a Virtual Machine Manager library share.
7. Click **Service Accounts**, and then configure the following:
 - **SQL Server service Run As account.** Use this required setting to specify the account for use with the SQL Server service.
 - **SQL Server agent service Run As account.** Use this required setting to specify the account for use with the SQL Server Agent service.
 - **Reporting services Run As account.** Use this optional setting to specify the account for use with Reporting Services.

Configuring Application Profiles

A *service* in VMM is a collection of virtual machines that you configure, deploy, and manage together. When you deploy a service by using VMM, that service will contain applications that integrate with Web Services or a SQL Server instance. You can configure and deploy application profiles to provide installation and configuration settings that VMM will use to deploy specific types of applications with a service.

Application profiles provide the instructions for installing applications to support a VMM-managed service, and these profiles support the following application types:

- SQL Server data-tier applications
- Server App-V applications
- Web applications
- Scripts

Application Profiles Overview

When you deploy a virtual machine as part of a service, application profiles provide configuration instructions for installing specific application types. Application profiles support the following application types:

- **SQL Server data-tier application.** SQL Server 2008 R2 and newer versions support a new package type called a data-tier application. A data-tier application contains the entire database and instance objects that the application uses. Typically, it is targets department-based applications.
- **Server App-V applications.** Microsoft Server Application Virtualization (Server App-V) is a technology that creates virtual application packages that you then deploy to servers that run the Microsoft Server Application Virtualization agent (Server App-V agent). A virtual application package does not require a local installation; however, the package runs as if it is a locally installed application. When you create a Server App-V package, the Server App-V Sequencer monitors a typical application installation and records information that is required for the application to run in a virtual environment.

Once you have created the Server App-V package, you can import it into the Virtual Machine Manager library so that it is accessible from an application profile.

- **Web application.** A *web application* is a package that is stored within the Virtual Machine Manager library, and that contains the content, websites, certificates, and registry settings of a web-based application. You can package and deploy web applications with the Microsoft Web Deployment Tool. VMM also uses this tool to deploy web applications as a service when deploying a web application as specified in an application profile.
- **Scripts.** When deploying a virtual machine as part of a service, you can use the application profile to run scripts. You use scripts during the preinstallation and the post-installation phases of a specific application. For example, you might need to copy updated configuration files to a deployed web application, or you may have to run specific virtual application commands to finalize a virtual application deployment.

Creating an Application Profile

To create an application profile, complete the following steps:

1. Open the VMM console, and then click the **Library** workspace.
2. In the navigation pane, expand **Profiles**, and then click **Application Profiles**.
3. On the ribbon, click **Create**, and then click **Application Profile**. The **New Application Profile** dialog box opens.
4. On the **General** page, provide a **Name** and **Description** for the application profile.

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

Deploying an Application Profile

As with a SQL Server profile, you can use an application profile only when you are deploying a virtual machine as part of a service.

To configure a service template for use in deploying an application with a service, use the Service Template Designer in VMM. Two options are available for specifying an application configuration within a service template:

- Create a virtual machine template, and then specify the application configuration settings. Then you can use the virtual machine template when you create the service template.
- Edit the properties of a service tier and specify the application profile manually.

Configuring Virtual Machine Templates

You use virtual machine templates to help you create new virtual machines. Then you can add the templates to tiers in a service template. The virtual machine template combines many of the settings that you would configure in hardware profiles, guest operating system profiles, application profiles, and SQL Server profiles.

Virtual Machine Template Overview

When you create a new virtual machine, you can derive the source of the new virtual machine from an existing virtual machine or hard disk, or you can base the new virtual machine on a virtual machine template. If you use a stored virtual machine, you can customize only the hardware settings; there is no option for adding additional information such as the operating system configuration or applications. However, a virtual machine template provides additional flexibility and efficiency for virtual machine deployment. The advantages of using a virtual machine template include the following:

- You can configure hardware, operating system, applications, and SQL Server specifications.
- You can create new virtual machines or service templates.
- You can share virtual machine templates with self-service roles to provide a consistent virtual machine deployment process.

Virtual machine templates provide an efficient way to deploy new virtual machines and services

These templates allow you to configure:

- Hardware
- Operating systems
- Applications
- SQL Server specifications

When you create a virtual machine template, you can configure the following:

- **Identity.** You can configure the template source, such as another virtual machine or an existing virtual hard disk. You provide a name and select either a Generation 1 or Generation 2 virtual machine.
- **Hardware profile.** You can configure the hardware settings directly in the virtual machine template, or you can specify a preconfigured hardware profile. You can save any modifications as a new hardware profile that the Virtual Machine Manager library stores. The main difference between the hardware configurations in the virtual machine template and in the hardware profile is that in the virtual machine template, you can create, remove, and configure disks as required.
- **Guest Operating System profile.** In the virtual machine template, you can configure the guest operating system profile settings manually, or you can import settings from a preconfigured guest operating system profile template. If you do not need to customize the operating system, on the Configure Operating System page of the Create VM Template Wizard, you can select None – customization not required.
- **Application profile.** You can configure application profile settings manually, import settings from a preconfigured application profile, or choose not to install any applications.
- **SQL Server profile.** You can configure SQL Server installation settings manually, import settings from a preconfigured SQL Server profile, or choose not to provide SQL Server configuration settings in the virtual machine template.

Creating a Virtual Machine Template

VMM provides several methods that you can use to create virtual machine templates. However, you must understand the implications of each method. The following table describes the methods and considerations for each method.

Method	Considerations
Create a virtual machine template from an existing virtual hard disk that the Virtual Machine Manager library stores.	Typically, the source virtual hard disk has an operating system that was installed and prepared by using the Sysprep tool. If you choose to use a source virtual hard disk that is not Sysprepped, you can configure a noncustomized virtual machine template in which the guest operating system profile is set to None – customization not required.
Create a virtual machine template from an existing virtual machine template that the Virtual Machine Manager library stores.	You can use the settings of a preconfigured virtual machine template as the basis for a new virtual machine template. All preconfigured and modified settings are saved in a new template. Then the Virtual Machine Manager library stores this template and makes it available.
Create a virtual machine template from an existing virtual machine that is deployed on a host.	You can only choose a source virtual machine that is deployed on a host and not a virtual machine that is stored in the library. You configure the virtual machine settings in the template and generalize the virtual disks of the virtual machine using Sysprep. Then you move the virtual machine into a Virtual Machine Manager library share, where it becomes no longer available on the host. You can further modify the virtual machine template as needed.
Import a preconfigured template.	You can use the Import Package Wizard to import preconfigured templates that have been configured in other virtualization platforms. You can start the Import Package Wizard by clicking the Import Template button on the ribbon.

Configuring Virtual Machine Templates

Use the following process to create a new virtual machine template based on a virtual hard disk that is stored in the Virtual Machine Manager library:

1. Open the VMM console, and then click the **Library** workspace.
2. In the navigation pane, expand **Templates**, and then click **Virtual Machine Templates**.
3. On the ribbon, click **Create Virtual Machine Template**. The Create Virtual Machine Template Wizard opens.
4. On the **Select Source** page, select one of the following options, and then click **Next**:
 - **Use an existing virtual machine template or a virtual hard disk stored in the library**
 - **From an existing virtual machine that is deployed on a host**
5. On the **Virtual Machine Template Identity** page, provide a Virtual Machine Template name and **Description**.
6. On the **Configure Hardware** page, configure the displayed hardware profile settings, or select a preconfigured hardware profile, and then edit as required.
7. On the **Configure Operating System** page, configure the displayed guest operating system profile settings, or select a preconfigured profile, and then edit as required.
8. On the **Configure Applications** page, configure the displayed application profile settings, or select a preconfigured profile, and then edit as required.
9. On the **Configure SQL Server** page, configure the displayed SQL Server profile settings, or select a preconfigured profile, and then edit as required.
10. On the **Summary** page, click **Create**.

Configuring Service Templates

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* encapsulates everything required to deploy and run a new instance of an application. Just as a private cloud user can create new virtual machines on demand, so can a user utilize service templates to install and start new applications on demand.

Information Included in the Service Template

The service template includes information about the virtual machines that are deployed as part of the service. In addition, the service template includes information about the applications to install on the virtual machines and the networking configuration needed for the service, including the use of a load balancer. The service template can also make use of existing virtual machine templates. While you can define the service without using any existing virtual machine templates, it is much easier to build a template if you have already created virtual machine templates. After creating the service template, you can configure it for deployment by selecting the template and then clicking the Configure Deployment option on the ribbon.

A service template encapsulates all necessary components for deploying and running a new instance of an application:

- Administrators create service templates in VMM
- Application owners deploy services based on the service template
- Users use App Controller or the VMM console to deploy a service based on template

Process for Deploying a New Service

When using service templates in VMM, the process of deploying a new service or application is as follows:

1. The system administrator creates and configures service templates in VMM by using the Service Template Designer.
2. The application owner (for example, a developer that needs to deploy the application environment), opens the Microsoft System Center 2012 R2 App Controller portal and requests a new service deployment based on available service templates that he or she can access. Then the application owner can deploy the service to a private cloud where a user has access. As an alternative to App Controller, the user can use the VMM console.
3. The VMM management server submits and evaluates a request. VMM searches for available resources in the private cloud, then calculates the user quota and verifies that the private cloud is capable of hosting the requested service deployment.
4. While the service is created automatically, the virtual machines and applications (if any) are deployed on the host that VMM has chosen.
5. The application owner gains control over service virtual machines through the App Controller portal or by using Remote Desktop Protocol (RDP).

If you require manual approval for resource creation, you can use Microsoft System Center 2012 R2 Service Manager to create workflows for this purpose.

Configuring Service Templates

Each service template that you create in VMM has several settings that you can configure. You access these settings by opening the Properties window of the service template that you are creating.

The most important service template properties include:

- **Name.** This service template will appear in the Virtual Machine Manager library. Self-service users will see this name, so it should be descriptive.
- **Release.** Release is a value indicating the version of the service template. This value is important when you update a service, because the release value helps you to identify the version of the service template. After you create a service template and use it to perform a deployment, you can make no further changes to it. If you want to make changes, you must create a new version.
- **Dependencies.** In the Dependencies setting, you can view objects that derive from a specific service template and library resources that the template references. You cannot make any changes to this setting.
- **Access.** You can define the owner of the service template in the Access setting. You can also list self-service users who can use this service template to deploy a service. If you want to allow self-service users to deploy services by using the VMM console or App Controller, you must add them to the access list for the service template.
- **VM network.** You can specify the VM networks to which a service can be deployed.

We recommend that you configure all service template settings before you begin deploying services based on that template.

Planning VMM Profiles and Templates

As part of your virtualization strategy and infrastructure design, consider the number of different templates, stored virtual disks, and hardware, guest application, and database profiles you will need. Plan how many different operating systems you will deploy and where you will store your files.

When planning operating system profiles, consider using automatic virtual machine activation. Automatic virtual machine activation allows you to deploy virtual machines that run Windows Server 2012 R2 on a virtualization host that runs Windows Server 2012 R2 Datacenter, without the need for managing individual product key. Activation of the virtual machine is bound to the host at the virtual machine start-up. The virtual machine retains this activation when you move or migrate the virtual machine to another host.

Considerations for working with profiles and templates include:

- Work with library items
- Keep the VMM library organized
- Do not create profiles that over-assign resources
- Keep offline virtual machines and virtual hard disks updated
- Keep licensing in mind when creating images
- Use revisions when updating templates
- Use Windows Server 2012 R2 autoactivation

Considerations for working with profiles and templates include:

- Working with library items. Consider the number of templates you think you will need, and then configure some or all of these before starting deployment. Consider keeping only the number of .vhd files that you require. In a mixed host environment, you may want to include both .vhd and .vhdx formats. Remove legacy and unused profiles and templates. Back up a library occasionally, and if you must recover an older image, you can retrieve it from the backup.
- VMM library organization. Try to keep the Virtual Machine Manager library organized and prevent virtual sprawl. Remove virtual machines and virtual hard disks that are unused. Virtual sprawl includes offline files, which can end up being stored across file and infrastructure servers other than hosts.
- Offline transfer. If you need to have the same templates and files across multiple Virtual Machine Manager libraries, you can send large files offline, and then import them where required. To avoid using small wide area network (WAN) links, set up equivalent objects at multiple locations for virtualization deployment resources that you need and do not want deployed over a WAN.
- Performance. Consider the impact of servicing many offline files. If updating them in a large organization, consider collaborating to ensure that one person is not servicing images while another is trying to deploy.
- A standard hardware profile. If you set the base configuration for all of your virtual servers with more memory, processors, and disk space than is needed, you are wasting resources. You are not achieving the full value of virtualization.
- Licensing. You can use the guest operating system profiles to help enforce licensing requirements. For example, you can preconfigure an image for Visual Studio, and then assign this image to the developers who have the Microsoft Developer Network (MSDN) agreement. Consider licensing when using a template that is based on another machine; ensure that only the people who should use a template are using it.
- Systems integration, automation, and self-service. VMM and its libraries are the source from which other applications deploy. If necessary, create multiple libraries with appropriate security and ensure the files and images deployed are up to date.
- Service Templates. When building services for applications that scale out, consider versions and revisions, and try to keep them consistent. For example, when you are updating the tiers of a three-tier application, you must remember to increment the revisions appropriately.

Lesson 3

Deploying Virtual Machines

The value of virtualization is that it brings high availability and portability to applications and services, reducing the cost of deployment and maintenance. Virtual machine deployment is a key component in an organization's virtualization strategy. Mastering virtual machine deployment will save time and reduce future workloads while improving the consistency of virtual machine configuration.

In this lesson, you will learn about the methods that you can use to deploy virtual machines from VMM. You will also learn about methods and tools for virtual-to-virtual (V2V) conversion.

Lesson Objectives

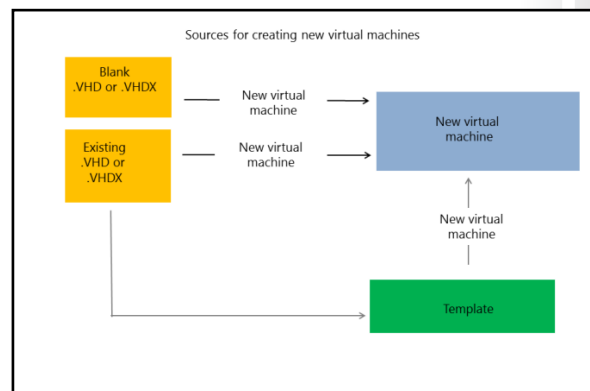
After completing this lesson, you will be able to:

- Deploy virtual machines by using VMM.
- Configure virtual machine placement in VMM.
- Describe how to perform V2V machine conversions in VMM.
- Describe considerations for V2V conversions.

Deploying Virtual Machines in VMM

Using VMM to manage a virtualized environment gives you the flexibility to create and deploy new virtual machines quickly. By using VMM, you can create a new virtual machine manually with new configuration settings and a new hard disk. Then you can deploy the new virtual machine from one of following sources:

- An existing .vhd or .vhdx file (blank or preconfigured)
- A virtual machine template
- A service template



You can create new virtual machines by converting an existing physical machine, cloning an existing virtual machine, or using a blank virtual hard disk or a preconfigured virtual hard disk that contains a Sysprepped operating system. VMM provides four blank VHD and VHDX templates that you can use to create new disks:

- VHD Blank Disk-Small
- VHD Blank Disk-Large
- VHDX Blank Disk-Small
- VHDX Blank Disk-Large

You can also use a blank VHD when you want to use an operating system with a Pre-Boot EXecution Environment (PXE). Alternatively, you can place an .iso image on a virtual DVD-ROM and then install an operating system on the empty drive. 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 machine, or initiate the guest operating system setup through a network service boot.

You can choose existing VHDs when deploying any operating system from which VMM cannot create a template, such as an operating system other than Windows.

When you create a new virtual machine using an existing VHD or VHDX file, you are essentially creating a new virtual machine configuration that is associated with the file. VMM will create a copy of the source file so that you do not have to move or modify the original file.

In this scenario, the source file must meet the following requirements:

- You must leave the Administrator password blank on the VHD as part of the Sysprep process.
- You must install integration services on the virtual machine.
- You must use Sysprep to prepare the operating system for duplication.

Deploying from a Template

Deploying from a template creates 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 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 computer name, local administrator password, and domain membership. Guest operating system profiles allow you to preconfigure the roles and features that will deploy for Windows Server 2008 R2 and newer Windows Server operating systems.

The deployment process does not modify the template, which you can reuse multiple times. 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 virtual hard disk used with the template.
- You must leave the Administrator password blank on the VHD or VHDX 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 or VHDX by removing computer identity information. For Windows operating systems, you can prepare the VHD or VHDX by using Sysprep.

Deploying to and from the Virtual Machine Manager Library

During the virtual machine deployment process, you can deploy to a library. If you deploy a virtual machine from the Virtual Machine Manager library, you remove the virtual machine from the library and place it on the selected host. When using this method, you must provide the following details in the Deploy Virtual Machine Wizard:

- The host for deployment. The template that you use provides a list of potential hosts and their ratings.
- The path of the virtual machine files on the host.
- The virtual networks used for the virtual machine. You will see a list of existing virtual networks on the host from which to choose.

Configuring Virtual Machine Placement in VMM

Virtual machine placement enables VMM to evaluate hosts' capacity and then select the most appropriate virtualization host for deployment. The most recent VMM extends this capability with over one hundred virtual machine placement checks. It also adds support for custom placement rules.

VMM Managed Virtual Machine Placement

You can define placement rules on a host group level to manage virtual machine placement on specific hosts inside a host group. In general, VMM tries to recommend the most appropriate host for virtual machine placement by calculating host ratings. However, by specifying custom placement rules, you can define your own rules for placement or placement blocking.

When deciding whether to deploy a virtual machine to a host or private cloud, you can configure the Expected Utilization settings, which further refine host ratings based on anticipated resource utilization. With these settings you can adjust the following:

- CPU percentage expected utilization
- Disk, physical disk space, and expected disk I/O per second
- Network, expected utilization in megabits per second (mbps)

Custom Placement Rules

Custom placement rules are based on host and virtual machine custom properties. On each host, you can define values for 10 predefined custom properties, and you can define your own new custom properties and their values. Similarly, you can define custom properties for each virtual machine. By defining custom placement rules on a host group level, you can define a rule that is using a custom property as a condition for allowing or blocking virtual machine deployment on a host in a host group. For example, you can define a rule specifying that a specific custom property value must match on both the host and the virtual machine. If it does not, the virtual machine will not deploy.

When performing virtual machine placement, you should:

- Recommend the most appropriate host for virtual machine placement
- Create custom placement rules
- Remember that custom placement rules are based on custom properties of virtual machines and hosts

V2V Machine Conversions

VMM allows you to convert existing VMware virtual machines to virtual machines running on the Hyper-V platform in a process called *V2V conversion*. V2V conversion enables administrators to consolidate a virtual environment that runs various virtual platforms, without having to move data or rebuild virtual machines.

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 VMware ESX server hosts or in Virtual Machine Manager libraries. Although we call V2V a conversion, V2V is a read-only operation that does

V2V machine conversion options include:

- VMM built-in capability
- Microsoft Virtual Machine Converter
- Microsoft Virtual Machine Converter Plug-in for VMware vSphere Client
- Disk2vhd
- Windows PowerShell
- Tools other than Microsoft

not delete or affect the original source virtual machine. In addition, the term *conversion* refers only to the process of converting VMware virtual machines. We use the term *migration* 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 VMM V2V Wizard creates matches VMware virtual machine properties, including name, description, memory, and disk-to-bus assignments. In addition, you can use the new V2V Windows PowerShell® cmdlet when scripting V2V conversions. Other tools or methods for V2V conversion are:

- Microsoft Virtual Machine Converter. Microsoft Virtual Machine Converter is a free, easy to use, stand-alone solution accelerator tool that is wizard-driven. The tool will make a copy of the source virtual machine, while leaving the source intact.
- Migration Automation Toolkit. Migration Automation Toolkit is a collection of Windows PowerShell scripts used to automate conversion by utilizing the Microsoft Virtual Machine Converter.
- Microsoft Virtual Machine Converter Plug-in for VMware vSphere Client. The Microsoft Virtual Machine Converter Plug-in for VMware vSphere Client extends the vSphere Client context menu to make it easier to convert the VMware-based virtual machine to a Hyper-V-based virtual machine.
- Disk2vhd, dedicated tools other than Microsoft tools, backup, and mirroring technologies. Disk2vhd will work, but may require drivers and changes to the source machine to start successfully. Tools other than Microsoft tools may also work with varying degrees of manual intervention.



For more information about the Microsoft Virtual Machine Converter, go to:

<http://go.microsoft.com/fwlink/?LinkID=285296>

Considerations for V2V Conversions

When considering V2V conversions, remember that the source server is already virtualized. This will help you estimate performance trends when planning to accommodate the virtual machine appropriately.

Some considerations for V2V conversions are:

- Ensure that the tool used for the conversion process supports the operating system, including the patch level. Also, ensure that Hyper-V supports the source operating system. Ensure that your system meets all prerequisites.
- If you perform the V2V conversion to resolve a performance issue, confirm that the issue is resource-related and not configuration-related.
- For virtual machines running on other hypervisors, use the Migration Automation Toolkit with the Microsoft Virtual Machine Converter or the new V2V Windows PowerShell cmdlet.
- Use Microsoft Assessment and Planning Toolkit (MAP) to compare the existing virtual machines to the baseline and then downsize or upsize where appropriate.

When planning a V2V machine conversion:

- Ensure that the tool used for the conversion process supports the operating system, including the patch level
- Ensure that Hyper-V supports the source operating system and that all prerequisites are met
- If you perform the V2V conversion to resolve a performance issue, confirm that the issue is resource-related and not configuration-related

Lesson 4

Planning and Implementing Hyper-V Replica

Hyper-V Replica is a disaster recovery feature included in Hyper-V. You can use it to replicate a running virtual machine to a secondary location, and in Windows Server 2012 R2, you can extend the replication to a third location. While the primary virtual machine runs, the Replica machine is turned off. Hyper-V Replica updates regularly, and you can perform failover from the primary virtual machine to a Replica virtual machine, when necessary. You can perform failovers manually, whether they are planned or unplanned. Planned failovers are without data loss, whereas unplanned failovers can cause a loss of the most recent changes made in a certain period of time, up to five minutes by default. In this lesson, you will learn how to implement and manage Hyper-V Replica, including how to perform both a test failover and a planned failover.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe Hyper-V Replica.
- Plan Hyper-V Replica.
- Plan Hyper-V replication.
- Plan Hyper-V Replica failover.
- Manage Hyper-V Replica with Windows Azure™ Hyper-V recovery manager.

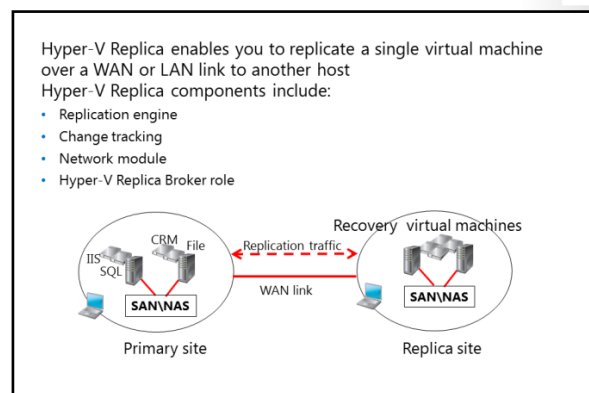
What is Hyper-V Replica

Hyper-V Replica enables virtual machines running at a primary site (or a location or host) to be replicated efficiently to a secondary site (or a location or host) across a WAN or local area network (LAN) link. With Hyper-V Replica, you can have two instances of a single virtual machine residing on different hosts: one is the primary (live) copy and the other is a Replica (offline) copy. These copies are synchronized in near real time and you can failover at any time.

If a natural disaster, power outage, or server failure causes a failure at a primary site, an administrator can use Hyper-V Manager to execute a failover of production workloads to Replica servers at a secondary location within minutes, thus incurring minimal downtime. Hyper-V Replica enables an administrator to restore virtualized workloads to a specific point in time, depending on the Recovery History selections for the virtual machine.

Hyper-V Replica technology consists of several components:

- Replication engine. This component is the core of Hyper-V Replica. It manages the replication configuration and tracks virtual machine and storage mobility events. The engine performs initial replication, delta replication, failover, and test failover operations.



- Change tracking. This component tracks changes on the primary copy of the virtual machine. It is designed to make the scenario work regardless of where the virtual machine virtual hard disk file(s) resides.
- Network module. This module provides a secure and efficient way to transfer virtual machine replicas between the primary host and the Replica host. Data compression is enabled by default. This communication is secure because it relies on HTTPS and certification-based authentication.
- Hyper-V Replica Broker role. Windows Server 2012 implemented this new role. It is configured in failover clustering. It provides Hyper-V Replica functionality even when the virtual machine being replicated is highly available and can move from one cluster node to another. The Hyper-V Replica Broker redirects all virtual machine-specific events to the appropriate node in the Replica cluster. The Hyper-V Replica 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 is executed.

The site configurations do not have to use the same server or storage hardware. However, it is important to ensure that sufficient hardware resources are available to run the replica virtual machine.

Windows Server 2012 R2 improves the Hyper-V Replica feature by adding the following enhancements:

- Replication frequency. In previous versions of Windows Server, Hyper-V Replica was set to a replication interval of five minutes, which you could not change. In Windows Server 2012 R2, you have the ability to set the replication interval to 30 seconds, five minutes, or 15 minutes. This means that you can configure your replication traffic based on your environment. However, keep in mind that replicating with a higher latency (such as 15 minutes) will generate more traffic when it happens.
- Extended replication. In Windows Server 2012, it is possible to have only one replica of an existing virtual machine. With Windows Server 2012 R2, you can replicate a single virtual machine to a third server. This means that you can replicate a running virtual machine to two independent servers. However, you cannot replicate from one server to two other servers. The server that runs an active copy of the virtual machine replicates to the Replica server, and the Replica server then replicates to the extended Replica server. You create a second replica by running the Extend Replication Wizard on a passive copy. In this wizard, you can set the same options that you configured when configuring the first replica.

Administrators can benefit from these improved features that optimize the usage of Hyper-V Replica and increase the availability of critical virtual machines.

Demonstration: Enabling Hyper-V Replica

In this demonstration, you will see how to enable Hyper-V Replica.

Demonstration Steps

1. On LON-HOST1 and LON-HOST2, configure each server to be a Hyper-V Replica server:
 - Use Kerberos (HTTP).
 - Enable replication from any authenticated server.
 - Create and use the folder **E:\VMReplica** as a default location to store Replica files. (Note: The drive letter might change depending on your host hardware configuration.)
2. Enable the firewall rule named **Hyper-V Replica HTTP Listener (TCP-In)** on both hosts.

Planning Hyper-V Replication

When planning to implement Hyper-V Replica technology, you must be sure to meet the following requirements:

- The server hardware supports the Hyper-V role on Windows Server 2012 or Windows Server 2012 R2.
- There is sufficient storage on planned primary and Replica servers to host the files that replicated virtual machines use.
- There is adequate network connectivity between the locations hosting the primary and Replica servers. This can be a WAN or LAN link.
- You have correctly configured firewall rules to enable replication between the primary and Replica sites (default traffic is going over TCP port 80 or 443).
- There is an X.509v3 certificate to support mutual authentication with certificates, if desired. Use this when virtualization hosts are not members of an AD DS domain.

When planning Hyper-V replication, ensure that:

- Sufficient storage exists on planned primary and Replica servers
- Adequate network connectivity exists between the locations hosting the primary and Replica servers
- Firewall rules are configured correctly to enable replication between the primary and Replica sites
- You have acquired and installed appropriate certificates, if using mutual authentication with certificates

You can deploy Hyper-V Replica on stand-alone servers or on servers that are part of a failover cluster, in which case you should configure the Hyper-V Replica Broker. Unlike failover clustering, a Hyper-V role is not dependent on AD DS. You can use it with Hyper-V servers that are stand-alone or that are members of different Active Directory domains, except when servers are part of a failover cluster.

To enable Hyper-V Replica technology, you should configure Hyper-V server settings first. In the Replication Configuration options, you should enable the Hyper-V server as a Replica server, and you should select authentication and port options. You should also configure authorization options. You can choose to enable replication from any server that authenticates successfully. This is convenient in scenarios where all servers are part of same domain. Alternatively, you can enter the 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 on a server level, you should enable replication on a virtual machine. During this configuration, you must specify both the Replica server name and the options for connection. When a virtual machine has more than one VHD, you can select which VHD drives you replicate. In addition, you can configure Recovery History and initial replication method. You also have the option of seeding content, which might be suitable when configuring replication between two sites that do not share a high-bandwidth link.

Specific to Windows Server 2012 R2, you can configure replication intervals; that is, for 30 seconds, five minutes, or 15 minutes. After you have configured these options, you can start replication. After you make the initial replica, in Windows Server 2012 R2, you can also make an extended replica to a third physical server running Hyper-V.



Note: Hyper-V Replica does not work between hosts with differing operating system versions. For example, you cannot replicate between Windows Server 2012 Hyper-V and Windows Server 2012 R2 Hyper-V.

Planning Hyper-V Failover

Hyper-V Replica supports three types of failovers: test failover, planned failover, and failover.

Test Failover

A test failover enables you to test a virtual machine on the Replica server while the primary virtual machine runs, without interrupting the replication. You can initiate a test failover on the replicated virtual machine, which will create a new checkpoint. You can use this checkpoint to select a recovery point, from which the new test virtual machine is created. The test virtual machine has the same name as the replica, but with - Test appended to the end. The test virtual machine does not start. It is disconnected by default to avoid potential conflicts with the running primary virtual machine.

After you finish testing, you can stop a test failover. This option is available only if test failover runs. When you stop the test failover, it stops the test virtual machine and deletes it from the Replica Hyper-V host. If you run a test failover on a failover cluster, you will have to remove the Test-Failover role from the failover cluster manually.

Planned Failover

You can initiate a planned failover to move the primary virtual machine to a Replica site, before site maintenance or before an expected disaster, for example. Because this is a planned event, there is no data loss, but the virtual machine will be unavailable for some time during its startup. A planned failover confirms that the primary virtual machine is turned off before executing the failover. During the failover, the primary virtual machine sends all the data that has not been replicated yet to the Replica server. Then the planned failover process fails over the virtual machine to the Replica server and starts the virtual machine at the Replica server. After the planned failover, the virtual machine will run on the Replica server and its changes are not replicated. If you want to establish replication again, you should reverse the replication. You will have to configure settings similar to when you enabled replication, and the existing virtual machine will be used as an initial copy.

Failover

In the event that an occurrence disrupts the primary site, you can perform a failover. You initiate a failover at the replicated virtual machine only if the primary virtual machine is unavailable or turned off. A *failover* is an unplanned event that can result in data loss. Data loss occurs because changes at the primary virtual machine have not replicated before the disaster happens (the Replication frequency setting controls how often changes are replicated). As with a planned failover, the virtual machine runs on a Replica server during a failover. If you need to start failover from a different recovery point and discard all the changes, you can cancel the failover. After you recover the primary site, you can reverse the replication direction to reestablish replication. This will remove the option to cancel failover.

Other Hyper-V replication–related actions include the following:

- Pause replication. This action pauses replication of the selected virtual machine.
- Resume replication. This action resumes replication of the selected virtual machine. This option is available only if you have paused replication for the virtual machine.
- View replication health. This action provides data about the replication events for a virtual machine.

- Test failover:
 - Is a nondisruptive test with zero downtime
 - Creates a new virtual machine in the recovery site from the Replica checkpoint; it is turned off and not connected
 - Allows you to stop a test failover
- Planned failover:
 - Failover moves a turned-off primary virtual machine to a Replica site
 - A primary virtual machine sends data that has not been replicated
 - Planned failover fails over the virtual machine to the Replica server and starts the Replica virtual machine
 - Replication should be reversed after the primary site is restored
- Failover:
 - Is performed in the event that an occurrence disrupts the primary site

- Extend replication. This action is available on replicated virtual machines. It is available on Windows Server 2012 R2 only, and it extends virtual machine replication from the Replica server to a third server, the extended Replica server.
- Remove recovery points. This action is available only during a failover. If you select it, all recovery points or checkpoints for a replicated virtual machine are deleted and their differencing virtual hard disks are merged.
- Remove replication. This action stops replication for the virtual machine.

Demonstration: Managing Hyper-V Replication and Failover

In this demonstration, you will see how to configure Hyper-V replication and failover.

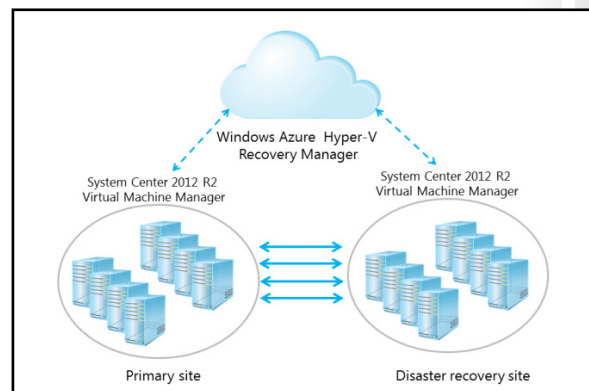
Demonstration Steps

1. On LON-HOST1, enable replication for the 20414C-LON-CORE virtual machine:
 - Use Kerberos Authentication (HTTP).
 - Select to have only the latest recovery point available.
 - Start replication immediately.
2. Wait for initial replication to finish, and make sure that the 20414C-LON-CORE virtual machine has appeared in the Hyper-V Manager console on LON-HOST2.
3. On LON-HOST2, view the replication health for 20414C-LON-CORE.
4. On LON-HOST1, perform a planned failover to LON-HOST2. Verify that 20414C-LON-CORE is running on LON-HOST2.
5. Shut down 20414C-LON-CORE on LON-HOST2.
6. On LON-HOST1, remove replication for 20414C-LON-CORE.

Managing Hyper-V Replicas with Windows Azure Hyper-V Recovery Manager

Windows Azure Hyper-V Recovery Manager manages the replication of large numbers of virtual machines between the primary site and disaster recovery sites. Recovery Manager manages the process of performing planned or unplanned failover of virtual machines from a primary to a disaster recovery site. Because a Windows Azure cloud hosts the Recovery Manager service, you can always access all of the components necessary to orchestrate the failover of virtual machines in one data center to another, even when one of the data center sites is

unresponsive. For example, if your organization's primary data center is in Sydney and the disaster recovery data center is in Melbourne, and the Sydney data center loses all power and connectivity suddenly, Recovery Manager, which exists independent of your organization's infrastructure, is able to



manage the failover process. Recovery Manager can ensure that Replica virtual machines in the Melbourne data center start in an orderly manner. Recovery Manager has the following requirements:

- System Center 2012 R2 VMM or VMM 2012 SP1 with cumulative update 3
- Windows Server 2012 with latest updates or Windows Server 2012 R2
- Creation of Recovery Manager vaults, including the installation of management certificates
- Recovery Manager provider installed on VMM servers
- VMM servers registered with Recovery Manager

Once you have installed and configured the prerequisite components, you can configure protection settings for VMM clouds, configure network mapping, and enable protection for specific Hyper-V virtual machines. Then you configure a recovery plan that details what steps Recovery Manager will take regarding the grouping of virtual machines when performing failover.

Lab: Planning and Implementing a Virtual Machine Deployment and Management Strategy

Scenario

Now that A. Datum Corporation has implemented the host, storage, and network components of the virtualization infrastructure, the next step is to plan and implement the virtual machine deployment.

One of the primary goals for the deployment project is to virtualize as many servers as possible. As part of the virtualization planning, you must identify which servers are candidates for virtualization, including which physical servers you can convert to virtual machines. In addition, you need to plan the virtual machine deployment and management strategy to simplify and standardize the deployment.

Objectives

After completing this lab, you will be able to:

- Plan virtual machines and service templates.
- Configure VMM profiles and templates.
- Deploy virtual machines by using VMM templates.

Lab Setup

Estimated Time: 60 Minutes

Virtual machines	20414C-LON-HOST1 20414C-LON-HOST2 20414C-LON-DC1 20414C-LON-VMM1
User name	Adatum\Administrator
Password	Pa\$\$w0rd

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

1. Ensure that you have started host computers LON-HOST1 and LON-HOST2 and signed in to each computer with the following credentials:
 - User name: **Administrator**
 - Password: **Pa\$\$w0rd**
 - Domain: **Adatum**
2. On LON-HOST1, open Hyper-V Manager from the Tools menu of the Server Manager console.
3. In Hyper-V Manager, click **20414C-LON-DC1**, and in the Actions pane, click **Start**.
4. Wait until the virtual machine starts, and then repeat step one for **20414C-LON-VMM1**.
5. Select **20414C-LON-VMM1**, and then in the Actions pane, click **Connect**.
6. Sign in using the following credentials:
 - User name: **Administrator**
 - Password: **Pa\$\$w0rd**
 - Domain: **Adatum**

Exercise 1: Planning Virtual Machine and Service Templates

Scenario

To simplify the deployment of virtual machines at A. Datum, you need to design virtual machine and service templates that you can use to deploy virtual machines. These templates should simplify the most common virtual machine deployments at A. Datum. You may wish to use a range of profiles. Initially, you will be working with the Toronto development team, who are handing over administration on their development virtualization hosts. They have provided configuration information based on a common setup for their development labs.

Development Service Template Reference Document	
Document Reference Number: BS0906/1	
Document Author	Ethan Rincon
Date	14th September
Requirements Overview Plan templates and profiles for the development team's virtualization objectives: Plan the details for a rapid deployment service for developers who frequently redeploy virtual machines for server and client operating systems.	
Task and Questions Fill in some of the attached profile items to help create the development service template. How many hardware, guest operating system, and SQL Server profiles will you require? Should you use service templates or virtual machine templates?	

Toronto Office					
Server name	Roles installed	Processor utilization	Processors and cores	Windows operating system version	Memory/ utilization
TOR-DDC1	AD DS	15%	One dual core processor	Windows Server 2008 R2 Standard SP1	2 GB/1 GB
TOR-SQLD1	SQL Server 2008 R2 Reporting Services Analysis Services	25%	One quad core processor	Windows Server 2008 R2 Standard SP1	6 GB/5 GB
TOR-WEBD1	IIS	15%	One dual core processor	Windows Server 2008 Enterprise SP2	4 GB/2 GB
TOR-CLD1	N/A	80%	One dual core processor	Windows XP SP3	2 GB/1 GB
TOR-CLD2	N/A	20%	One quad core processor	Windows 7 SP1	4 GB/2 GB

Sample:

Hardware profile		
Option/setting	Value	Additional value
Name	SQL Server	
Description	Sample development server	
Capability	Hyper-V	Citrix XenServer
Virtual processor	2	
Memory	2,048	Static

Hardware profile		
Option/setting	Value	Additional value
Name		
Description		
Capability		
Virtual processor		
Memory		

Hardware profile		
Option/setting	Value	Additional value
Name		
Description		
Capability		
Virtual processor		
Memory		

Hardware profile		
Option/setting	Value	Additional value

Hardware profile		
Option/setting	Value	Additional value

MCT USE ONLY. STUDENT USE PROHIBITED

Guest operating system profile		
Option/setting	Value	Additional value
Name	Development SQL Servers (WS2008R2SP1)	
Description	MSDN Windows 2008 R2	
Computer name	TOR-SQLD#	
Admin password	Pa\$\$w0rd	
Product key	XXXX-XXXX-XXXX-XXXX-XXXX	
Operating system	64-bit edition of Windows Server 2008 R2 Standard	
Roles		
Features	Microsoft .NET Framework 3.5.1	
Domain/workgroup	Adatum.com	
Domain user/password	Adatum\Administrator	Pa\$\$w0rd

Guest operating system profile		
Option/setting	Value	Additional value
Name		
Description		
Computer name		
Admin password		
Product key		
Operating system		
Roles		
Features		
Domain/workgroup		
Domain user/password		

MICROSOFT USE ONLY. STUDENT USE PROHIBITED

Guest operating system profile		
Option/Setting	Value	Additional value
Name		
Description		
Computer name		
Admin password		
Product key		
Operating system		
Roles		
Features		
Domain/workgroup		
Domain user/password		

SQL Server profile		
Option/setting	Option/setting	Option/setting
Name		
Description		
SQL Server deployment name		
Instance name		
Instance ID		
Product ID		
Run As		
Media source		
Administrators		
Security mode		
System administrator		
TCP/IP / Named pipes		
Service accounts		

The main tasks for this exercise are as follows:

1. Read the supporting documentation
2. Update the proposal document with your planned course of action
3. Examine the suggested proposals in the Lab Answer Key
4. Discuss your proposed solution with the class, as guided by your instructor

► **Task 1: Read the supporting documentation**

Read the documentation and scenario provided.

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

Answer the questions in the proposals section of the Development Service Template Reference Document.

► **Task 3: Examine the suggested proposals in the Lab Answer Key**

Compare your proposals with those in the Lab Answer Key.

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

Be prepared to discuss your proposals with the class.

Results: After completing this exercise, you should have planned a basic service template based on an existing infrastructure.

Exercise 2: Configuring VMM Profiles and Templates

Scenario

Now that you have designed the VMM profiles and templates, you will create a guest operating system profile, a hardware profile, and a SQL Server profile. Then you will create a virtual machine and service template.

The main tasks for this exercise are as follows:

1. Configure a guest operating system profile
2. Configure a hardware profile
3. Configure a SQL Server profile
4. Configure a virtual machine template
5. Configure a service template

► **Task 1: Configure a guest operating system profile**

Configure a guest operating system profile

1. In the VMM console, navigate to the **Library** workspace, expand **Profiles**, and then create a Guest OS Profile with the following properties:
 - Name: TOR-WEB OS Profile
 - Description: **Guest OS Profile for new development Web Server**
2. Click **Guest OS Profile**.

3. On the **Guest OS Profile** page, under **General Settings**, click **Identity Information**, and then configure the following settings:
 - Computer name: **TOR-WEB#**
 - Specify the password of the local administrator account: **Pa\$\$w0rd**
 - Operating System: **Windows Server 2012 R2 Standard**
 - Networking: **Domain/Workgroup**
 - Domain: **Adatum.com**
 - Domain credentials: **Specify credentials to use for joining the domain**
 - Domain user: **Adatum\Administrator**
 - Password: **Pa\$\$w0rd**

► **Task 2: Configure a hardware profile**

Configure a hardware profile

1. In the Library workspace of the VMM console, expand **Profiles**, and then click **Hardware Profiles**.
2. Create a new hardware profile with
3. the following settings:
 - Name: **WsStd2012R2**
 - Description: **Hardware Profile for new Windows Server 2012 R2 Servers**
4. On the **Hardware Profile** page, under **Compatibility**, click **Cloud Capability Profile**, and then select the **Hyper-V** check box.
5. In the General section, click **Processor**, and then enable **Allow migration to a virtual machine host with a different processor version**.
6. In the Memory section, set the **Virtual machine memory** option to **1024 MB**.
7. In the Network Adapters section, configure **Network Adapter 1** to connect to the **External VM** network.

► **Task 3: Configure a SQL Server profile**

Configure a SQL Server profile

1. In the Library workspace of the VMM console, expand **Profiles**, and then click **SQL Server Profiles**.
2. Create a SQL Server Profile with the following settings:
 - Name: **SQLDev1**
 - Description: **Template for new SQL servers**
3. On the **SQL Server Configuration** page, click **SQL Server Deployment**, and then configure the following settings:
 - Name: **SQLDev1**
 - Instance name: **MSSQLSERVER**
 - Instance ID: **DefaultInstance**
 - Installation Run As account: **Administrator**

4. On the **Configuration** page, configure the following settings:

- Media source: C:\SQLInstall
- SQL Server administrators: **Adatum\Administrator**
- Security mode: Windows Authentication
- Use TCP/IP for remote connections: **Enabled**
- SQL Server service Run As Account: **Administrator**
- SQL Agent Service Run As Account: **Administrator**
- Reporting Services service Run As Account: **Administrator**

► Task 4: Configure a virtual machine template

Configure a virtual machine template

1. On LON-VMM1, use File Explorer to copy Base14A-WS12R2.vhd from \\LON-HOST1\e\$\Program Files\Microsoft Learning\Base to \\lon-vmm1\MSSCVMLibrary\vhds\.



Note: You may need to substitute f\$ for e\$ in the above path depending on your environment.

2. In the VMM console, navigate to Library Servers\LON-VMM1.Adatum.com\MSSCVMLibrary\ VHDs.
3. Navigate to **Templates\ VM Templates** and create a new VM template.
4. Use **Base14-WS12R2.vhd** as the template source.
5. On the **VM Template Identity** page, configure the following options:
 - VM Template name: **Adatum Web Application Server**
 - Description: **Web Server hosting the Adatum Web Application**
6. On the **Configure Hardware** page, set the Hardware profile to **WsStd2012R2**.
7. On the **Configure Operating System** page, set the **Guest OS profile** to **TOR-WEB OS Profile**, and then select the Web Server (IIS) role check box.
8. On the **Configure Applications** page, set **Application profile** to **None – do not install any applications**.
9. On the **Configure SQL Server** page, set the **SQL Server profile** to **None - no SQL Server configuration settings**.

► Task 5: Configure a service template

Configure a service template

1. In the VMM console, click the **Library** workspace, expand **Templates**, and then click **Service Templates**.
2. Create a service template called **Adatum Web Service**, and assign the **Adatum Web Application Server** to the application tier.
3. Save and validate the service template.

Results: After completing this exercise, you should have configured Microsoft System Center 2012 R2 Virtual Machine Manager (VMM) profiles and templates.

Exercise 3: Implementing Hyper-V Replica

Scenario

One of the options for providing high availability for a virtual machine is to configure Hyper-V Replica. A. Datum has decided to implement Hyper-V Replica for one of the critical servers located in the Toronto data center. However, they first want to test Hyper-V Replica functionality, so they provided the 20414C-LON-CORE virtual machine as a test platform to evaluate Hyper-V Replica functionality.

The main tasks for this exercise are as follows:

1. Configure a replica on both host machines
2. Configure replication for the virtual machine
3. Validate a planned failover to the Replica site
4. Prepare for the next module

► Task 1: Configure a replica on both host machines

1. On LON-HOST1 and LON-HOST2, configure each server to be a Hyper-V Replica server:
 - Use Kerberos Authentication (HTTP).
 - Enable replication from any authenticated server.
 - Create and use folder **E:\VMReplica** as a default location to store Replica files. (Note: The drive letter might change depending on your host hardware configuration.)
2. Enable the firewall rule named **Hyper-V Replica HTTP Listener (TCP-In)** on both hosts.

► Task 2: Configure replication for the virtual machine

1. On LON-HOST1, enable replication for the 20414C-LON-CORE virtual machine:
 - Use Kerberos Authentication (HTTP).
 - Select to have only the latest recovery point available.
 - Start replication immediately.
2. Wait for initial replication to finish, and make sure that the 20414C-LON-CORE virtual machine has appeared in the Hyper-V Manager console on LON-HOST2.

► Task 3: Validate a planned failover to the Replica site

1. On LON-HOST2, view the replication health for 20414C-LON-CORE.
2. On LON-HOST1, perform planned failover to LON-HOST2. Verify that 20414C-LON-CORE is running on LON-HOST2.
3. On LON-HOST1, remove replication for 20414C-LON-CORE.
4. On LON-HOST2, shut down 20414C-LON-CORE.

► Task 4: Prepare for the next module

1. On LON-HOST1, remove the Toronto Logical Network virtual switch.
2. Do not revert the virtual machines, as you will need them for the next module.

Results: After completing this exercise, students will have implemented Hyper-V® Replica.

Module Review and Takeaways

Review Questions

Question: You must implement a Hyper-V extended replication. You have confirmed that the initial replication from host A to host B has completed. However, when you right-click the virtual machine, there is no option to extend replication. Why might this occur?

Question: Which two bare-metal system types can you preconfigure and provision by using the Virtual Machine Manager console?

Tools

- Microsoft Assessment and Planning Toolkit
- Microsoft Virtual Machine Converter
- Migration Automation Toolkit
- Microsoft Virtual Machine Converter Plug-in for VMware vSphere Client
- Disk2vhd

Best Practice:

Just as with most information technology (IT) projects, good planning and change control will help you achieve overall success during virtualization. Performance and resource monitoring are essential to maintaining the virtualization infrastructure. Therefore, your implementation and conversion strategy should include these factors to maintain performance levels and avoid scenarios where you run out of resources.

 **For more information and to download the Virtual Machine Servicing Tool 2012, go to:**

<http://go.microsoft.com/fwlink/?LinkID=285295>

Module 5

Planning and Implementing a Virtualization Administration Solution

Contents:

Module Overview	5-1
Lesson 1: Planning and Implementing Automation with System Center 2012	5-2
Lesson 2: Planning and Implementing System Center 2012 Administration	5-7
Lesson 3: Planning and Implementing Self-Service Options in System Center 2012	5-14
Lesson 4: Planning and Implementing Updates in a Server Virtualization Infrastructure	5-22
Lab: Planning and Implementing an Administration Solution for Virtualization	5-28
Module Review and Takeaways	5-38

Module Overview

This module will prepare you for designing an administrative model that you can use to manage virtualization using Microsoft® System Center 2012. You will learn about System Center 2012 components including System Center 2012 Virtual Machine Manager (VMM), System Center 2012 Orchestrator, and System Center 2012 App Controller. Use these components to delegate administrative functions, plan for basic self-service, and design and implement automation. The skills you will gain in this module are the foundation for operating an information technology (IT) infrastructure that is similar or equal to that of cloud computing.

Objectives

After completing this module, you will be able to:

- Plan and implement automation in System Center 2012.
- Plan and implement System Center 2012 administration.
- Plan and implement self-service options in System Center 2012.
- Plan and implement updates in a server virtualization infrastructure.

Lesson 1

Planning and Implementing Automation with System Center 2012

This lesson provides a high-level overview of the System Center 2012 options for automating administrative workflows for virtualization.

Lesson Objectives

After completing this lesson, you will be able to:

- Design automation in System Center 2012.
- Automate virtualization management with Orchestrator.
- Integrate Orchestrator and VMM.
- Create a basic runbook.

Designing Automation in System Center 2012

The following list summarizes further considerations for System Center 2012 R2 process automation:

- **Security.** Document the permissions that automation requires. Additionally, decide who will create, edit, and run runbooks. Also, decide who will audit what has been run and determine the permissions that this auditor will need.
- **Schedules.** You can automate many repetitive tasks, such as daily operational tasks. Automation can benefit many teams. Consider if your current manual processes fit the automation model to determine whether you will benefit from automation. For example, a manual 20-minute task could become an automatic two-minute task, which is a big benefit.
- **Infrastructure.** Consider what tasks you can automate, and then gather the information necessary to implement them manually. Then you can determine how you could automate and preapprove this process. For example, you could obtain a list of IP addresses that you can preapprove for use.
- **Business processes.** Find out who needs to approve deployments and costs, and then try to get preapproval so that automation can occur. Then determine whom you must notify regarding deployments.
- **Integration.** Integrate the System Center 2012 R2 components, and review what you can take advantage of in the integration packs. Use sample runbooks and the Cloud Services Process Pack, which provide many default offerings that you can customize.
- **Windows PowerShell®.** If the activities required do not exist for the automation runbooks, you need to deploy a third-party service. Then you can utilize Windows PowerShell scripting within the runbooks.

Consider the following for process automation:

- **Security:** consider whether you will use role-based access and ensure that you can audit activities
- **Schedules:** determine how and when automation should run
- **Infrastructure and planning requirements:** ensure that adequate disk space is available on storage area networks
- **Business processes:** identify the stakeholders and configure them accordingly
- **Windows PowerShell:** use cmdlets when activities do not achieve the workflow's requirements
- **Quotas:** use quotas to control resource allocation
- **Deployment order:** deploy the most cost effective runbooks first

- **Quotas.** If you have self-service offerings, you should set quotas at sensible levels, and then monitor capacity and reclaim unused resources.
- **Performance.** Consider the number and frequency of automation tasks that run. For example, if you have 50 runbooks that are monitoring various events, what is the performance impact on the runbook server and on each of the servers that you are automating? For the server that you are manipulating, consider the change from a manual process to an automated process, because automation can behave differently.
- **Process list.** Create a list of possible automation candidates, and then consider the time, effort, and complexity of each. Try to order the list so that the best return on investment (ROI) is at the top, and then determine how long a task takes and how much it costs.
- **Costs.** If your organization is new to process automation, you will probably have a list of possibilities for automation. Compare the cost of each of these possibilities to the return on investment. This can help you prioritize which runbooks to create first.

Getting Started with System Center 2012 - Orchestrator

<http://go.microsoft.com/fwlink/?LinkID=285301>

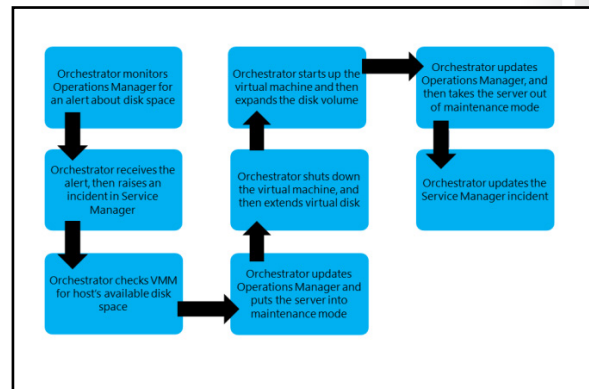
System Center 2012 Components

<http://go.microsoft.com/fwlink/?LinkID=285297>

Automating Virtualization Management with Orchestrator

Orchestrator has a series of integration packs that include virtualization activities that help you design workflows for administrators or self-service users. There are integration packs for each of the System Center 2012 R2 components and several third-party vendors' packs, including Hewlett-Packard, IBM, and VMware packs.

These virtualization activities enable you to create a series of tasks. For example, you could deliver a new virtual machine and then add the virtual machine to a System Center 2012 R2 Data Protection Manager (DPM) backup job.



If you cannot find the activity that you want to perform in an integration pack, then you can use the Run VMM Windows PowerShell Script activity to call a script and carry out any function that you want.

The four activities below are examples of runbook activities that come with the VMM integration pack. They are in the order that you might use for a simple workflow:

- Shut Down Virtual Machine
- Create Checkpoint
- Update Network Adapter
- Start Virtual Machine

When using Orchestrator with VMM, you can create runbooks that you can invoke manually or automatically. The following is an example of a runbook that you could create with steps that interact with several System Center 2012 R2 components. To use this runbook example, you would need to install the System Center 2012 R2 integration packs for System Center 2012 R2 Service Manager, System Center 2012 R2 Operations Manager, and VMM. Here are the sample runbook steps:

1. Orchestrator monitors Operations Manager for an alert regarding low disk space on a virtual machine.
2. Orchestrator sees this alert and raises an incident in Service Manager.
3. Orchestrator checks VMM to confirm that the host server has available disk space. If there is not enough space, it will update the incident and send a notification email. If there is enough space, the automation continues running predefined activities.
4. Orchestrator updates Operations Manager and places the virtual machine in maintenance mode to prevent further alerts regarding this server.
5. Orchestrator shuts down the virtual machine and extends the virtual disk.
6. Orchestrator starts the virtual machine and expands the disk volume.
7. Orchestrator updates Operations Manager and takes the server out of maintenance mode.
8. Orchestrator updates the Service Manager incident.

A service desk analyst could run this runbook interactively as one or several runbooks. For example, you could create the runbook so that it prompts the service desk analyst to run it at a specific time. Then you could configure it to pause until that time and continue when the analyst runs it. This is practical on a production system that you cannot shut down during working hours.

Demonstration: Integrating Orchestrator and VMM

In this demonstration, you will see how to deploy and configure the VMM integration pack.

Demonstration Steps

Install the System Center Integration Pack for VMM

1. On LON-OR1, run the following file: `\\lon-dc1\e$\labfiles\system_center_2012_orchestrator_integration_packs.exe`. Extract content in the default directory.
2. On LON-OR1, start **Deployment Manager**, and then register the integration pack at `\\lon-dc1\e$\Labfiles\SC2012_Virtual_Machine_Manager_Integration_Pack.oip`.
3. Deploy the System Center Integration Pack for System Center 2012 Virtual Machine Manager to LON-OR1.
4. Review the log entries, and then close the Orchestrator Deployment Manager.

Set the Windows PowerShell execution policy

1. On LON-OR1, on the task bar, right-click **Windows PowerShell**, and then under **Tasks**, click **Run as Administrator**.
2. At the Windows PowerShell prompt, type `set-executionpolicy remotesigned`, press Enter, type **Y**, and then press Enter.
3. Close the Windows PowerShell window.

4. On LON-VMM1, on the task bar, right-click **Windows PowerShell**, and then under **Tasks**, click **Run as Administrator**.
5. At the Windows PowerShell prompt, type **set-executionpolicy remotesigned**, press Enter, type **Y**, and then press Enter.
6. Close the Windows PowerShell window.

Enable Remote Management Trusted Hosts

1. On LON-OR1, open **gpedit.msc**.
2. Edit the Local Computer Policy\Computer Configuration\Administrative Templates\Windows Components\Windows Remote Management (WinRM)\WinRM Client\Trusted Hosts policy, enabling the policy and adding LON-VMM1 to the list of trusted hosts.
3. Close the Group Policy editor.

Configure the System Center Integration Pack for VMM

1. On LON-OR1, open **Runbook Designer**.
2. In the **Options** menu, click **SC 2012 Virtual Machine Manager**.
3. On the **Configurations** page, click **Add**, on the **Add Configuration** page, in the **Name** field, type **LON-VMM1**, and then click the browse (...) button.
4. On the **Item Selection** page, click **System Center Virtual Machine Manager**, and then click **OK**.
5. On the **Add Configuration** page, under **Properties**, in the **VMM Administrator Console** field, type **LON-VMM1**. In the **VMM Server** field, type **LON-VMM1**, and then in the **User** field, type **Adatum\Administrator**. Delete the text in the **Domain** field. In the **Password** field, type **Pa\$\$w0rd**.
6. In the **Authentication Type (Remote only)** field, click **Browse**, click **Negotiate**, and then click **OK**.
7. Click **OK**, and then on the **Configurations** page, click **Finish**.

In the Runbook Designer console, in the Activities section on the right, click **SC 2012 Virtual Machine Manager**, and then review the activities.

Demonstration: Creating a Basic Runbook

In this demonstration, you will review the Orchestrator components required for creating basic runbooks.

Demonstration Steps

Create a basic runbook

1. On LON-OR1, start **Runbook Designer**.
2. Create a folder named **20414 Runbooks** to store your new runbooks.
3. Create a new runbook in this folder named **VMM Library Monitor**.
4. In the right pane, under **Activities**, click **File Management**, and then drag the **Monitor Folder** activity to the center of the central pane.
5. Rename the **Monitor Folder** activity **VMM Library Monitor**.
6. Edit the **VMM Library Monitor** activity's properties. On the **General Information** page, in the **Description** field, type **This Runbook monitors the VMM library for new virtual hard disks**.
7. On the left, click **Details**, in the folder to monitor section in the **Path** field, type **\\LON-VMM1\MSSCVMLibrary**, and then click **Include sub-folders**.

8. In the File Filter section, click **Add**, on the **Filter Setting** page, click the **Name** drop-down list box, click **File Name**, in the **Value** field, type ***.vhd**, and then click **OK**.
9. On the left, click **Triggers**, and in the Trigger if section, select the check box next to **Number of files is**. Click the drop-down list box under **Number of files is**, select **greater than**, and then in the field next to **greater than**, type **0**.
10. On the left, click **Authentication**, in the **User name** field, type **Adatum\Administrator**, in the **Password** field, type **Pa\$\$w0rd**, and then click **Finish**.
11. On the right, under **Activities**, click **Notification**, and then click and drag the **Send Event Log Message** activity to the central pane and to the right of the VMM Library Monitor activity.
12. Connect the VMM Library Monitor activity to the **Send Event Log Message**. A link with an arrow should now appear between the two activities.
13. Right-click the link between the two activities, click **Properties**, when the **Link Properties** page appears, review the filter, and then click **Finish**.
14. Right-click the **Send Event Log Message**, click **Properties**, and on the **Details** page, in the Properties section, in the **Computer** field, type **LON-OR1**. In the **Message** field, type **A virtual hard disk file was created or updated in the LON-VMM1 library**. In the Severity section, click **Warning**, and then click **Finish**.
15. On the ribbon, click **Check In**, and then click **Runbook Tester**. In **Confirm Check Out** dialog box, click **Yes**.
16. In the Runbook Tester, click **Run**.
17. In File Explorer, navigate to **\\lon-vmm1\MSSCVMMLibrary\VHDs**.
18. Copy any of the **Blank Disk – Large.vhd** files, switch to the Runbook Tester, and then wait until the **Activity name Send Event Log Message** appears.
19. In **Event Viewer**, expand **Warning**, and then you should see an Event ID with the ID of **1** and a Source of **Orchestrator Runbook**. Review the event.
20. Close the Event Viewer, the File Explorer window, and the Runbook Tester.

Lesson 2

Planning and Implementing System Center 2012 Administration

This lesson provides a high-level overview of the options within System Center 2012 for delegating virtualization administration using VMM, App Controller, and Orchestrator.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe System Center 2012 Run As accounts.
- Delegate administrative options in VMM.
- Delegate administrative options in App Controller.
- Delegate administrative options in Service Manager
- Delegate administrative options in Orchestrator.
- Describe considerations for delegating administration in System Center 2012.
- Create a delegated administrator in VMM.

Overview of Run As Accounts in System Center 2012

In VMM, Run As accounts serve as containers for stored credentials that allow users to run administrative tasks. For example, a Run As account can store a user's credentials for connecting to Citrix XenServer, or for discovering and adding third-party storage.

Run As accounts have the following benefits:

- Streamline administrative tasks
- Reduce the likeliness of error
- Allow delegated administrators to perform tasks without requiring important passwords

Run As accounts store credentials used to run administrative tasks and include the following benefits:

- Streamline administrative tasks
- Reduce the likeliness of error
- Allow delegated administrators to perform tasks without requiring important passwords

Only administrators and delegated administrators can create and manage Run As accounts

Only administrators and delegated administrators can create and manage Run As accounts. Read-only administrators can only see the account names associated with Run As accounts that are in the scope of their user role. To create a Run As account, following this procedure:

1. Open the VMM console.
2. Click the **Settings** workspace.
3. On the ribbon, click **Create Run As Account**.
4. Type a name and description for the Run As credentials that you want to store.
5. Type the credentials in the **User name** and **Password** fields.
6. Remove the following selection **Validate domain credentials**. This is optional.
7. Click **OK** to create the Run As account.

Credential Security for Run As Accounts

VMM uses the Windows Data Protection Application Programming Interface (DPAPI) to provide data protection services at the operating system level, so that you can store and retire Run As account credentials.

During the installation of a VMM management server, you can use distributed key management to store encryption keys in Active Directory® Domain Services (AD DS), rather than storing the encryption keys on the computer on which you install the VMM management server. However, please note that if you install a clustered VMM management server, you must store the encryption keys in AD DS.

Delegating Administration Options in VMM

VMM provides four administrative profiles that you can assign to clouds. In VMM, a cloud is an on-premises logical grouping of resources, such as storage, networks, hosts, load balancers, and libraries. Once you create a cloud in VMM, you can delegate administrator permission based on responsibilities. When you create a cloud in VMM, you define the scope of administrator access, which you can group into departmental hosts or geographical groups.

You can assign the following administrative roles within VMM, including a read-only role that is suitable for auditors:

In VMM, you can assign the following administrative roles:

- Fabric administrator or Delegated administrator
 - Performs all tasks within an assigned scope
- Read-only administrator
 - Views properties, status, and job status within assigned host groups, clouds, and library servers
- Tenant administrator
 - Manages self-service users and creates and manage virtual machines and services
 - Assigns quotas on resources and virtual machines
- Application administrator or self-service user
 - Creates and manages virtual machines and services
 - Assigns quotas on resources and virtual machines

- Delegated administrator (fabric administrator). Delegated administrators can perform tasks on objects within their scope. Delegated administrators cannot modify VMM settings or add and remove members from the Administrator user role.
- Read-only administrator. Read-only administrators can view properties, status, and job status within their assigned host groups, clouds, and library servers. Read-only administrators cannot modify objects.
- Tenant administrator. Tenant administrators manage self-service users and virtual machine networks. They can create, deploy, and manage their own virtual machines and services by using the VMM console or App Controller. A tenant administrator user role specifies which tasks self-service users can perform on their virtual machines and services. Tenant administrators can place quotas on computing resources and virtual machines, effectively controlling the resources a self-service user can deploy.
- Application administrator or self-service user. Self-service users create, deploy, and manage their own virtual machines and services by using the VMM console or App Controller. A self-service user role specifies which actions the user can perform on his or her virtual machine and services. Application administrators can also place quotas on computing resources and virtual machines.

When you create user roles, you can assign groups of users from AD DS by selecting multiple resources, such as host groups, clouds, and library servers, and then assign Run As accounts. The ability to assign Run As accounts is very useful, especially if you want to delegate fabric administration and allow virtualization administrators to add storage from a storage area network (SAN), but do not wish to reveal the actual user name and password for the storage hardware.

To create a user role in VMM, follow this procedure:

1. In the VMM console, click **Settings**.
2. Expand **Security**, and then click **Create User Role**.
3. On the **Name and description** page, in the **Name** field, type the User Role name. In the **Description** field, type a description for the new user role, and then click **Next**.
4. On the **Profile** page, click a profile for the role, and then click **Next**.
5. On the **Members** page, click **Add**, type an account name, click **OK**, and then click **Next**.
6. On the **Scope** page, select the check boxes for the resources to include in the scope, and then click **Next**.
7. On the **Library Servers** page, click **Add**, select the Library servers to include with the role (to select more than one Library server, hold down the Shift key when selecting), click **OK**, and then click **Next**.
8. On the **Run As accounts** page, click **Add**, select the Run As accounts to include with the profile (to select more than one Run As account, hold down the Shift key when selecting), click **OK**, click **Next**, and then click **Finish**.

Delegating Administration Options in App Controller

App Controller has two user types: administrators and self-service users. The built-in administrators group can perform all administrative actions on all App Controller objects. This role is created during setup and applies to members of the local administrators group for the server on which you install App Controller.

Administrators can create one or more self-service roles and then delegate access to Windows Azure™ subscriptions. You can grant read-only permissions to the self-service role.

App Controller has the following user roles:

- Administrators
 - Built-in administrators can perform all actions on all App Controller objects
- Self-service user
 - Administrators can create self-service roles and then delegate access to Windows Azure subscriptions to users
 - You can designate self-service user roles as read-only

Roles created in App Controller are used only for delegation in Windows Azure



Note: Creating roles in the App Controller console is specifically for Windows Azure subscriptions or service provider connections. To create delegated user roles for a VMM connection, you must create the role in VMM.

To create a user role in App Controller, follow this procedure:

1. In the App Controller console, double-click **Settings**, and then double-click **User Roles**.
2. On the **User Roles** page, click **New**, and then in the General section, type a name and description. Optionally, you can click **Read only user role**.
3. In the Members section, click **Add**, in the **Enter the object name to select** field, type a user or group name in the format **domain\user**, and then click **Add**.
4. In the Scope section, select **Windows Azure Subscriptions** or **Service Provider Connections** for this role.

Delegating Administration Options in Service Manager

Service Manager can track all of your configuration items and provide service requests. It can enable you to manage business processes and release management for your virtualized infrastructure and applications.

When you are designing your environment's service management function, you may want to delegate different roles for different teams and users. You can assign security in Service Manager by using *role profiles*, which are collections of access rights that determine what actions or tasks a user may perform. You create a scope to define what and where these user roles can perform tasks.

The following is a list of some important Service Manager roles:

- Change managers. Create and edit change requests and activities.
- Release managers. Create releases and activities.
- Systems administrator. Configures and manages Server Manager.
- Incident resolvers. Resolve incidents.
- End users. Create work items.

To assign a user to a user role in Service Manager:

1. In the Service Manager console, click **Administration**.
2. In the Administration pane, expand **Security**, and then select **User Roles**.
3. In the User Roles pane, double-click **Advanced Operators**.
4. In the **Edit User Role** dialog box, click **Users**.
5. On the **Users** page, click **Add**.
6. In the **Select Users or Groups** dialog box, type the name of a user or group that you want to add to this user role, click **Check Names**, and then click **OK**.
7. In the **Edit User Role** dialog box, click **OK**.



Appendix A - List of User Role Profiles in System Center 2012 - Service Manager

<http://go.microsoft.com/fwlink/?LinkID=392391>



IPD Guide for Service Manager

<http://go.microsoft.com/fwlink/?LinkId=245470>

When you assign Service Manager roles, you must:

- Review the business and IT users' roles, and then map the most obvious ones to existing Service Manager roles, such as:
 - Change managers
 - Release managers
 - Systems administrator
 - Incident resolvers
 - End users
- Create new roles where required

Delegating Administration Options in Orchestrator

Orchestrator has two types of user roles, Runbook Authors and Runbook Operators. When you install Orchestrator, System Center 2012 prompts you to select a local Windows group or an Active Directory domain group to function as runbook authors. This group will enable you to assign permissions for the Runbook Designer and Deployment Manager, which users can utilize to deploy integration packs to runbook servers. Users in this group can:

- View, change, and run existing runbooks, and create new runbooks.
- Deploy new runbook servers.
- Deploy new Runbook Designers.
- Register and deploy integration packs.

Orchestrator also allows you to grant the following permissions at the runbook level, including:

- Read: Enable running of the runbook.
- Write: Enable running and changing of the runbook.
- Full Control: Enable reading, changing, and assigning of user permissions for the runbook.

Runbook Operators. Runbook Operators are able to read and invoke runbooks through the Runbook Designer or web console

You can assign permissions to operators by performing the following steps:

1. In the Runbook Designer, in the Connections pane, click the **Runbooks** folder.
2. In the Runbooks folder, right-click a runbook or a folder containing multiple runbooks, and then click **Permissions**. The Permission page opens. From the Permission page, you can click **Add** to give another user or security group access to the runbook and then select the user or security group from the local computer or from the domain.
3. From the Permissions page, you can assign the following runbook permissions:
 - Allow the user or security group to view and run the runbook by selecting the **Allow** check box next to **Read**.
 - Allow the user or security group to change the runbook by selecting the **Allow** check box next to **Write**.
 - Allow the user or security group to change permissions for the runbook by selecting the **Allow** check box next to **Full Control**.
4. When you have added and assigned permissions, click **OK** to close the Permissions for Runbook dialog box.



Orchestrator Security Planning:

<http://go.microsoft.com/fwlink/?LinkID=285302>

When using Orchestrator to delegate administration, remember that:

- You can grant Runbook Authors administrative access to the Runbook Designer and Deployment Manager
- You can grant Runbook Operators access to read and invoke runbooks through the Runbook Designer or web console
- You cannot grant Runbook Operators administrative access

Considerations for Delegating Administration in System Center

You should consider the following best practices carefully when you are delegating administration in System Center 2012:

- **Security.** Consider business needs and existing roles, and determine whether they map to System Center 2012 roles or whether you need to create new roles. You should document requirements carefully and then plan delegation of administration accordingly. If possible, create a lab environment to test the abilities of delegated administrators. Make sure that you define and set the scope of delegated administrators clearly in each application. Review the delegated administration periodically and remove any accounts that are not required. Also, delete any temporary access that you have granted that the user no longer requires. Additionally, define the logical boundaries and create host groups, clouds, and library resources. Delegate administrative responsibilities only when a user clearly requires them.
- **Infrastructure.** Consider what functions delegated administrators need to perform. Should they be able to add load balancers and SANs? Should they be able to create their own monitoring packs or just assign existing ones? You can use Run As accounts to mask crucial system credentials.
- **Training.** Provide appropriate training for staff to whom you grant administrative roles
- **Updates.** Always read the update rollup information. If new functionality becomes available, you need to determine whether it is available through existing privileges.

When delegating administration in System Center 2012, you must:

- Document requirements carefully and then plan delegation of administration accordingly
- Define and set the scope of delegated administrators clearly in each application
- Consider what functions delegated administrators need to perform
- Provide appropriate training for staff to whom you grant administrative roles

Demonstration: Create a Delegated Administrator in VMM

In this demonstration, you will create a delegated administrator in VMM.

Demonstration Steps

Configure delegated administration in VMM

1. On LON-VMM1, from the desktop, open the Virtual Machine Manager console.
2. In the **Connect to Server** dialog box, ensure that the **Use current Microsoft Windows session identity** check box is selected, and then click **Connect**. The Virtual Machine Manager (VMM) console opens.
3. In the VMM console, click the **Settings** workspace, and then on the ribbon, click **Create User Role**.
4. Enter **DevAdmin** as the name and **Development team administrators** as the description, and then click **Next**.
5. Click **Fabric Administrator (Delegated Administrator)**, and then click **Next**.
6. Click **Add**, type **Rob Cason**, click **OK**, and then click **Next**.
7. Click **All Hosts**, click **Next**, click **Add**, select **LON-VMM1.Adatum.com**, click **OK**, and then click **Next**.
8. Click **Add**, select **Administrator**, click **OK**, and then click **Next**. Review the summary, and then click **Finish**.
9. Close the Jobs window.

Lesson 3

Planning and Implementing Self-Service Options in System Center 2012

This lesson provides a high-level overview of the System Center 2012 options that you can use to provide self-service capabilities to end users and delegated administrators. The lesson presents the self-service options within the System Center 2012 components, including VMM, Orchestrator, and App Controller. You will learn how you can enable self-service for different groups within your organization, such as branch-office administrators or business-group administrators.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the self-service options in VMM.
- Create VMM private clouds.
- Describe the self-service options in Service Manager.
- Configure self-service options in Orchestrator.
- Describe the self-service options in App Controller.
- Describe the considerations for implementing self-service in System Center 2012.

Options for Self-Service in VMM

In System Center 2012 Service Pack 1 (SP1) and earlier versions, VMM included an optional self-service portal. However, the VMM self-service portal has been discontinued and replaced by App Controller.

VMM 2012 SP1 includes the following self-service options:

- VMM console. Users can deploy resources within their assigned scope by using the VMM console. You can deploy the console to self-service users or to a Windows Server® 2012 RemoteApp installation, and then publish it to end users. A benefit to this approach is that you may have fewer consoles to manage VMM cmdlets. You can use Windows PowerShell cmdlets to perform any steps that you perform by using VMM. This enables you to use your own custom applications that can call the Windows PowerShell cmdlets to deliver resources and administer a virtualized environment.
- App Controller. Users can access App Controller to see and manage scoped resources that you assign in VMM. Beginning with System Center 2012 SP1 VMM, App Controller is the replacement for the VMM self-service portal.

You can consume VMM resources by using the following self-service options:

- App Controller
- Service Manager/Cloud Services Process Pack
- The hosting company's bespoke portal, by using Service Provider Foundation or VMM Windows PowerShell cmdlets

System Center 2012 SP1 VMM and newer versions do not have their own self-service portal

- **Service Manager.** For more comprehensive self-service solutions, users can access the Service Manager self-service portal. The self-service portal in Service Manager is based on Microsoft SharePoint® 2013, and you can customize it easily. In addition to Service Manager's out-of-box portal, you can use the Cloud Services Process Pack, an infrastructure as a service-enabled (IaaS-enabled) process and automation pack that maximizes the capabilities of an integrated System Center deployment. The Cloud Services Process Pack and utilizes VMM, Service Manager, and Orchestrator.
- **Service Manager/Cloud Services Process Pack.** By deploying the Cloud Services Process Pack, you can make use of preconfigured forms and request and service offerings that provide automation and self-service for an IaaS computing model. The Cloud Services Process Pack provides a number of request offerings for Service Manager and the corresponding runbooks for Orchestrator, so you can provision these requests automatically by using VMM.
- **Hosting Resources.** Since the System Center 2012 SP1 Orchestrator release, Microsoft offers an optional technology known as the Service Provider Foundation. This technology hosts companies that have an existing self-service portal. It also provides the required integration with System Center 2012 components to deliver resources through existing portals. Additionally, the new Service Management Portal and application programming interface (API) enable you to offer a platform similar to that of Windows Azure, which allows website and virtual machine offerings.

The following table lists the actions that you can assign to a self-service user role in VMM.

Name	Description
Author	Authors virtual machine and service templates.
Checkpoint	Creates and manages virtual machine checkpoints.
Checkpoint (Restore Only)	Restores to but cannot create virtual machine checkpoints.
Deploy	Creates virtual machines and services from virtual hard disks (VHDs) or templates.
Deploy (From template only)	Creates virtual machines from templates only.
Local Administrator	Grants local administrator rights on virtual machines.
Pause and resume	Pauses and resumes virtual machines and services.
Receive	Receives resources from other self-service users.
Remote connection	Connects to virtual machines remotely.
Remove	Removes virtual machines and services.
Save	Saves virtual machines and services.
Share	Shares resources with other self-service users.
Shut down	Shuts down virtual machines.
Start	Starts virtual machines and services.
Stop	Stops virtual machines and services.
Store and redeploy	Stores virtual machines in the library and redeploys those virtual machines.

Demonstration: Creating VMM Private Clouds

In this demonstration, you will see how to create a cloud and delegate administration for clouds in VMM.

Demonstration Steps

Create a private cloud

1. On LON-VMM1, in the VMM console, click the **VMs and Services** workspace, click **Create Cloud**, and then create a cloud with the following settings:
 - Name: **London Development**
 - Description: **London Development Cloud**
 - Resources: **All Hosts**
 - Logical Networks: **External Network**
 - Load Balancers: **Microsoft Network Load Balancing (NLB)**
 - Port Classifications:
 - **Network load balancing**
 - **Medium bandwidth**
 - **Low bandwidth**
 - Read-only library shares: **MSSCVMMLibrary**
 - Capacity:
 - 8 virtual central processing units (CPUs)
 - 12 gigabytes (GB) memory
 - 250 GB storage
 - 15 Custom quota (points)
 - 8 virtual machines
 - Capability Profiles: **Hyper-V**

Assign access to the London Development cloud

1. Create a User Role with the following properties:
 - Name: **DevSS**
 - Description: **London Development Team Self-service Role**
 - Profile: **Application Administrator (Self-Service User)**
 - Members: **anat**
 - Scope: **London Development**
 - Networking: **External Network**
 - Permissions: **Select All**
 - Run As accounts: **Administrator**

Connect App Controller to VMM

1. On LON-VMM1 start **App Controller**.
2. Sign on as Adatum\Administrator with the password Pa\$\$w0rd.
3. Add a new VMM connection with the following properties:
 - Connection name: LON-VMM1.adatum.com
 - Description: London VMM Server access
 - Server name: LON-VMM1.adatum.com
4. Sign out of App Controller

Options for Self-Service in Service Manager

Service Manager Self-Service Portal

Service Manager provides a rich self-service web experience and a number of options that you can configure. The portal is comprised of a SharePoint website, Microsoft Silverlight® applications, and a web content server. You can deploy these components separately or on the same server.

- The Service Manager self-service portal allows you to publish service request offerings built from one or more System Center 2012 components to users
- This can include publishing service request offerings that allow self-service users to trigger Orchestrator runbooks



Self-Service Portal Deployment Scenarios for System Center 2012 - Service Manager

<http://go.microsoft.com/fwlink/?LinkID=392392>

With the portal in place, you enable end users to sign in via the web browser to make service requests, report issues, or search the organization's knowledge-based articles. You can create and customize these self-service items by integrating Service Manager with the other System Center 2012 products. For example, you can build request offerings that map to activities or tasks that you can deliver by using Orchestrator's automation capability. You might use this capability to deliver a new virtual machine or conduct a backup of a server.

Service Manager and the self-service portal enable you to automate the logging of processes for requests, notifications, and approvals for the deployment and management of virtualization resources. You will learn more about automation in the next lesson.

The role-based access and self-service portal in Service Manager enable you to:

- Report incidents. Instead of relying on service desk analysts to gather information for incident troubleshooting, you can build request forms that prompt the user to include all of the necessary information to submit a request.
- Use request offerings. A request offering enables you to offer a service, such as:
 - The creation of a user account.
 - The addition of a mailbox.
 - The deletion of a virtual machine.
 - The expansion of a virtual disk drive.
 - An increase of memory.
 - Access to a system or file.

These are just a few of the sample offerings that you can build and publish to self-service users by using the service catalog.

- Use service offerings. A service offering is a collection of request offerings, such as a backup service offering, which includes multiple backup requests. Service offerings can include requests for:
 - A one-off backup.
 - Addition of a server to a nightly backup rotation.
 - Inclusion of a weekly offsite tape backup.
 - Removal from a backup.
 - Restore of data or a system.
 - A report on the data that backups are consuming.

By grouping multiple requests into a service offering, you can provide self-service access, in addition to cost and service level agreement (SLA) information to users.

- Use IaaS offerings. The Cloud Services Process Pack for Service Manager provides request and service offerings and Orchestrator runbooks. We recommend deploying this offering, as it can reduce the time necessary to create your own customized request or service offerings.

 **Customizing the Self Service Portal:**

<http://go.microsoft.com/fwlink/?LinkID=392393>

 **Technical Documentation Download for System Center 2012 - Service Manager**

<http://go.microsoft.com/fwlink/?LinkID=392394>

Configuring Self-Service Options in Orchestrator

Orchestrator has a web-based console that requires Silverlight 4 or a newer version. This console provides administrators and operators with the ability to:

- View the list of runbooks and runbook servers.
- View the history and running status of runbooks.
- View a high-level definition of runbooks.
- Start and stop runbooks.
- Review events that runbooks and management servers create.

The self-service Orchestrator console provides a place for runbook operators and authors to:

- Start, stop, and review runbooks
- Access information on the history and current running status of runbooks
- Review events logged by the runbook and management servers

The Orchestrator Console offers a subset of the functionality that the Runbook Designer offers. You can access the Orchestrator Console by:

- Opening a browser, and then in the address bar, typing `http://<computer name>:<port number>`, where computer name is the name of the server where the web service is installed, and port number is the port number selected during the web service configuration. By default, the port is 82.
- Using the Runbook Designer Console, which you can access by clicking Orchestration Console on the toolbar.

Options for Self-Service in App Controller

App Controller gives the application owners a self-service experience across the VMM infrastructure. It provides a unified view that enables application owners to manage applications and service across private clouds and Windows Azure. With App Controller, users can manage application components in the context of the comprehensive service that it provides to the organization.

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

- Configure, deploy, and manage services through a single interface, while using a library of standard templates.
- Provide self-service application management, visibility, and control across both the Microsoft private cloud services and the Microsoft public cloud services, such as Windows Azure.
- Create, manage, and move services using a web-based interface that presents a customized view of resources based on your role in the organization. It also enables you to manage services rather than servers. This means that application owners can view virtual machines and both private and public cloud services. They can control components at each layer, track jobs, and maintain a detailed history of changes.
- Manage up to five VMM servers and 20 Windows Azure subscriptions.

- A self-service tool for end users that allows them to manage, deploy, and view resources in clouds
- Web-based, uses Internet Explorer to connect
- Can manage Windows Azure subscriptions
 - Allows end users to manage, deploy and view resources in Windows Azure
- Allows administrators to delegate authority to application administrators for certain cloud resources

Each App Controller instance supports up to 75 concurrent self-service users. App Controller also enables data center administrators to delegate authority to application owners. Predefined templates help ensure compliance with the company's IT standards and policies. By using App Controller, data center administrators can provide application owners with a customized, role-based view of private and public cloud services, and a view of consumed and available resources. In addition, application owners can customize all service components, including virtual machines, network resources, and load balancing.

The App Controller console is a portal that is accessible through a web browser. You should install Silverlight 4 before connecting to the App Controller portal. We also recommend that you add the App Controller portal to Trusted sites or intranet sites on the computer from where you are making a connection. To use single sign-on, you will have to add the portal to intranet sites in the Internet Explorer® settings, so that Internet Explorer will allow delegation of default credentials. If you do not want to log on by using the same credentials that you use to log on to your computer, you should not enable Windows Authentication on the /api virtual directory.

The default path for connecting to the App Controller console is <https://AppControllerServerFQDN/>, where FQDN is your server fully qualified domain name (FQDN). Ensure that the certificate for App Controller is issued to the same name that you are using to connect. After you connect to App Controller, you can use it to deploy and manage services, private clouds, and virtual machines. However, unlike the VMM console that provides a full set of options for these tasks, App Controller provides a limited set of options that focus on private clouds and services. For example, you can use App Controller to deploy new virtual machines and new services, but only those based on existing templates. Additionally, App Controller enables you to connect to and manage both public cloud and private cloud resources from the same place.

If you log on to App Controller as a VMM administrator, you will be able to create connections, view resource usage, and manage user roles. However, if you log on to the App Controller console as a self-service user, your set of available options will be limited to resources to which you have permissions. For example, in App Controller, on the Clouds tab, a self-service user can view both private clouds and public clouds to which he or she has appropriate permissions. On this tab, a self-service user will also see an option to deploy resources to cloud services. Based on templates provided in the VMM library that are available to the self-service user, it is possible to deploy a new service or virtual machine. Self-service users can also access a library view, where they can view available templates, shares, and other resources. From this view, it is also possible to deploy a new service or virtual machine. However, unlike VMM where a new virtual machine or service deployment requires several steps and several decisions, the App Controller process is quite straightforward. From App Controller, each self-service user can see his or her active jobs, job progress, and state.

Considerations for Implementing Self-Service in System Center 2012

Now that you are familiar with self-service and delegated administration in System Center 2012, you can begin mapping business requirements. VMM provides options for self-service functionality, and App Controller, Service Manager, and Orchestrator extend self-service capabilities even further. However, before you determine which of these components to deploy to offer self-service capabilities in your environment; you need to determine your business requirements.

When implementing self-service functionality, you should:

- Plan the number and type of self-service users
- Consider performance and availability
- Review resources, and apply quotas for individuals and groups
- Consider automation, notification, and approval
- Review permissions, assign only what is needed, and review your assignments periodically
- Consider backing up your portals, libraries, and the System Center 2012 components
- Consider testing a recovery plan

You should consider the following factors before implementing self-service capabilities:

- **Planning.** When you are deploying or upgrading applications, you must always plan the deployment first. Deploying self-service functionality is no exception to this rule. Because of the highly customizable and extensible nature of the System Center 2012 components, deploying self-service capabilities can take significant time and effort. If you want to begin planning a self-service deployment, you should consider the latest information and resources available on TechNet, such as product implementation and system requirements. Additionally, you should review any reference architecture and use the information and related job aids in the infrastructure and planning guides for each component that you plan to deploy.
- **Recommendations.** Look for guidance related to the Microsoft private cloud, which may be relevant to infrastructure and application self-service. Cloud Services Process Pack allows you to map your business requirements to the product capabilities, and consider a phased deployment if you must perform a large amount of configuration during your deployment. Measure the projected ROI for the key self-service areas, and determine whether what you want to achieve is available out-of-box. Furthermore, you should look for expert consultants or developers who can create and extend the management packs. Perform end-to-end planning for each self-service request, and consider using flow charts to establish both the business and technical processes you will need to employ. Lastly, consider the benefits of automating self-service requests.
- **Security.** You should include security concerns in the documented planning for your systems' deployment. Use role-based access control (RBAC) and assign the minimum required privileges for each role to meet your needs. Provide users with access or visibility to only the tools that they require. For example, you may not need to make every service portal catalog or offer visible and accessible to every user. If your organization undertakes security reviews of system access, you may want to review the self-service access periodically against your offerings to ensure that permissions have not changed. For example, make sure that users are not able to delete virtual machines if they should not have that capability.
- **Performance and resources.** Consider the performance and resource impacts that may occur when you move the deployment capability from the server team to a service desk team, or to developers and project and support professionals. You may have many people requesting and deploying services simultaneously, depleting resources rapidly. Consider how to moderate the deployment and lifecycle management of your systems. For example, if you allow users to utilize self-service backups of servers, and you automate agent deployment, you need to consider what might happen if someone requests a backup for every server on one specific night. This could cripple a network or cause storage resources to run out before other, more critical systems are backed up.
- **Backup.** Make sure that you back up your System Center 2012 product regularly.
- **Disaster recovery.** Consider the potential impact should your self-service offering become unavailable. Ensure that you have a plan to both recover the self-service offering and deploy services manually, if necessary.

Lesson 4

Planning and Implementing Updates in a Server Virtualization Infrastructure

Maintaining the infrastructure in VMM includes tasks such as adding new Windows Server 2012 Hyper-V® host servers and ensuring that the infrastructure components contain the latest approved software updates.

VMM integrates the functionality that Windows Server Update Services (WSUS) provides to ensure that all servers are compliant with the latest update baseline requirements.

Lesson Objectives

After completing this lesson you will be able to:

- Describe update management in VMM.
- Configure update management in VMM.
- Describe considerations for planning an update baseline.
- Describe considerations for deploying software updates.
- Manage updates using System Center 2012 Configuration Manager.

Overview of Update Management in VMM

Microsoft provides a number of solutions for deploying software updates and scanning computers for compliance. However, some network clients, such as cluster-based server nodes or other highly available server roles, present complexities that can make it difficult and time-consuming to maintain a standard update management process.

VMM integrates with WSUS to provide on-demand compliance scanning and remediation of servers that comprise an infrastructure, including Hyper-V hosts, library servers, Pre-Boot EXecution Environment (PXE) servers, and the VMM management server.

- The Update Server role manages updates for servers that make up the private cloud infrastructure
- A private cloud infrastructure includes:
 - Hyper-V hosts and clusters
 - VMM Library servers
 - PXE servers
 - VMM management servers
 - Infrastructure servers
- Updates in VMM do not apply to virtual machines

System Center 2012 R2 allows you to add a VMM agent to any Windows Server that meets the agents' prerequisites and to use this server to generate a compliance baseline. You can use the compliance baseline generated from a physical server to evaluate the compliance status of virtual machines.

Integrating WSUS with VMM also provides you with the ability to perform orchestrated updates of Hyper-V host clusters. When you remediate a host cluster, VMM places one cluster node at a time in maintenance mode, and then installs the approved updates. For clusters that support Live Migration, intelligent placement moves virtual machines off the cluster node that you are updating. If a cluster does not support Live Migration, then VMM saves the state of the virtual machines before updating the cluster node.



Note: You must have Windows Server 2008 R2 or Windows Server 2012 installed on a Hyper-V cluster node for live migration support.

You can use the Update Server role in VMM to manage more complex update tasks for servers in your private cloud infrastructure. These servers include:

- Hyper-V hosts
- Hyper-V clusters
- VMM Library servers
- PXE servers
- VMM management servers
- Infrastructure servers



Note: You use the Update Server role for updating servers that make up the VMM infrastructure only. You cannot use this solution to update VMM-managed virtual machines. For updating virtual machines, you should use a solution such as WSUS or Configuration Manager, which this module presented in earlier sections.

For all server roles within the VMM infrastructure, you scan against a baseline of approved updates to determine compliance status. For any servers that are noncompliant, you can perform update remediation tasks to install missing updates and restart servers, if necessary.

Configuring Update Management in VMM

The process for implementing update management within your VMM environment is as follows:

1. Use VMM to manage updates by first enabling update management. You enable update management by adding an existing WSUS server to VMM, or by installing a dedicated WSUS server and then adding the new update server to VMM.
2. Configure and manage update baselines, which specify a set of updates to be deployed to a host group, a stand-alone host, a host cluster, or a VMM server. The next topic provides more detail about update baselines. Once you add an update server, you can perform the following tasks from within the VMM console:
 - Perform on-demand synchronization of WSUS with Windows Update.
 - Configure proxy server name and port settings, which are required for connecting to the Internet for WSUS synchronizations.
 - Specify update classifications to synchronize.
 - Specify products to synchronize.
 - Specify supported languages to synchronize.
3. Start a scan to determine compliance status, once you assign an update baseline. During a compliance scan, WSUS checks each update in the assigned update baseline to determine whether the update is applicable and installed on the target server. The target server then reports a compliance status for each update.

Process for implementing update management in VMM:

1. Enable update management
2. Configure and manage update baselines
3. Start a scan to determine compliance status
4. Perform an update remediation
5. Specify update exemptions

4. Perform an update remediation to bring a managed server or Hyper-V host cluster into compliance. You can choose to remediate all update baselines assigned to a computer, all noncompliant updates in a specific update baseline, or a single update, as necessary.
5. Specify which update exemptions will prevent a specific update from being installed on a server. The computer will remain accountable for the assigned baseline, even if you exempt a specific update from being installed.

Considerations for Planning an Update Baseline

After you add a WSUS server to VMM to perform the Update Server role, your next step is to determine which updates you should install on each server within the private cloud infrastructure. VMM adds the updates that you select to an update baseline, against which each server scans. Then VMM remediates any server that does not meet the baseline, to install the missing updates immediately.

Once you determine which updates your VMM infrastructure servers require, you must create a list for VMM to use as a baseline. VMM uses the update baseline as the list, to which you can add or remove updates, as necessary.

- An update baseline is a set of required updates assigned to a scope of infrastructure servers within the private cloud
- If you move a host or host cluster to a new host group, the object will inherit the baseline associated with the target host group
- If you assign a baseline specifically to a stand-alone host or host cluster, the baseline will stay with the object when it moves from one host group to another

What Is an Update Baseline?

All updates from a specific product and category display within the VMM console when you synchronize with WSUS. To specify only the updates necessary for your requirements, you create an update baseline, which is a set of required updates that are assigned to a scope of infrastructure servers within the private cloud. You can assign an update baseline to the following:

- All hosts within all host groups
- Specific host groups
- A specific stand-alone server within a host group
- A specific host cluster within a host group
- Library servers
- PXE servers
- The VMM server
- The update server
- Infrastructure servers

Planning for Update Baselines

You should consider the following factors carefully when you are planning update baselines:

- If you assign a baseline to a host group, any host or host cluster within that group will be assigned to that baseline. If you move a host to a new host group, WSUS removes the original baseline, and the host will inherit the baseline associated with the new host group.
- If you assign a baseline specifically to a stand-alone host or host cluster, the baseline will stay with the object when it moves from one host group to another.

- When you first add the update server, two built-in update baselines are provided. The Sample Baseline for Critical Updates contains all of the critical updates that synchronize initially, and the Sample Baseline for Security Updates contains all of the security updates that synchronize initially. If you plan to use these built-in update baselines, you will need to maintain the updates as you perform subsequent synchronization tasks. You must assign computers to the baseline before you are able to use the baseline for compliance scanning and remediation.
- You can create a new update baseline that contains the updates that you require, in addition to those that you assign to the servers for which you want to maintain update compliance.

Considerations for Deploying Software Updates

If you design your update deployment plan properly, it should involve four broad phases:

- Determining the updates that you must deploy. Microsoft releases regular security bulletins that you can review to determine the scope and impact of any new updates.
- Testing updates prior to deployment. Organizations need to develop a test infrastructure that represents the diverse software configurations on organizational computers. You must determine whether you will test the infrastructure in a virtual environment, or whether you will test updates by using limited deployments to live computers.
- Deploying updates to all affected computers. Your update management plan should include determining how you will deploy updates to affected computers. You should choose a deployment method that minimizes effort while deploying updates reliably. You must have a plan that includes time for administrators to deploy updates that Microsoft releases regularly on the second Tuesday of each month. Additionally, you must have contingency plans for the rapid deployment of out-of-band updates, which Microsoft releases prior to the usual scheduled release of updates.
- Verifying successful update deployment. Your update management plan should include a post-deployment verification procedure, so that you can ensure that updates have deployed to the necessary computers. Additionally, your verification strategy must include use of an appropriate technology that can verify installation of updates, and a schedule for when you will check your network's computers to determine if they are missing critical updates.

An update deployment plan should include the following:

1. Specification of which updates to deploy
2. Test of updates prior to deployment
3. Deployment of updates to affected computers
4. Verification of successful update deployment

Ensure that your test group represents your production environment

As a network administrator, you must balance prompt deployment of updates that mitigate security risks with ensuring that updates do not conflict with existing configurations. When developing your update deployment strategy to support this requirement, consider the following guidelines:

- Understand the purpose and impact of the update. Microsoft releases notification bulletins prior to the release of updates, which include detailed information about the updates' content. Administrators should take time to study update bulletins and determine what impact the updates will have on their existing organizational configuration. This information includes the update's severity rating, such as Critical or Important, and the software that the update will affect.



Note: It is important to realize that an update that you deploy to your computers might cause unexpected consequences. Be very careful when you are considering deploying updates to server computers. If a server computer becomes unavailable because of update-related issues, the impact on your network infrastructure could be more significant than it would be on a single client computer.

- Test updates prior to general deployment. Your administrators should create a test group of computers or servers on which to test updates. Testing should:
 - Include groups that are representative of your organization's configuration as a whole.
 - Take sufficient time to ensure that no critical problems arise. However, the test period should not be so long that it delays update deployment substantially.
 - Not delay the deployment of critical updates significantly. Typically, there are approximately six to ten days between an update's release for an undisclosed vulnerability and an exploit code being released that targets that vulnerability.

Managing Updates by Using System Center 2012 R2 Configuration Manager

Organizations often use WSUS as the primary solution to stay current with and manage security updates, software upgrades, and service packs. Configuration Manager incorporates the WSUS engine, which provides the foundation for a comprehensive update solution for virtual machines that have the Configuration Manager clients installed.

You might consider using Configuration Manager to manage updates for virtual machines for the following reasons:

- The software updates feature in Configuration Manager scans, analyzes, and then deploys software updates to virtual machines with the Configuration Manager client installed
- Configuration Manager updating supports the following:
 - Seamless and flexible update deployment
 - Automatic deployment rules
 - Enhanced monitoring and reporting
 - Support for NAP

- Seamless and flexible update deployment. The software updates feature provides functionality to scan for required updates, analyze results, and deploy updates to client virtual machines. Additionally, you can integrate the software updates feature with other Configuration Manager functionality, such as compliance baselines. Compliance baselines allow you to determine whether a virtual machine meets a specific configuration standard, such as having specific updates installed.
- Collection-based maintenance windows help ensure that updates are applied only during approved maintenance periods, to meet organizational SLAs.
- Automatic deployment rules. Configuration Manager supports automatic approval and deployment of software updates based on automatic deployment rules. This ensures that updates that meet specific criteria are deployed as quickly as possible to client virtual machines, with minimal administrative effort.
- Enhanced monitoring and reporting. Configuration Manager provides extensive monitoring capabilities, such as detailed state messages, status updates, and alerts for key software update issues. Configuration Manager also provides an extensive number of reports to show deployment status and compliance statistics of updates throughout the organization.
- Support for Network Access Protection (NAP). With the integration of Windows Server 2012 NAP and the System Health Validator point site system role, you can define which software updates virtual machines require to connect to and communicate with the network resources. Configuration Manager software updates can provide remediation services to virtual machines that do not meet NAP policies, and deliver the software updates necessary to bring the client into compliance.

Lab: Planning and Implementing an Administration Solution for Virtualization

Scenario

A. Datum Corporation has deployed several virtualized servers. It has become increasingly clear that they will need to streamline their environment's management. Some business groups are requesting additional access to the virtual environment, and they require more control over the deployment and configuration of their virtual machines. Furthermore, they want to be able to initiate management tasks in the virtual environment.

As the number of physical hosts and virtual machine grows, administrative tasks are overwhelming the System Center 2012 administrators at A. Datum. Therefore, to decrease the time that administrators spend on regular tasks, they would like to delegate some of the most common and predictable tasks that they perform.

Objectives

After completing this lab, you will be able to:

- Configure process automation in System Center 2012.
- Plan administration delegation and self-service in System Center 2012.
- Configure delegated administration and self-service in VMM.
- Implement host updating in VMM.

Lab Setup

Estimated Time: 60 minutes

Virtual machines	20414C-LON-HOST1 20414C-LON-DC1 20414C-LON-VMM1 20414C-LON-OR1 20414C-LON-WSUS
User name	Adatum\Administrator
Password	Pa\$\$w0rd

For this lab, you will use the available virtual machine environment. This environment should still be running from the previous module. If it is, you will also need to start **20414C-LON-WSUS** and **20414C-LON-OR1**. If no virtual machines are active, you must complete the following steps:

1. On 20414C-LON-HOST1, start **Hyper-V Manager**.
2. In Hyper-V Manager, click **20414C-LON-DC1**, and then in the Actions pane, click **Start**.
3. In Hyper-V Manager, click **20414C-LON-DC1**, and then in the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Sign in by using the following credentials:
 - User name: **Administrator**
 - Password: **Pa\$\$w0rd**
 - Domain: **Adatum**
5. Repeat steps two to four for 20414C-LON-VMM1, 20414C-LON-WSUS, and 20414C-LON-OR1.

Exercise 1: Configuring Process Automation in System Center

Scenario

The development administration team has installed the Orchestrator console and they want you to configure it for creating workflows that will automate virtual machine deployment.

The main tasks for this exercise are as follows:

1. Install and configure System Center Integration Packs for VMM
2. Configure automation in Orchestrator

► Task 1: Install and configure System Center Integration Packs for VMM

Install the System Center Integration Pack for VMM

1. On LON-DC1, from the task bar, click **File Explorer**.
2. In File Explorer, browse to **E:\Labfiles**.
3. Double-click **System_Center_2012_R2_Integration_Packs.EXE**. Click **OK**, and then click **OK** again to close the Extraction complete message.
4. Close File Explorer.
5. On LON-OR1, move the pointer to the bottom left corner until the Start icon appears, click the **Start** icon, and then when the Start page appears, click **Deployment Manager**.
6. By using the Orchestrator Deployment Manager console, register and deploy the **Virtual Machine Manager Integration Pack**, which is located in **\\lon-dc1\efs\Labfiles\SC2012R2_Integration_Pack_for_Virtual_Machine_Manager.oip**.
7. Deploy the pack to LON-OR1.
8. Review the Log Entries, and then close the Orchestrator Deployment Manager.

Set the Windows PowerShell execution policy to RemoteSigned

1. On LON-OR1, on the task bar, right-click **Windows PowerShell**, and then under **Tasks**, click **Run as Administrator**.
2. At the Windows PowerShell prompt, type **set-executionpolicy remotesigned**, press Enter, type **Y**, and then press Enter.
3. Close the Windows PowerShell window.
4. On LON-VMM1, on the task bar, right-click **Windows PowerShell**, and then under **Tasks**, click **Run as Administrator**.
5. At the Windows PowerShell prompt, type **set-executionpolicy remotesigned**, press Enter, type **Y**, and then press Enter.
6. Close the Windows PowerShell window.

Enable Remote Management Trusted Hosts

1. On LON-OR1, move the pointer to the bottom left of the screen. When the Start icon appears, right-click the **Start** icon, and then click **Run**. On the **Run** page, in the **Open** field, type **gpedit.msc**, and then click **OK**.
2. On the left side, under **Local Computer Policy\Computer Configuration**, click to expand **Administrative Templates**, double-click **Windows Components**, and then double-click **Windows Remote Management (WinRM)**.
3. Click **WinRM Client**, and then on the right, under **WinRM Client**, double-click **Trusted Hosts**.

4. On the **Trusted Hosts** page, click **Enabled**, in the **TrustedHostList** field, type * and then click **OK**.
5. Close the Local Group Policy editor.
6. Repeat steps one through five on LON-VMM1.

Configure the System Center Integration Pack for VMM

1. On LON-OR1, click to the Start screen, and then click **Runbook Designer**.
2. On the menu at the top, click **Options**, and then click **SC 2012 Virtual Machine Manager**.
3. On the **Configurations** page, click **Add**, on the **Add Configuration** page, in the **Name** field, type **LON-VMM1**, and then click the browse (...) button.
4. On the **Item Selection** page, click **System Center Virtual Machine Manager**, and then click **OK**.
5. On the **Add Configuration** page, add the following details:
 - VMM Administrator Console: **LON-VMM1**
 - VMM Server: **LON-VMM1**
 - User: **Adatum\Administrator**
 - Domain: delete the text
 - Password: **Pa\$\$w0rd**
 - Authentication Type (Remote only): **Negotiate**
6. Click **OK**, and then on the **Configurations** page, click **Finish**.

► Task 2: Configure automation in Orchestrator

Create a basic runbook

1. On LON-OR1, open **Runbook Designer**.
2. In Runbook Designer, create a new folder named **20414 Runbooks**.
3. In this folder, create a new runbook named **VMM Library Monitor**.
4. Drag the Monitor Folder activity to the center of the central pane and rename it **VMM Library Monitor**.
5. Edit the properties of **VMM Library Monitor**. On the **General Information** page, in the **Description** field, type **This Runbook monitors the VMM library for new virtual hard disks**.
6. Click **Details**, in the folder to monitor section, in the **Path** field, type **\\LON-VMM1\MSSCVMLibrary**, and then click **Include sub-folders**.
7. In the **File Filters** section, click **Add**, on the **Filter Settings** page, click the **Name** drop-down list box, click **File Name**, in the **Value** field, type ***.vhd**, and then click **OK**.
8. On the left, click **Triggers**, in the **Trigger if** section, select the **Number of files is** check box, click the **Number of files is** drop-down list box, select **greater than**, and then in the **greater than** field, type **0**.
9. Click **Authentication**, in the **User name** field, type **Adatum\Administrator**, in the **Password** field, type **Pa\$\$w0rd**, and then click **Finish**.
10. Under Activities, click **Notification**, and then click and drag the **Send Event Log Message** activity to the central pane and to the right of the **VMM Library Monitor** activity.
11. Connect the **VMM Library Monitor** activity to the **Send Event Log Message** activity.

12. Right-click **Send Event Log Message**, click **Properties**, on the **Details** page, in the Properties section, in the **Computer** field, type **LON-OR1**, and then in the **Message** field, type **A virtual hard disk file was created or updated in the LON-VMM1 library**. In the **Severity** section, click **Warning**, and then click **Finish**.
13. On the ribbon, click **Check In**. On the ribbon, click **Runbook Tester**.
14. In the Runbook Tester, click **Run**.
15. Use File Explorer to navigate to `\\lon-vmm1\MSSCVMMLibrary\VHDs`.
16. Copy any of the **Blank Disk – Large.vhd** files.
17. Switch to the Runbook Tester. In the log section, wait until the Activity name **Send Event Log Message** appears.
18. In **Event Viewer**, expand **Warning**, and find an event with an Event ID with the ID of **1** and a Source of **Orchestrator Runbook**. Review the event.
19. Close the Event Viewer, File Explorer, and Runbook Tester.

Results: After this lab, you will have installed the Microsoft® System Center 2012 Integrations Pack for Virtual Machine Manager (VMM), created a basic runbook in the System Center 2012 Orchestrator Runbook Designer, and reviewed the Orchestrator web console.

Exercise 2: Planning Administrative Delegation and Self-Service in System Center 2012

Scenario

The Toronto-based development team at A. Datum would like to have more control over the virtual environment where their servers are deployed. They also want to collect advanced monitoring information and customize the monitoring configuration. They want to be able to create and manage workflow events related to their virtual environments. They also want to be able to delegate the responsibility for creating new virtual machines to users within their organization.

Virtualization Administration Strategy	
Document Reference Number: BS0907/1	
Document Author	Brad Sutton
Date	20th September
<ul style="list-style-type: none"> • Requirements Overview: • Plan a virtualization administration strategy to support the following objectives: • Provide the development group's administrative team with administrative control over their host groups and associated resources. • Allow the development team to create self-service accounts for developers and contractors to deploy machines rapidly. • Provide the development team with server monitoring to include the option to configure specific advanced settings. 	
<ul style="list-style-type: none"> • Additional Information: • The development team advises that the department is very dynamic and that they deploy virtual machines daily for internal systems development. 	

Virtualization Administration Strategy

- Currently, they have external staff on a short-term basis testing two large systems. The developers that maintain the virtual host servers normally are pressed for time due to their various administrative tasks. Their policy is to monitor and back up all virtual machines, including test systems.

Proposals:

1. Which administrative role would best meet the primary control objectives of the administrative teams in the development group?
2. Most developers will require the ability to create self-service accounts. Which administrative roles will be best for them?
3. What can you do to help reduce the administrative burden on the development administrators?

The main tasks for this exercise are as follows:

1. Read the supporting documentation
2. Update the proposal document with your planned course of action
3. Examine the suggested proposals in the Lab Answer Key
4. Discuss your proposed solution with the class, as guided by your instructor

► Task 1: Read the supporting documentation

Read the documentation and scenario provided.

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

Answer the questions in the proposals section of the Virtualization Administration Strategy document.

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

Compare your proposals with those in the Lab Answer Key.

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

Be prepared to discuss your proposals with the class.

Exercise 3: Configuring Delegated Administration and Self-Service in VMM

Scenario

To address the administration requirements for one of the business groups, you will implement delegated administration and self-service in VMM. You need to configure the delegated administration roles and the self-service roles, and then verify the configuration.

The main tasks for this exercise are as follows:

1. Configure a delegated administrator role in VMM
2. Configure self-service administration in VMM
3. Validate the configuration by using VMM
4. Validate the configuration by using App Controller

► Task 1: Configure a delegated administrator role in VMM

Create a private cloud

1. On LON-VMM1, from the desktop, open the Virtual Machine Manager console.

2. In the **Connect to Server** dialog box, ensure that the **Use current Microsoft Windows session identity** check box is selected, and then click **Connect**.
3. Click the **VMs and Service** workspace, and then create a cloud named **London Development**, with the description **London Development Cloud**.
4. Follow the Create Cloud Wizard, and then on the **Resources** page, select **London Hosts**.
5. On the **Logical Networks** page, select **External Network**.
6. On the **Load Balancers** page, select **Microsoft Load Balancing (NLB)**.
7. On the **Port Classifications** page, select **Network load balancing, Medium Bandwidth, and High Bandwidth**.
8. On the **Library** page, select **MSSCVMMLibrary**.
9. On the **Capacity** page, review the capacity options. Clear the check box next to each selected resource, and then assign the following:
 - 8 virtual CPUs
 - 12 GB memory
 - 250 GB storage
 - 15 quota points
 - 4 virtual machines
10. On the **Capability Profiles** page, select **Hyper-V**.
11. Review the **Summary** page, click **Finish**, and then close the Jobs window.

Configure delegated administration in VMM

1. In the VMM console, click the **Settings** workspace, and then create a User Role with the name **DevAdmin** and the description **Development team administrators**.
2. Follow the Create User Role Wizard, on the **Profile** page, select **Fabric Administrator (Delegated Administrator)**, and then on the **Members** page, add **Rob Cason**.
3. On the **Scope** page, select the **London Development** cloud and the **London Hosts** host group.
4. On the **Library servers** page, select **LON-VMM1.Adatum.com**.
5. On the **Run As accounts** page, select **Administrator** account.
6. Review the summary, click **Finish**, and then close the Jobs window.

► Task 2: Configure self-service administration in VMM

Configure self-service in VMM

1. In the VMM console, click the **Settings** workspace, and then create a User Role named **DevContractors** with the description **Development team contractors**.
2. Follow the Create User Role Wizard, and then on the **Profile** page, select **Application Administrator (Self-Service User)**.
3. On the **Members** page, add **Adam**.
4. On the **Scope** page, select **London Development**.
5. On the **Networking** page, select **External Network**.
6. On the **Resources** page, select the **Adatum Web Application Server**.

7. On the **Permissions** page, assign the following permitted actions:
 - **Deploy**
 - **Remote connection**
 - **Shut down**
 - **Start**
 - **Stop**
8. On the **Run As accounts** page, select **Administrator** Account.
9. On the **Summary** page, review the settings, click **Finish**, and then close the Jobs window.

► **Task 3: Validate the configuration by using VMM**

Verify delegation of administration

1. Open a new connection in the Virtual Machine Manager console.
2. Sign in to the Virtual Machine Manager console by using the following credentials:
 - User name: **Adatum\Rob**
 - Password: **Pa\$\$w0rd**
3. Create a new virtual machine in the London Development cloud. Follow the Create Virtual Machine Wizard, and then on the **Select Source** page, select **Create the new virtual machine with a blank virtual hard disk**.
4. On the **Specify Virtual Machine Identity** page, name the virtual machine **RobVM**.
5. On the **Configure Hardware** page, select a **Hyper-V** capability profile.
6. On the **Select Destination** page, select **Deploy the virtual machine to a private cloud**.
7. On the **Select Cloud** page, select **LON-HOST1.adatum.com**.
8. On the **Summary** page, create the virtual machine. Close the DevAdmin instance of the Virtual Machine Manager console.
9. In the Administrator instance of the Virtual Machine Manager console, click the **VMs and Services** workspace, and then click the arrow next to **Clouds**. You should see the London Development cloud.
10. Click the **London Development** cloud, and then on the ribbon, click **Overview**. Note that you can see User roles and Virtual Machine Owners. Confirm that you can see the DevAdmin and DevContractors roles. Click to expand these roles and view their assigned users. Review the details on this page.

► **Task 4: Validate the configuration by using App Controller**

Connect App Controller to VMM

1. On LON-VMM1, open **App Controller**.
2. On the **App Controller Credentials** page, sign in with the following credentials:
 - User name: **Adatum\Administrator**
 - Password: **Pa\$\$w0rd**

3. Add a new connection to **LON-VMM1.adatum.com**, with the **Description** set to **London VMM Server access** and the **Server name** set to **LON-VMM1.adatum.com**.
4. On the **Overview** page, in the top right corner of the browser, click **Sign Out**.

Verify self-service in App Controller

1. On the **App Controller Credentials** page, in the **User name** field, type **Adatum\Adam**, in the **Password** field, type **Pa\$\$w0rd**, and then click **Sign In**.
2. From the **Overview** page, deploy a new service from the **Adatum Web Service template**.
3. In the Service section, click **Configure**, on the **Properties of Adatum Web Service** page, name the service **Contractor Service**, and then add a cost center called **London Development**.
4. In the Instance section, click **Configure**, in the **Computer Name** field, type **LON-WEB1**, and then click **OK**.
5. Note that, by clicking **Deploy**, a new virtual machine would be deployed. This would take between 10 and 20 minutes.
6. Click **Cancel** to cancel the deployment.
7. Close the App Controller window.

Exercise 4: Implementing Host Updating in VMM

Scenario

A. Datum has deployed VMM to manage the Hyper-V host machines in its London data center. A. Datum administrators are planning to use VMM to manage the updates that they must apply to the Hyper-V hosts. Because Hyper-V hosts will control several critical resources in the virtual machines, it is vital that they are up-to-date with all security and functionality fixes.

The main tasks for this exercise are as follows:

1. Configure VMM integration with WSUS
2. Configure a software update baseline in VMM
3. Verify baseline compliance
4. To prepare for the next module

► Task 1: Configure VMM integration with WSUS

1. On LON-VMM1, switch to the Virtual Machine Manager console, where you are signed in as Administrator, and then click the **Fabric** workspace.
2. In the navigation pane, expand the **Servers** node, expand **Infrastructure**, and then click **Update Server**.
3. Add an update server with the following configuration:
 - Computer name: **LON-WSUS**
 - TCP/IP port: **8530**
 - User name: **Adatum\Administrator**
 - Password: **Pa\$\$w0rd**
4. In the Jobs window, select the **Add Update Server** job. On the **Summary** and **Details** tabs, monitor the status of the configuration job.
5. When the job displays as **Completed w/info**, close the Jobs window.

► Task 2: Configure a software update baseline in VMM

1. On LON-VMM1, in the Virtual Machine Manager console, click the **Library** workspace.
2. In the navigation pane, expand **Update Catalog and Baselines**, and then click **Update Catalog**.
3. Verify that various updates display.
4. Create a new update baseline with the following settings:
 - Name: **Server Baseline**
 - Update: **Update for Windows Server 2012 R2 (KB2883200)**
 - Assignment Scope:
 - Library Servers: **LON-VMM1.Adatum.com**
 - Update Server: **LON-WSUS.Adatum.com**
 - VMM Server: **LON-VMM1.Adatum.com**
5. Verify that **Create new baseline** has completed successfully.

► Task 3: Verify baseline compliance

1. On LON-VMM1, click the **Fabric** workspace.
2. In the navigation pane, expand **Servers**, expand **Infrastructure**, and then click **Library Servers**.
3. On the ribbon, click **Compliance**.
4. In the results pane, note the compliance and operational status of **LON-VMM1.Adatum.com**.
5. Scan **LON-VMM1.Adatum.com** and verify its compliance status. Note: It should be compliant. If the compliance status is reported as **Non-Compliant**, then right-click **LON-VMM1.Adatum.com** and select **Remediate**, in the Update Remediation window. Select the **Do not restart servers after remediation** check box, and then click **Remediate**. After a couple minutes, the compliance status should change to **Compliant**.

► Task 4: To prepare for the next module

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

1. On the host computer, start Hyper-V Manager.
2. On the Virtual Machines list, right-click **20414C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps two and three for 20414C-LON-VMM1, 20414C-LON-SVR1, 20414C-LON-SVR2, 20414C-LON-WSUS, and 20414C-TOR-SVR1.

Results: After completing this exercise, you will have configured VMM host updating.

Question: If you receive a call from the Toronto developers informing you that there are no resources left, where can you see who is consuming all of the resources?

Question: Can Adam deploy a virtual machine?

Module Review and Takeaways

Review Questions

Question: What are some of the benefits of automation and self-service?

Module 6

Planning and Implementing a Server Monitoring Strategy

Contents:

Module Overview	6-1
Lesson 1: Planning Monitoring in Windows Server 2012	6-2
Lesson 2: Overview of Operations Manager	6-10
Lesson 3: Planning and Configuring Monitoring Components	6-25
Lesson 4: Configuring Integration with VMM	6-32
Lab: Implementing a Server Monitoring Strategy	6-37
Module Review and Takeaways	6-42

Module Overview

Monitoring is an essential part of server management and maintenance. If you monitor your servers, you can often identify potential problems and resolve them before there is any impact on users. In this module, you will see how to use the monitoring tools included in Windows Server® 2012 and Microsoft® System Center 2012 – Operations Manager.

Objectives

After completing this module, you will be able to:

- Plan monitoring in Windows Server 2012.
- Describe Operations Manager.
- Plan the configuration of monitoring components.
- Configure integration of Operations Manager and the Microsoft System Center 2012 - Virtual Machine Manager (VMM).

Lesson 1

Planning Monitoring in Windows Server 2012

A number of monitoring tools are included in the Windows Server 2012 operating system. These tools are useful for monitoring one or more servers. You can use tools such as Event Viewer, Performance Monitor, Reliability Monitor, Resource Monitor, and Windows Server 2012 Server Manager to ensure that servers are performing as expected. They can also be used for troubleshooting.

There are special considerations for monitoring servers in a virtualized environment. You need to be aware of how virtualization affects the data that you can collect. By using monitoring tools, you can ensure that you have accurate information.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the monitoring tools in Windows Server 2012.
- Monitor multiple servers by using Server Manager.
- Develop a monitoring strategy.
- Describe considerations for monitoring server roles.
- Describe options for monitoring a server-virtualized environment.
- Describe considerations for monitoring virtual machines.
- Implement a Windows PowerShell® Desired State Configuration (DSC) solution.
- Implement a Microsoft System Center 2012 global monitoring solution.

Overview of Monitoring Tools in Windows Server 2012

Monitoring server performance is an important way to identify potential resource issues before they affect application performance. When you monitor resources, such as usage over time, you can identify trends and predict when to update servers with additional resources. Also, when you monitor resources you can determine when services need to be moved to new servers. Failure to monitor server performance can result in reduced user productivity. The most commonly monitored resources are processor utilization, disk activity, memory utilization, and network utilization.

Windows Server 2012 includes the following tools for monitoring performance:

- Event Viewer
- Resource Monitor
- Performance Monitor
- Reliability Monitor
- Task Manager

Windows Server 2012 includes the following tools to monitor server performance:

- **Event Viewer.** You can use this tool to view event logs, which are a valuable source of information for monitoring and troubleshooting Windows Server 2012 and applications. The event logs contain informational, warning, and error events. You can use the contents of the events to identify the source of a problem. You can use subscriptions in Event Viewer to centralize the collection of events from multiple servers.

- **Resource Monitor.** You can use this tool to monitor the current state of resource utilization on your servers. You can monitor CPU utilization, disk activity, memory utilization, and network utilization. Resource utilization is identified for each process that is running on the computer. This tool provides information similar to performance data from Task Manager, but the information is much more detailed.
- **Performance Monitor.** You can use this tool to monitor performance counters that are installed on the server. This tool provides much more detailed information about system activity than Resource Monitor. It also has the ability to record logs for evaluation at a later time, and trigger alerts when counters are outside of acceptable boundaries. You can also use built-in reports for System Diagnostics and System Performance to perform troubleshooting.
- **Reliability Monitor.** You can use this tool to view an overview of system stability, problem history, and software installation. This allows you to view when stability problems began to occur, and correlate them with any changes to your system, such as new software installations or driver updates that may be causing the issue.
- **Windows Task Manager.** You can use this tool to perform real time monitoring of processes and services running on a server. The performance tab offers a graphical view of CPU, memory, and networking resources, as well as the ability to open the Resource Monitor for more in-depth analysis.

Monitoring Multiple Servers by Using Server Manager

Server Manager is redesigned for Windows Server 2012, and is now the central tool for managing servers running Windows Server 2012. Two significant changes that simplify administration are the ability to remotely manage individual servers the ability to manage multiple servers as a group from a single location. By default, Server Manager automatically opens when an administrator signs on and then displays a rollup status of performance, services, and events. Any issues relating to a role or service will display as red text and each role and service has a hyperlink that you can select to provide additional details.

You can use Server Manager to:

- Monitor one or more servers from a central console
- View Best Practice Analyzers
- View Event Logs

The two items in Server Manager that are useful for monitoring are:

- **Best Practice Analyzers.** Windows Server 2012 includes Best Practice Analyzers for server roles. You can use the Best Practice Analyzers to identify whether server roles are configured correctly or experiencing issues. As part of the analysis, event logs are scanned for relevant information.
- **Event Logs.** You can use Server Manager to view event logs from the local server or groups of servers. Event logs also display by server role. You can use the event logs to identify errors or warnings that are being generated by applications or server roles.

You may be able to identify potential problems and performance issues before they affect users if you regularly use Best Practice Analyzers and review Event Logs. For example, hardware events like a degraded disk drive or memory module will prominently display with red text on the main page of Server Manager.

Developing a Monitoring Strategy

As a best practice, you should develop a monitoring strategy so that you can identify performance and configuration issues before they begin affecting users. Your monitoring strategy should include the following:

- The servers that you should monitor.
- The applications that you should monitor.
- Performance characteristics to monitor.
- Roles and services running on the server.

One of the biggest monitoring concerns that administrators have is how to analyze the large amount of generated information. It is not feasible to monitor every possible characteristic of a server. Instead, you should identify the specific information that you need to identify performance concerns accurately. This includes identifying which resources are of greatest concern for a particular server role.

When monitoring, you also need to understand the normal performance of a server or application. The normal performance of a server or application is referred to as a *baseline*. A baseline helps you identify what has changed when performance problems occur. For example, if the baseline for a server indicates that it normally uses 8 gigabytes (GB) of memory, you would investigate memory utilization when reports indicate that the same server is using 12 GB of memory.

A monitoring strategy needs to include:

- Servers to monitor
- Applications to monitor
- Performance characteristics to monitor
- Roles and services on a server

Optimize monitoring by:

- Limiting data gathering to only relevant information
- Identifying a baseline

Options for Monitoring a Server Virtualized Environment

Virtualized environments provide the flexibility of running multiple virtual machines with different operating systems on the same physical computer. This allows for maximizing your hardware, and even lowering energy costs. However, virtualization also introduces some complexity when it comes to monitoring your virtualized environment.

For monitoring purposes, you can divide a virtualized environment into two separate groups:

- Virtualization hosts.
- Virtual machines.

• Use the following counters to monitor processor usage on a Hyper-V host:

- Hyper-V Hypervisor\Partitions
- Hyper-V Hypervisor\Logical Processors
- Hyper-V Hypervisor Logical Processor\% Total Run Time
- Hyper-V Hypervisor Logical Processor\% Hypervisor Run Time
- Hyper-V Hypervisor Root Virtual Processor
- Hyper-V Hypervisor\Virtual Processors
- Hyper-V Hypervisor Logical Processor\% Guest Run Time

• Monitor the following memory-related counters for your virtualization host:

- Memory\Available Mbytes
- Hyper-V Dynamic Memory Balancer\Average Pressure

Virtualization Hosts

Your virtualization environment will be composed of one or more servers running the Hyper-V® role in Windows Server 2012. You can use Performance Monitor to monitor memory, network, and disk utilization on your host servers in the same way you monitor any given server. However, you can measure processor utilization differently in a virtualized environment, because the processors are used by the host operating system and the guest operating systems, which run in the virtual machines that the host supports. You use the following counters to monitor processor usage on a Hyper-V host:

- Hyper-V Hypervisor\Partitions. Monitors the number of virtual machines.
- Hyper-V Hypervisor\Logical Processors. Monitors the number of logical processors.
- Hyper-V Hypervisor\Virtual Processors. Monitors the number of virtual processors.

- Hyper-V Hypervisor Logical Processor\% Total Run Time. Monitors the total nonidle time of the logical processors. This value should be below 75 percent.
- Hyper-V Hypervisor Logical Processor\% Hypervisor Run Time. Monitors the nonidle time of the logical processors for the host operating system only.
- Hyper-V Hypervisor Root Virtual Processor. All counters in this object measure the CPU usage for the host operating system only.

You also should consider monitoring the following memory-related counters for your virtualization host:

- Memory\Available Mbytes. Ensure your virtualization host has at least 10 percent free memory at any given time. If the host's memory falls below that reading, consider adding memory to the virtualization host or reducing its load.
- Hyper-V Dynamic Memory Balancer\Average Pressure. This counter measures the percentage of memory available to a given virtual machine as compared to the amount of memory that the same virtual machine requires. This means that a virtual machine with an average pressure of 100 has the exact amount of memory it needs, whereas a virtual machine with an average pressure of 120 needs 20 percent more memory than what you allocate. Ideally, you should have a value of 80 or less for this counter.

Virtual Machines

Your virtual machines compete for host resources, and a single bottleneck can undermine their performance. Resource bottlenecks are not always obvious, and monitoring virtualization host performance helps identify the sources of bottlenecks that need correction.

When monitoring the performance of virtual machines, you should rely on tools designed for virtual environments. Because many tools, such as Windows Performance Monitor, are not aware of the underlying virtualization layer, they can provide inaccurate results on certain counters and measurements. Because of this, it is a best practice to monitor the virtualization host rather than the individual guests.



Reference Links:

You can review the Performance Tuning Guidelines for Windows Server 2012 document at <http://go.microsoft.com/fwlink/?LinkID=285313>.

Considerations for Monitoring Virtual Machines

Organizations typically monitor virtual machine usage for two reasons:

- Performance Monitoring. You can measure usage of network, memory, processor, disk, and other resources to ensure the virtual machine is performing within a predefined baseline.
- Resource Metering. You can measure bandwidth, memory, processor, and storage utilization for resource planning, and accounting purposes.

Performance Monitoring Counters

- Hyper-V Hypervisor Logical Processor\% Guest Run Time
- Hyper-V Hypervisor Virtual Processor

Resource Metering cmdlets

- Enable-VMResourceMetering
- Disable-VMResourceMetering
- Reset-VMResourceMetering
- Measure-VM
- Get-VM

Performance Monitoring

Performance monitoring within a virtual machine is similar to performance monitoring on a computer that is running the Windows® operating system. You can retrieve memory, disk, and network utilization information by using performance monitoring. However, if you use the same technique to gather processor-utilization data, the data that you collect may not be accurate. When you monitor from within a virtual machine, you do not see information related to the overall virtualization host.

If you have a virtual machine that is running a processor-intensive application that displays 100 percent CPU utilization in Task Manager, this indicates that the virtual machine is using all of the CPU resources allocated to it. If you start a second machine on the same virtualization host that has 100 percent CPU utilization, the resources that the second virtual machine is using are removed from the first virtual machine, and they will share the host's processing power equally. Even though Task Manager shows a consistent value, you must remember that you will reduce the real processing power of the first virtual machine by starting the second virtual machine.

To monitor processor usage within a virtual machine, you need to use Performance Monitor in the host operating system, and then monitor the following counters:

- Hyper-V Hypervisor Logical Processor \% Guest Run Time. This counter monitors the nonidle time of a guest operating system. This value should be below 75 percent.
- Hyper-V Hypervisor Virtual Processor. All counters in this object measure the CPU usage for the guest operating system only.

Resource Metering

Metering is very common for Information Technology (IT) organizations that report resource consumption by business unit within the company they support and/or for companies that provide hosted services to their customers. With metering, companies and business units pay only for the resources they use.

Whichever scenario your company requires, Hyper-V in Windows Server 2012 provides a feature called Resource Metering that allows you to monitor resource consumption per virtual machine over time. The data that you collect remains with the virtual machine, even if you move the virtual machine using live, offline, and storage migration.

Resource Metering collects the following data from each monitored virtual machine:

- Average CPU usage, measured in megahertz (MHz) over a specific period.
- Average physical memory usage, measured in megabytes (MB).
- Minimum memory usage (lowest amount of physical memory).
- Maximum memory usage (highest amount of physical memory).
- Maximum amount of disk space allocated to the virtual machine.
- Total incoming network traffic, measured in megabytes, for a virtual network adapter.
- Total outgoing network traffic, measured in megabytes, for a virtual network adapter.

You can configure Resource Metering and retrieve the data collected by using the Windows PowerShell command-line interface. The following cmdlets are used for Resource Metering:

- **Enable-VMResourceMetering.** Enables resource metering for a given virtual machine.
- **Disable-VMResourceMetering.** Disables resource metering for a given virtual machine.
- **Reset-VMResourceMetering.** Resets resource-metering counters for a given virtual machine.
- **Measure-VM.** Displays resource-metering data for a given virtual machine.
- **Get-VM.** Displays the properties of a virtual machine, which allows you to view the **ResourceMeteringEnabled** property to determine whether resource metering is enabled for the virtual machine.

Implementing a Windows PowerShell Desired State Configuration Solution

The Windows PowerShell Desired State Configuration (DSC) is a new tool available as part of the Windows Management Foundation (WMF) 4.0. The WMF 4.0 consists of several support technologies for Windows server and client operating systems. WMF 4.0 enables some of the newer features of Windows Server 2012 R2 and the Windows 8.1 operating system on older systems starting from Windows 7 SP1 and Windows Server 2008 R2 SP1.

The DSC enables you script how you want to configure the software environment on target computers. DSC allows you to set server roles, manage registry settings, and manage user accounts and groups. DSC contains several Windows PowerShell language extensions that provide you with a way to maintain and manage computer configurations. In DSC, individual computers are called nodes. DSC enables you to:

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

DSC relies heavily on *remoting*, communication between different operating system processes, regardless of whether they are on the same computer. Therefore, it is important to have the Windows Remote Management (WinRM) command-line tool enabled on all nodes that you manage. You can enable WinRM for all nodes in the domain via a Group Policy Object (GPO). Also note that you will still need to enable WinRM when using DSC locally.

DSC includes several new Windows PowerShell cmdlets. Cmdlets are language extensions and declarative resources that allow you to configure your software environment and maintain and manage existing software environment configurations.

You can also use DSC to set up a central configuration server in which you can store the configurations of the various nodes in your environment. The central configuration server can also store custom DSC resources that you can use to control a large number of target nodes that you need to configure. It is especially important to configure target nodes as they come online. You can also use the central configuration server to check for configuration updates periodically on the target nodes.

DSC makes use of a new element keyword called Configuration. You use the Configuration keyword by defining it in a Windows PowerShell script block. Within a Configuration script block, you can have a node keyword to specify the node that the configuration applies to. There are many additional parameters you can apply in the Configuration script block.

- You can simplify configuration by gathering the existing configuration of running systems
- To prevent configuration drift, you can apply configuration resets against systems that are no longer configured as desired or as an originally set
- To enable continuous deployment, you can:
 - Run DSC scripts periodically as scheduled tasks
 - Check current configurations, if out of compliance, and apply the desired configuration

The following code sample checks to see if the Internet Information Services (IIS) is installed on the target computer named LON-WEB:

```
configuration IISWebsite
{
  node LON-WEB
  {
    windowsFeature IIS
    {
      Ensure = "Present"
      Name = "Web-Server"
    }
    WindowsFeature ASP
    {
      Ensure = "Present"
      Name = "Web-Asp-Net45"
    }
  }
}
IISWebsite
```

There are many new keywords and parameters associated with DCS. Because these parameters can control, install, manage, run, start, and stop many different functions within a node, it is not possible to list them all here.



Additional Reading:

For more information on Windows PowerShell, refer to the Windows PowerShell Desired State Configuration Overview at <http://go.microsoft.com/fwlink/?LinkID=392395>.

An online Hands on Lab about the DSC with Windows Server 2012 R2 is available from Microsoft TechEd North America 2013 at the following address: <http://go.microsoft.com/fwlink/?LinkID=392396>.

Implementing a System Center Global Service Monitor Solution

System Center 2012 offers a subscription-based cloud service called the System Center Global Service Monitor. The Global Service Monitor enables you to monitor the availability of external web-based applications from multiple locations. Rather than focusing on the infrastructure or an individual URL, the Global Service Monitor concentrates primarily on the external, web-based application. The Operations Manager console can be integrated with the Global Service Monitor functionality. When the subscription is activated, the Global Service Monitor becomes a node in the Administration node in the Operations Manager console. This integration allows you to monitor web applications both internally and externally. You can use your own management group and watcher agents and also obtain access to Microsoft agents in the cloud from 15 different locations.

- Availability monitoring monitors your externally facing websites regardless of where they are hosted
- Efficiency:
 - Is managed by Microsoft
 - Extends your Operations Manager infrastructure to the cloud – at no additional cost
- Integration can:
 - Deliver GSM data seamlessly into your Operations Manager environment
 - Facilitate the development operations cycle

There are two kinds of monitoring types, web application availability monitoring and Microsoft Visual Studio® web tests. With web application availability monitoring you perform a test on one URL from one location. For the Visual Studio web tests you can run tests from the 15 external locations provided by Microsoft as part of the subscription.

The following table lists the test parameters for the web application availability monitoring type.

Test Parameter	Description
Total tests	This provides the number of tests multiplied by the number of locations.
Trial subscription	A trial subscription limits the total tests to 25 per subscription and 10 tests for each location.
Paid subscription	A paid subscription allows up to five subscriptions per tenant, however the total number of tests is limited to 25 per subscription and 10 tests for each location.
Minimum interval per test	Greater than or equal to five minutes.
Global test timeout	30 seconds.

The following table lists the test parameters for the Visual Studio web tests:

Test Parameter	Description
Total tests	This provides the number of .webtest files multiplied by the number of locations.
Trial subscription	A trial subscription limits the total number of tests to 25 per subscription and a maximum of three tests for each POP location.
Paid subscription	A paid subscription allows up to five subscriptions per tenant, but the total tests are limited to 25 per subscription and a maximum of three tests for each POP location.
Minimum interval per test	Greater than or equal to five minutes.
Maximum number of requests per web test	100
Maximum web test file size	100 KB
Download/response size limit per request	500 KB

If you have tests that exceed the test maximums shown in the tables above the console will display these tests as "not monitored". The web application name is mapped to the performance counter instance name. If this web application is renamed, the performance counter instance name changes also. This can split the data returned in the performance of your tests across different counter names. It will no longer display any data previously appearing in a dashboard.

In order to use the Global Service Monitor, you should have a Windows account, a full or trial Windows Azure™ subscription, and preferably Visual Studio 2012 Ultimate, to replay scenarios against apps.



To learn more about the Global Service Monitor, please see <http://go.microsoft.com/fwlink/?LinkID=392397>.

Lesson 2

Overview of Operations Manager

Operations Manager is a server-monitoring tool that is suitable for use in enterprise environments where a high level of automation is required. Similar to Performance Monitor and Event Viewer, you can use Operations Manager to view server performance information and perform troubleshooting. However, Operations Manager centralizes this functionality to allow you to monitor hundreds of servers from a single central console.

Lesson Objectives

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

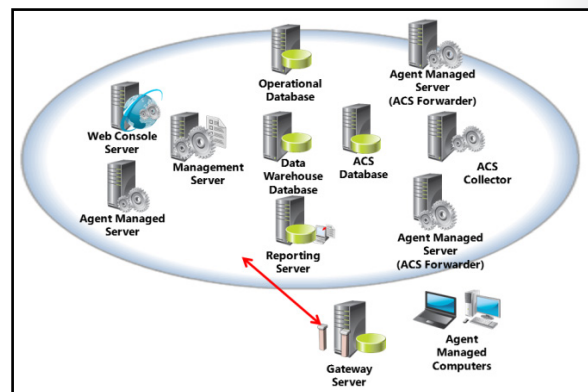
- Describe the components of Operations Manager.
- Use the Operations console.
- Describe resource pools.
- Describe the operations management agent.
- Install an Operations Management agent.
- Monitor a virtualized environment using Operations Manager.
- Describe the components of Audit Collection Services (ACS).
- Explain how to deploy ACS.
- Describe considerations for deploying Operations Manager.

Operations Manager Components

Operations Manager is a cross-platform monitoring and alerting solution that provides application and infrastructure monitoring. You can integrate Operations Manager with VMM, Service Manager, and Microsoft System Center 2012 - Orchestrator (Orchestrator) to provide automated remediation in response to errors, performance issues, and outages. Operations Manager also provides management packs to monitor other systems, including many non-Microsoft hardware and software components.

Operations Manager includes the following features:

- Network monitoring. Operations Manager supports the discovery of network routers and switches. This provides a platform for you to monitor network communications between computers.
- Application code monitoring. Operations Manager provides detailed monitoring information for applications, including Microsoft .NET Framework, and Java 2 Platform, Enterprise Edition applications. Operations Manager also provides the ability to identify problems with applications.



- End-to-end monitoring. Operations Manager can perform end-to-end monitoring of applications. This means that it can monitor the application, the operating system on which it is running, the hardware that the operating system relies on, and the network devices that provide access to the application. If an application is distributed across multiple systems, you can configure Operations Manager to show the application topology in a single pane. This allows administrators to see instantly where in the topology a problem occurs.
- Dashboard widgets. Operations Manager offers predefined and easily customizable dashboard widgets.

A number of components in Operations Manager work together to monitor and manage the IT infrastructure environment. You can install these components on separate systems across the network to provide scalability, availability, and better performance of the Operations Manager environment.

The collection of Operations Manager components that you use to monitor a given environment is called a *management group*. Large corporations may have several management groups that connect in a hierarchy. A local management group can receive consolidated data from connected management groups.

The following components make up an Operations Manager management group:

- Operations Manager server roles:
 - Management server
 - Gateway server
 - Operational database
 - Microsoft SQL Server® Reporting Services (SSRS) database
 - Reporting data warehouse database
 - Web console server
 - ACS collector
 - ACS database
 - ACS forwarder (on agent-managed devices)
- Agents
- Operations console
- Management packs
- Resource pools
- Application Advisor and Application Diagnostic consoles

Management Server

The Operations Manager management server is the first component installed in a management group. It provides communication between the agents and the management group, and serves as an administration point for configuration of the management group.

Agents send monitoring data back to a management server for inclusion into the Operations Manager database. If a management server fails, another management server can take over the load automatically, providing automatic failover for communication between agents and the management group.

Resource Pools

You can use resource pools to load balance between multiple management servers that are part of the same pool, providing automatic load balancing and failover to the Operations Manager environment. As soon as you add a server to a resource pool, it starts servicing requests automatically, and it reduces the load on the other servers. Only one management server can manage an individual object at any given time.

Gateway Server

Gateway servers enable the management of agent-managed systems that reside outside the Kerberos v5 protocol trust boundary of management groups. Agents in domains that are not trusted communicate with the gateway, and the gateway server then communicates with one or more management servers. Because the gateway server resides in an untrusted domain, you must use certificates to establish the identity of the managed systems, gateway server, and management servers. Furthermore, the gateway server performs discovery and installation, and relays ongoing administration traffic on behalf of the management server to the agents.



Note: All communication between gateway servers and management servers occurs through a single port (TCP 5723) that must be open on any firewall between these servers. Gateway servers also can discover and manage computers running UNIX and Linux operating systems over TCP port 1270 and, as needed, TCP port 22 (Secure Shell).

Operational Database

The operational database stores all monitoring, alerting, and performance data collected from agents. This includes event and alert data, performance data, and state change data. You must install the operational database on a Microsoft SQL Server 2008 R2 server. System Center 2012 Service Pack 1 (SP1) will support SQL Server 2012.

Reporting Data Warehouse Database

The reporting data warehouse database stores aggregated data that is used in reporting. Data collected from agents is stored in both the operational database and the reporting data warehouse database at the same time. Having a separate database for generating reports reduces the load on the operational database. The data warehouse database also stores data for an extended period.

Microsoft SQL Server Reporting Database

System Center 2012 uses SSRS to host and display reports. SSRS uses its own database to store report definitions and to cache report data.

Web Console Server

The Web console allows remote users to access most of the same functionality provided in the Operations console. You install the Web console on a computer that is running the Web Server role. You then secure the Web console using certificates.

Agents

The Operations Manager agent is a service that you can install on computers that Operations Manager will manage. The agent collects data based on rules and monitors, which are part of the management packs that you deploy to the agent.

Management Packs

Management Packs provide the basic elements that you use for monitoring objects managed by Operations Manager. Management Packs are created as XML files that can contain one or more of the following elements:

- **Object Discoveries.** Object discoveries discover objects that you want to monitor by examining data retrieved from an agent-managed computer. The most common form of discovery is to check for data in the registry or by using Windows Management Instrumentation (WMI). For instance, you can read the registry to find the version of the operating system that is running on a given computer.
- **Monitors.** You use monitors to determine the health state of a monitored object. There are different types of monitors, such as Unit monitors, Dependency Rollup monitors, and Aggregate Rollup monitors. Monitors can track a service, receive data from Simple Network Management Protocol (SNMP) devices, event logs, or performance counters, and return the health state of the object they are monitoring, based on criteria that you predefine.
- **Rules.** You use rules to generate alerts based on predefined criteria, collect data for reports, or run a scheduled task.
- **Tasks.** You use tasks to perform actions on monitored objects such as restarting a service, by running the **ipconfig** command-line tool on a system or recycling IIS application pool.
- **Views.** In the Operations console, views display specific information about monitored objects. You can create the following views:
 - **Alert view.** Displays alerts for a group of computers.
 - **Event view.** Displays specific events for a computer or group of computers.
 - **State view.** Displays the health state for an application, server, or group of servers.
 - **Performance view.** Displays the performance counters gathered for a given server.
 - **Diagram view.** Displays a diagram containing all servers used to perform a task or host an application.
 - **Task Status view.** Displays all tasks scheduled to run for a given server.
 - **Web page view.** Contains a view that you can access on a web browser.
 - **Dashboard view.** Contains multiple workspaces, each with its own view, which you can publish to Microsoft SharePoint® Server.
- **Knowledge Base.** Knowledge bases provide documentation specific to alerts generated by a given management pack. Knowledge bases contain detailed information on how to fix the problem related to the alert. You can customize knowledge bases to include specific information that relates to your own IT infrastructure environment.
- **Overrides.** Overrides allow you to customize the default monitoring provided by management packs. For instance, a given management pack may have a monitor stating that a server changes its state to unhealthy when its Available Memory falls below 10 percent. With an override, you can specify that a given server can have less available memory for a specific period without changing its status to unhealthy.

Application Advisor and Application Diagnostic Consoles

After you configure .NET Application Performance Monitoring, you can use the Application Advisor and Application Diagnostic to troubleshoot .NET Framework applications. Application Advisor identifies the applications that generated the highest number of alerts to help you start your troubleshooting. The Application Diagnostics console displays performance and reliability events for .NET Framework applications. These events help you identify how to resolve problems with the .NET Framework application. A Java application performance monitor is also available.

Audit Collection Services

You use ACS to collect events that an audit policy generates and then store them in a centralized database by using SQL Server. Once you have them in a centralized database, users can filter and analyze events by using SQL Server's data-analysis tools.



Reference Links: For more information about distributed deployments of Operations Manager, see <http://go.microsoft.com/fwlink/?LinkID=285319>.

Demonstration: Using the Operations Manager Console

The Operations console is a Microsoft Silverlight® application that administrators use to manage most of the typical functions performed in an Operations Manager environment. The Operations Console has five individual workspaces.

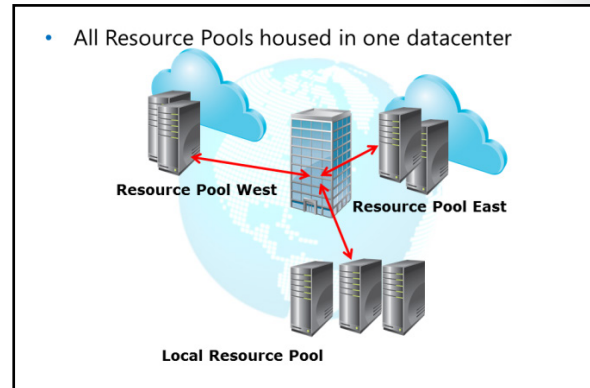
- **Monitoring.** The Monitoring workspace contains several views, including Active Alerts, Task Status, and Windows Server State. New views are added automatically as you install management packs. You also can create dashboard views, which contain multiple workspaces, and which can display data from a separate view. You then can publish these dashboards to Microsoft SharePoint Server.
- **Authoring.** You use the Authoring pane to customize your Operations Manager environment by creating and modifying rules, monitors, and overrides. You also can create your own management packs in the Authoring workspace.
- **Reporting.** You use the Reporting workspace to view and create reports.
- **Administration.** You use the Administration workspace to configure your Operations Manager environment and setup notifications, security, and resource pools.
- **My Workspace.** You use My Workspace to access your favorite views and saved searches. Each user has his or her separate My Workspace settings.

Demonstration Steps

1. Start the Operations Console.
2. In the Monitoring workspace, note the different views available under the Operations Manager node.
3. Identify the different nodes in the Authoring workspace.
4. Identify the different nodes in the Reporting workspace.
5. Identify the different nodes in the Administration console.
6. Identify the different nodes in the My Workspace console.

What Are Resource Pools?

The purpose of the resource pool is to use multiple management servers to distribute monitoring processes and take over for any failed member, especially in a large operation with distributed data centers that need a considerable monitoring workload. The monitoring workload is fast-paced and multiple management servers sharing this workload need to communicate quickly. Therefore, you should put all of your management servers in one central data center and then use gateway servers at the other sites. This will ensure the data shared between the management servers is replicated quickly. If for business reasons servers are in separate datacenters, ensure that fast and reliable wide area network links connect those datacenters.



For high availability of management servers, you need a minimum of two servers configured in a single resource pool. This will allow for failover and easy restoration if any one management server fails. By default all management servers are members of the All Management Servers resource pool, which provides automatic failover for monitoring. As new management servers are added, the All Management Servers resource pool balances the monitoring load of this management group. Two other default resource pools are created: the Notifications resource pool and the AD Assignment resource pool. All management servers belong to these pools as well. Because not all management servers will have access to specific communication devices, the Notifications resource pool allows you to specify which management servers will participate in the notification function. AD DS integration workflows are consolidated into the AD Assignment resource pool.

You can create your own resource pools to do such things as monitoring network devices in a specific geographical area or monitoring for other specific purposes. Management servers so designated can share monitoring workflows within the same resource pool. In the event that a particular management server fails, the other management servers in that resource pool will take over the monitoring function.

You can view and modify the membership of any resource pool in the Operations Manager console in the Administration workspace. Initially, all management servers are in all three default resource pools, and the membership of each pool is set to Automatic. If you remove any management servers from the default resource pools the membership type will change to manual. When the membership type is set to manual, any additional management servers must be added manually and do not automatically go into the default resource pools.

Options for Operations Manager Agent Installation

The Operations Manager agent is available for the Windows® XP operating system and newer clients, or on Windows Server 2003 or newer Windows Server operating systems. An Operations Manager agent is also available for UNIX and Linux operating systems.



Reference Links: For more information on the supported operating systems for the Operations Manager agent, see <http://go.microsoft.com/fwlink/?LinkID=285321> and search for "Operations Manager Agent – Windows-Based Computers" in the browser.

The following are options for Operations Manager agent installation:

- Discovery Wizard
- Manual install
- Scripted install
- Windows PowerShell Install-SCOMAgent

You can install the Operations Manager agent for Windows operating systems in three different ways. You can use the Discovery Wizard, perform a manual install, or create a script.

Discovery Wizard

You can use the Discovery Wizard to deploy Operations Manager agents. When you start the wizard, you can choose whether to discover computers running Windows operating systems, computers running UNIX or Linux operating systems, or network devices. After selecting the type of object to discover, you can choose between automatic discovery and advanced discovery.

- **Automatic discovery.** Automatic discovery scans the domain for all Windows computers that do not have an installed agent that is reporting to the management group from where the discovery is initiated. Once the scan is complete, you can select to which computers you want to deploy the agent.
- **Advanced discovery.** Advanced discovery allows you to change the scope of discovery, whether Servers or Clients or both, and specify a Lightweight Directory Access Protocol (LDAP) query to use for the scan. For example, you could scan for all computers whose name starts with LON-. You can also manually browse for computers, select a particular management server and verify that discovered computers can be contacted.

You can provide credentials for the discovery process or use the management server action account. You also use these credentials to install the agent, or you can use the local administrator account to perform the agent installation. In this case, you use the local administrator account to install the agent, and you use the management server's action account for discovery.

After discovery is complete, you can select from the list of discovered computers. The computers that you select are the computers that Operations Manager will monitor. You can choose to make the computers agent managed or agentless. Agent managed computers have the Operations Manager agent installed, while the management server performs agentless monitoring remotely. In general, agentless monitoring provides less functionality than agent managed monitoring, and it is not as scalable. Some management packs, such as the Active Directory® Domain Services management pack, do not support agentless monitoring.

In the Discovery wizard, you also can specify the agent installation directory and the agent action account that the agent will use when performing actions. You can change the agent action account after installation, but it should be set to the default of local system in most cases. Some management packs function properly only when you use the local system as the agent action account. The documentation for a management pack will include any requirements for the agent action account.

Manual Installation

In certain circumstances, it might be necessary to install the agent manually. This could be due to firewall restrictions or trust boundaries. For example, if you have a perimeter network with a web server that is available to the public, but which you have secured with firewalls. Furthermore, this web server is in its own domain or workgroup, which is kept separate from the domain and network in which the Operations Manager management group is located.

In this scenario, you use certificates and/or gateways to provide communication to the management server. Additionally, you must install the agent manually, because the firewall ports that you would typically use to install the agent by using the Discovery Wizard are blocked.

To install the agent manually, you can run `setup.exe` from the Operations Manager media, and select the **Local agent** option. Alternatively, you can run the `MOMAgent.msi` from the Agent folder on the Operations Manager media. With this second method, you can copy only the required agent files to the computer that will run the agent instead of making the full media available to the computer.

The agent folder has the agent installation files for AMD64, i386, and ia64 versions of the agent. You can run the `MOMAgent.msi` file from the appropriate folder to start the agent installation process.

Scripted Installation

You can pass parameters to `MOMAgent.msi` to script the installation and bypass the installation wizard. To install the agent using a scripted installation, use the following command:

```
%WinDir%\System32\msiexec.exe /i path\Directory\MOMAgent.msi /qn  
USE_SETTINGS_FROM_AD={0|1} USE_MANUALLY_SPECIFIED_SETTINGS={0|1}  
MANAGEMENT_GROUP=MGroupName  
MANAGEMENT_SERVER_DNS=MName MANAGEMENT_SERVER_AD_NAME =MName SECURE_PORT=PortNumber  
ACTIONS_USE_COMPUTER_ACCOUNT={0|1} ACTIONSUSER=UserName ACTIONSDOMAIN=DomainName  
ACTIONS_PASSWORD=Password
```

You can use the command above with a login script or even with a Microsoft System Center 2012 Configuration Manager program to push the installation to existing machines.

In Windows PowerShell, you can also use the **Install-SCOMAgent** cmdlet to deploy Operations Manager agents by using a client push installation.

Demonstration: Installing the Operations Manager Agent

In this demonstration, you will see one method you can use to install the Operations Manager agent.

Demonstration Steps

1. On LON-OM1, run the Discovery Wizard to push the Operations Manager client to network servers using the following options
 - Select **Windows computers**
 - Use **Advanced discovery**
 - **Browse** for, or type-in computer names: **LON-SVR1, LON-SVR2**
 - **Select All** and the select **Finish**
2. On LON-SVR1, open **Administrative Tools**, open **Services**, and note that **Microsoft Monitoring Agent** is running.

Monitoring a Virtualized Environment Using Operations Manager

You can use Operations Manager to monitor both virtual machines and virtual machine host computers. We cover the monitoring differences that virtualization introduces in depth. There are two major management packs that cover virtualization. They are the System Center 2012 Management Pack for Windows Server Hyper-V 2012 and the Monitoring Pack for System Center 2012 Virtual Machine Manager.

The System Center 2012 Management Pack for Windows Server Hyper-V 2012 monitors the Hyper-V role running on a server running Windows Server 2012 and newer operating systems. With this monitoring pack, you can provide essential monitoring of the host computer, the virtual machines, and the various virtual components running on the host. You can also monitor the use of Microsoft RemoteFX® technologies and virtual machine replication.

The Monitoring Pack for System Center 2012 VMM monitors the entire VMM infrastructure. Monitored components of a System Center 2012 VMM infrastructure include:

- The VMM management servers.
- Assigned physical hosts in VMM.
- Clustered VMM management servers and hosts.
- Virtual machines managed by VMM.
- The virtual environment to include fabric monitoring for System Center 2012 R2 VMM and Operations Manager.

VMM- managed physical hosts include Hyper-V, VMware ESX and Citrix XenServer. It is important to note that the Operations Manager management pack for System Center 2012 is specific to VMM 2012 and cannot be used on earlier versions. Agentless monitoring of hosts in virtual machines is not supported by the Monitoring Pack for VMM. You can install Operations Manager agents on the Hyper-V host and the virtual machines. Also note that you cannot install the Monitoring Pack for VMM if VMM is not already installed. If your organization does not include VMM you can still use the Hyper-V Monitoring Pack.

You can integrate Operations Manager and VMM by installing the Operations Manager console on the VMM management server and then adding the Monitoring Pack for System Center 2012 VMM to the Operations Manager. Once this is done you can, on the VMM console, in the Settings workspace to the left-hand side, click on System Center Settings, and then right-click on Operations Manager Server and select Properties. This then brings up the Add Operations Manager wizard, which will walk you through the configuration steps to allow Operations Manager to monitor your VMM infrastructure.

- Operations Manager monitors both virtual machines and virtual machine hosts
- The System Center 2012 Management Pack for Windows Server Hyper-V 2012 can monitor Windows Server 2012 and newer operating systems running the Hyper-V role
- The Monitoring Pack for System Center 2012 VMM monitors the entire VMM infrastructure
- Add Operations Manager Wizard has specific configuration steps for monitoring a VMM infrastructure

What Is Audit Collection Services?

ACS collects the records that an audit policy generates, and then stores them in a centralized database. By default, when you implement an audit policy on a computer that is running a Windows operating system, that computer automatically saves all events that are generated by the audit policy to its local Security log. This is true for Windows client workstations and Windows servers. In organizations that have strict security requirements, audit policies can quickly generate large volumes of events.

When you use ACS, you can consolidate individual Security logs into a database that you can manage centrally, and then filter and analyze events by using the data analysis and reporting tools that Microsoft SQL Server provides. With ACS, only a user who has been given the specific right to access the ACS database can run queries and create reports on the collected data.

ACS requires the following components:

- ACS forwarders.
- ACS collector.
- ACS database.

ACS Forwarder

The ACS Forwarder service is part of the Operations Manager agent, and is installed by default on all computers where the agent is deployed. However, this service is not enabled until you run the Enable Audit Collection task from the Operations Console. After enabling this service, all security events are sent to the ACS collector and the local Security log.

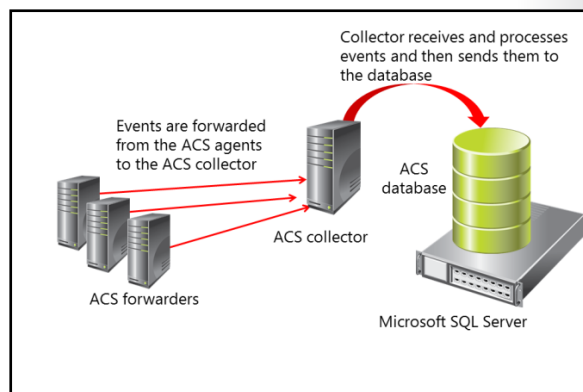
ACS Collector

The ACS collector receives and processes events from ACS forwarders and then sends this data to the ACS database. This processing includes disassembling the data so that it can be placed in several tables within the ACS database, minimizing data redundancy, and applying filters so that unnecessary events are not added to the ACS database.

The number of ACS forwarders that can be supported by a single ACS collector and ACS database can vary, depending on the following factors:

- The number of events that your audit policy generates.
- The role of the computers that the ACS forwarders monitor, such as domain controller versus member server.
- The level of activities on the computer.
- The hardware on which the ACS collector and ACS database run.

If your environment contains too many ACS forwarders for a single ACS collector, you can install more than one ACS collector. However, each ACS collector must have its own ACS database.



The minimum requirements for an ACS collector are as follows:

- An Operations Manager management server.
- A member of an Active Directory domain.
- A minimum of 1 GB of RAM, with 2 GB recommended.
- At least a 1.8 gigahertz (GHz) processor, with a 2.8 GHz processor recommended.
- At least 10 GB of hard disk space available, with 50 GB recommended.

On each computer on which you plan to install the ACS collector, you must download and install the latest version of the Microsoft Data Access Components (MDAC) from the Microsoft Download Center website.

ACS Database

The ACS database is the central repository for events that are generated by an audit policy within an ACS deployment. The ACS database can be located on the same computer as the ACS collector, each should be installed on a dedicated server for better performance.

The requirements for an ACS database are as follows:

- System Center 2012 Operations Manager with SQL Server 2005 or SQL Server 2008. For the best performance, use the Enterprise version of SQL Server. You can choose an existing or new installation of SQL Server.
- System Center 2012 SP1 and R2 Operations Manager with SQL Server SQL 2008 R2 SP1, SQL Server 2008 R2 SP2, SQL Server 2012, or SQL Server 2012 SP1. For the best performance, use the Microsoft SQL Server, Enterprise Edition. You can choose an existing or new installation of Microsoft SQL Server.
- A minimum of 1 GB of RAM, with 2 GB recommended.
- At least a 1.8 GHz processor, with 2.8 GHz recommended.
- At least 20 GB of hard drive space, with 100 GB recommended.



Reference Links:

- For information on how to filter ACS events for UNIX and Linux computers, see <http://go.microsoft.com/fwlink/?LinkID=290803>
- For more information on the use of Dynamic Access Control with ACS in Operations Manager, see <http://go.microsoft.com/fwlink/?LinkID=285318>

If you use Microsoft SQL Server Standard Edition, the database must pause during daily maintenance operations. This may cause the ACS collector queue to fill with requests from ACS forwarders. A full ACS collector queue then causes ACS forwarders to be disconnected from the Planning and Configuration Notifications and Reporting CS collector. Disconnected ACS forwarders reconnect after the database maintenance completes, and then the database processes the queue backlog. To ensure no audit events are lost, allocate a sufficient amount of hard-disk space for the local security log on all ACS forwarders. However, this pause provides a brief period that can be considered a security risk because events are lagging behind actual activity. If you use Microsoft SQL Server Enterprise edition, you can avoid this pause because the database does not pause during maintenance.

Deploying Audit Collection Services

To use ACS, you must deploy at least one ACS collector, one ACS database per collector, ACS forwarders, and ACS reporting.


Installing ACS Collectors and Databases

Follow these steps to install an ACS collector and an ACS database:

1. Log on to the server by using an account that has local administrative credentials.
2. On the Operations Manager installation media, run **Setup.exe**, and then click **Audit collection services**. For monitoring UNIX and Linux computers, click **Audit collection services for UNIX/Linux**.
3. On the **Welcome** page of the Audit Collection Services Collector Setup Wizard, click **Next**.
4. On the **License Agreement** page, read the licensing terms, click **I accept the agreement**, and then click **Next**.
5. On the **Database Installation Options** page, click **Create a new database**, and then click **Next**.
6. On the **Data Source** page, in the **Data source name** box, type a name that you want to use as the Open Database Connectivity (ODBC) data source name for your ACS database (by default, this name is **OpsMgrAC**), and then click **Next**.
7. On the **Database** page, if the database is on a separate server than the ACS collector, click **Remote Database Server**, and then type the computer name of the database server that will host the database for this installation of ACS. Otherwise, click **Database server running locally**.
8. In the **Database server instance name** field, type the name of the database that will be created for ACS, and then click **Next**. If you leave this field blank, the default name **OpsMgrAC** will be used. In the **Database** name field, the default database name of **OperationsManagerAC** is entered automatically. You can choose to highlight the name and then type in a different name, or leave the default name.
9. On the **Database Authentication** page, select one of the authentication methods. If the ACS collector and the ACS database are members of the same domain, select **Windows authentication**. Otherwise, select **SQL authentication**, and then click **Next**.

To deploy ACS:

- Install the ACS collector
- Specify the database server
- Set the retention schedule
- Enable ACS forwarders

 **Note:** If you select **SQL authentication** and then click **Next**, the **Database Credentials** page displays. In the **SQL login name** box, enter the name of the user account that has access to the computer running SQL Server and the password for that account in the **SQL password** box, and then click **Next**.

10. On the **Database Creation Options** page, click one of the two following options, and then click **Next**:
 - **Use SQL Server's default data and log file directories** to use SQL Server's default folders.
 - **Specify directories** and enter the full path, including drive letter, of the location you want to use for the ACS database and log file. For example, C:\Program Files\Microsoft SQL Server\MSSQL.1\MSSQL\Data.
11. On the **Event Retention Schedule** page, click **Local hour of day to perform daily database maintenance**. Choose a time of day when the number of expected security events is low because during the database maintenance period, database performance will be impacted. In the **Number of days to retain events** box, type the number of days ACS should keep events in the ACS database before the events are removed during database grooming. The default value is 14 days. Click **Next**.
12. On the **ACS Stored Timestamp Format** page, choose **Local** or **Universal Coordinated Time**, formerly known to as Greenwich Mean Time, and then click **Next**.
13. The **Summary** page displays a list of actions that the installation program will perform to install ACS. Review the list, and then click **Next** to begin the installation.



Note: If a **SQL Server login** dialog box displays and the database authentication is set to **Windows authentication**, click the correct database and verify that the **Use Trusted Connection** check box is selected. Otherwise, clear the check box, type the SQL login name and password, and then click **OK**.

14. When the installation completes, click **Finish**.

Enabling ACS Forwarders

1. Log on to the computer with an account that is a member of the Operations Manager Administrators role. This account must also have the rights of a local administrator on each agent computer that you want to enable as an ACS forwarder.
2. In the Operations console, click **Monitoring**.
3. In the Monitoring workspace, expand **Operations Manager**, expand **Agent Details**, and then click **Agent Health State**. This view has two workspaces, and the actions in this procedure are performed in the right-hand workspace under Navigation..
4. In the details workspace, click to select all agents that you want to enable as ACS forwarders. You can make multiple selections by pressing Ctrl or Shift as you click.
5. In the Actions workspace, under **Health Service Tasks**, click **Enable Audit Collection** to open the **Run Task - Enable Audit Collection** dialog box.
6. In the **Task Parameters** section, click **Override** to open the **Override Task Parameters** dialog box.
7. In the **Override the task parameters with the new values** section, click the **CollectorServer** parameter. In the **New Value** column, type the fully qualified domain name (FQDN) of the ACS collector; and then click **Override**.



Note: If you are enabling ACS on a gateway or management server and you do not specify the CollectorServer parameter, the task will fail with a "Type Mismatch Error." To avoid this, provide a value for the override.

8. In the **Task credentials** section, click **Other**. In the **User Name** box, type the name of a user account that belongs to the local Administrators group on the agent computers. In the **Password** box, type the password for this user account. Click to expand the **Domain** list to view the available domains, and then click the domain of the appropriate user account.
9. Click **Run Task**. The **Task Status** dialog box displays, tracking the progress of the task.
10. When the task completes successfully, click **Close**.

**Reference Links:**

- For more information on how to configure certificates for ACS Collector and Forwarder, see <http://go.microsoft.com/fwlink/?LinkID=285320>
- For more information on the ACS administration and the AdtAdmin.exe tool, see <http://go.microsoft.com/fwlink/?LinkID=285317>

Monitoring Distributed Applications

Operations Manager can display the health state of your business-critical applications graphically by using a Distributed Application Diagram. This diagram enables the application owners to view the overall health state of all components that make up the distributed application quickly.

Many of the Microsoft Management Packs include an automatically generated Distributed Application

Diagram for the application that is monitored. This includes AD DS and Exchange. From inside the diagram, operators can use the Health Explorer to view the state of all relevant monitors that are currently running against the Distributed Application.

To create a Distributed Application Diagram, you must understand the components of a distributed application, and how you can add them to the Distributed Application Designer.

When you monitor distributed applications, you should make sure that all components of the application are discovered and monitored. This helps make sure that if an issue occurs with any component, the distributed application diagram will update correctly.

This is also the case for the .NET Framework applications. When you configure Application Performance Monitoring in Operations Manager, several objects are created in the Operations Manager database.

These objects help you understand the end-to-end health and performance of the application. Similarly, when you create a Web Application Availability monitor, objects are created that help you understand the performance and availability of the application from the end-user.

You can use the .NET 3-Tier Application template in Operations Manager to create a diagram view that monitors the performance and availability of .NET Framework applications. To create a diagram view that you can use to monitor a .NET application, you must first understand how to configure the .NET 3-Tier Application template.

- When you use Management Packs to automatically generate a distributed application diagram, you can include AD DS and Microsoft Exchange Server
- Ensure all application components are discovered and monitored
- Use the .NET 3-Tier Application template in Operations Manager to create and monitor diagram views that monitors the performance and availability of .NET Framework applications

Considerations for Deploying Operations Manager

You must consider several factors when you design your Operations Manager environment. Factors you need to consider include the architecture of AD DS, firewalls, and bandwidth availability between sites.

The following list summarizes these factors and how they relate to your design decisions:

- AD DS architecture. You can use a management group to manage any computer within a trusted forest. If you need to manage computers outside of your Kerberos authentication boundary, you need to use a gateway server and certificates.
- Physical infrastructure. The minimum bandwidth required for communications between an agent-managed computer and a management server is 64 kilobits per second (Kbps). Between an agentless monitored computer and the management server, the minimum bandwidth required is 1,024 Kbps. If your bandwidth falls below these settings, you must consider using multiple management servers.
- Administration. If you have different sets of monitoring needs in different physical locations, you should consider deploying multiple management groups. As a best practice, the minimum bandwidth between tiered management groups should be 1,024 Kbps.
- Number of Operations consoles. Operations consoles can be used by multiple users within your organization. The recommended maximum number of Operations consoles that should be open simultaneously is 50 consoles per management group.
- Number of agent-managed computers per management server. A management server should manage no more than 3,000 computers.
- Number of agent-managed computers per gateway server. As a best practice, a gateway server should manage no more than 2,000 computers.
- Number of agentless-managed computers per management server. As a best practice, a management server should manage no more than 10 agentless-managed computers.
- Number of Agentless Exception Monitoring (AEM) computers per dedicated management server. As a best practice, a dedicated management server should manage no more than 25,000 AEM computers.
- Collective client monitored computers per management server. As a best practice, a management server should manage no more than 2,500 client computers.
- Number of agentless-managed computers per management group. As a best practice, a management group should manage no more than 60 agentless-managed computers.
- Number of agent-managed computers per management group. As a best practice, a management group should manage no more than 6,000 computers when using up to 50 open Operations consoles, or 15,000 computers when using 25 Operations consoles.
- Number of AEM computers per management group. As a best practice, a management group should manage no more than 100,000 AEM computers.
- Management servers per agent for multihoming. As a best practice, a managed computer is recommended to use multihoming on no more than four different management servers.

Operations Manager deployment considerations:

- AD DS architecture
 - Use a gateway server and certificates outside of AD DS trusts
- Physical infrastructure
 - Bandwidth requirements
- Administration
 - Consider deploying multiple management groups
- Number of consoles
 - Recommended no more than 50 per management group
- Number of managed computers
 - Depends on whether managed computers are agent, agentless, or Agentless Exception Monitoring



Reference Links: System Requirements: System Center 2012 R2- Operations Manager
<http://technet.microsoft.com/en-us/library/dn249696.aspx>

Lesson 3

Planning and Configuring Monitoring Components

Management packs store the rules for monitoring the health of your servers and generating alerts when the health status of a server requires attention. You need to understand how to import management packs for specific server roles and applications. You also need to understand how to tune those management packs for your environment to ensure that only relevant alerts are generated.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe management packs.
- Describe management pack installation and tuning.
- Describe management pack creation.
- Review implementation considerations for management packs.
- Import management packs.
- Describe and configure notifications.
- Describe reporting using SSRS.
- Explain configuring reports and reporting.
- Describe the notifications and reporting options in Operations Manager.
- Discuss planning monitoring with Operations Manager.

Overview of Management Packs

As discussed earlier in this module, management packs provide the basic elements used for monitoring objects managed by Operations Manager. Management packs can contain one or more of the following elements:

- Object discoveries. Use object discoveries to discover objects that you want to monitor.
- Monitors. Use monitors to determine the health state of a monitored object.
- Rules. Use rules to generate alerts based on predefined criteria, collect data for reports, or run a scheduled task.
- Tasks. Use tasks to perform actions on monitored objects.
- Views. Use views in the Operations Console to view specific information about monitored objects.
- Knowledge. Knowledge bases provide documentation specific to alerts generated by a given management pack, and contain detailed information on how to fix the problem related to the alert.
- Overrides. Overrides allow you to customize the default monitoring provided by management packs.

Management pack components:

- Object Discoveries
- Monitors
- Rules
- Tasks
- Views
- Knowledge
- Overrides

Installing and Tuning Management Packs

Once you install Operations Manager, you will have access to dozens of management packs directly from Microsoft. You do not need to install every management pack, however. Keeping the deployed management packs to a minimum will ensure the best results. Each management pack contains a large number of objects. Some of these objects collect data from agents in small intervals of time—data that is sent over your network, and stored in the Operations Manager database.

You must consider what management packs are applicable to your environment based on the applications and services that you need to monitor. A more detailed approach to tuning management packs is listed below:

1. Identify the applications and devices in your infrastructure that you want Operations Manager to monitor.
2. Read the management pack guide for each management pack and identify which monitors, rules, and discoveries need to be overridden for your environment.
3. Download and import the management packs into a test or lab environment.
4. Create a separate management pack to store the customizations for each management pack that you import.
5. Create your groups, overrides, and any other objects that you need in your management packs.
6. Test the management pack in your test or lab environment.
7. Export your custom management packs.
8. Install the original management packs and custom management packs in your production environment.

In order to evaluate, adjust, and tune the management pack as necessary, you should first consider importing new management packs into a test environment. Effective management pack tuning involves the entire team, including the service owner or experts and the operations team members who monitor the alerts and events. Depending on the monitored service that a particular management pack is intended for, you might also include networking and security experts. The default settings of a particular management pack might not match your environment. You can start with defaults and create overrides to prevent alerts that are not relevant for your environment.

To customize management packs, you must:

1. Identify what to monitor
2. Import management packs to test the environment
3. Study management pack guides
4. Customize management packs
5. Test the management packs
6. Export customizations
7. Import customizations into production
8. Install the original management packs and custom management packs

Tools for Creating a Management Pack

You can create your own management packs on the Operations Manager console in the Authoring workspace. The Authoring workspace contains several templates and wizards that allow you to create various monitoring scenarios with minimal knowledge of authoring concepts. This allows you to create a management pack file, create predefined monitoring scenarios and custom classes and discoveries if the template provides for a specific case. There are several components within the Authoring workspace that allow you to create custom management packs. Management pack templates allow you to create a management pack quickly. You can use a single template to make different monitors, rules, and new target classes without having access to specific details. This is usually the simplest and easiest way to make a management pack. You can also use the Distributed Application Designer, which enables you make a single application that is comprised of several objects. However, you cannot create new monitoring for these objects if the health of these objects is based on monitors for these objects that are already running. Finally, the Authoring workspace of the Operations Management console includes several Authoring wizards. Wizards are available for different data sources and monitoring scenarios, but you will need to supply the details of the data the workflow will use and what you will want done with that data.

The Authoring workspace:

- Contains several templates and wizards
- Allows you to create various monitoring scenarios with minimal knowledge of authoring concepts

When you edit directly using any XML editor:

- Requires extensive XML knowledge
- Can configure elements not available in other methods

Management packs are XML files, so you can also use an XML editor to edit a management pack directly. Note that this type of editing is complex and beyond the scope of this lesson.

In many cases, you can create a management pack to store overrides related to an imported management pack. To simplify management, a best practice in this case would be to create a management pack for overrides for each management pack that you import.

Considerations for Implementing Management Packs

Careful consideration and planning should be part of your management pack strategy. Testing is highly important for a successful deployment of the management packs. You should thoroughly test not only the implementation, but also the tuning of new and already in production management packs. Very often, after deploying a management pack, you might want to make changes based on the results the management packs help collect, and this will lead to further tuning. You should also follow these guidelines:

When you implement management packs, you should:

- Deploy or tune one management pack at a time
- Keep it simple
- Do not deploy management packs unless you need the specifics they collect
- Editing a management pack directly with an XML editor:
*Requires extensive XML knowledge

- Deploy or tune one management pack at a time. Because there are so many management packs, and your comprehensive management solution might include many of the packs, it is tempting to deploy as them as quickly as possible. However, it is important to proceed carefully and make sure the proper settings are made and the results returned reflect the data you wish to collect. By only deploying or tuning one management pack at a time, you can give the process your full attention.

- Keep it simple and build upon success. Management pack tuning can lead to fairly complex designs. If you incrementally enhance your designs, you will have better control over the output you want, without having to go back to the design phase if your results are not what you want.
- Do not deploy management packs unless you need the specifics they collect. There are many free management packs available from Microsoft. However, you do not want to deploy a management pack unless the data collected has meaningful information for your organization. Management packs consume resources like any other processes, such as CPU cycles, memory, and bandwidth. Keeping consumption down will enhance your overall operations.
- Make use of the available tools. For example, you can use the Visio Management Pack Designer. While the Visio Management Pack Designer is a free download from Microsoft, Visio is not. Still, the Visio Management Pack Designer is an excellent tool that allows you to graphically map out aspects of a management pack design. Alternatively, you can use an XML editor to do very granular and concise tuning of existing management packs.

Demonstration: Importing Management Packs

Demonstration Steps

1. On **LON-OM1**, import the SQL Server management pack to LON-OM1, using the following options:
 - a. Administrative node, Management Packs, select **Import Management Packs**
 - b. Add from disk (no for Online Catalog Connection): and expand through to the following location: **C:\Program Files (x86)\System Center Management Packs\System Center Monitoring Pack for SQL Server**, and select all SQL Server Management Packs and then click **Install**.



Note: After about 15 of the SQL Server Management Packs are installed, navigate to the Monitoring workspace and look for Computers in the Microsoft SQL Server node. **LON-OM1** should appear in that node..

2. Revert virtual machines
3. When you finish the demonstration, revert all virtual machines to their previous state.

Overview and Configuring of Notifications

Operations Manager users can use the Operations console to view alerts generated by the several rules and monitors deployed to the Operations Manager environment. Although the Operations console is a perfect place to access these alerts, sometimes you might need to have alerts sent to a different group from within your organization. This could include groups that might not have access to the console, or might not have an Operations Manager administrator who visits the console frequently. For instance, you might want to notify the database administrators of a critical e-commerce database if the database they manage changes its status to offline. There might be one or more alerts from different monitors or rules that monitor the database and the server on which the e-

- You can send messages to users based on predefined criteria with Notifications
- Users do not need to use the Operations console to view alerts
- Configure the following:
 - Channel
 - E-Mail (SMTP)
 - Instant Message (IM)
 - Text Message (SMS)
 - Command
 - Criteria
 - Subscriber
 - User or group
 - Schedule
 - Addresses
 - Subscription

commerce database resides. Furthermore, the database administrators might not check the Operations console on a regular basis. In this scenario, you can use a notification to send an alert to the database administrators. Notifications can be sent based on sets of criteria by using email, instant messaging, text messaging, or by running a custom command.

Configuring Notifications

To use notifications, you must create and configure the following objects in the Operations console:

- Channel. Channels specify the type of messaging solution used for the notification. There are four channel types that you can use:
 - E-Mail. Send email messages by Simple Mail Transfer Protocol (SMTP).
 - Instant message. Use to send instant messages (IMs).
 - Text Message. Use to send text messages by SMS (Short Message Service).
 - Command. Use to run a custom command.
- Subscriber. Subscribers represent users or groups that can receive a notification. Each subscriber contains the following settings:
 - Subscriber name. We recommend using the account name for the user or group that the subscriber represents.
 - Schedule. Specifies what days and times the subscriber can receive notifications.
 - Addresses. Specifies a list of addresses that can use any of the preconfigured channels to deliver a notification to the subscriber.
- Subscriptions. A subscription specifies who receives an alert based on a set of criteria. Subscriptions are composed of:
 - Criteria. A set of predefined conditions, similar to Microsoft Outlook® rules, that can be combined to determine when a notification is sent.
 - Subscriber. A list of subscribers that should receive notifications for this subscription.
 - Channels. The list of channels used to send the notification.

Overview and Configuring of Reports and Reporting

Operations Manager uses SSRS to create, manage, and store reports. You add reports to the Operations console either by importing a management pack, or by designing a new report in the Reporting workspace.

When you design your first report from the Reporting workspace, you will be prompted to install the Microsoft SQL Server Reporting Services Report Builder tool from SQL Server. Report Builder is a web-based application that allows you to build tabular, matrix, and chart reports based on models created by a database administrator.

Reporting components:

- SSRS
- Reporting Workspace

When viewing reports, you can:

- Run reports from the Reporting workspace
- Schedule reports to be delivered automatically



Note: Creation of report models and reports using Report Builder are beyond the scope of this course. Learn about these reports in detail in course *10778 Implementing Data Models and Reports with SQL Server 2012*.

Configuring Reports and Reporting

During the installation of Operations Manager, you are required to configure the SSRS server. You should install the reporting component on the SSRS server. At any time after installation, you can change which server SSRS uses by accessing the Administration workspace, and, from the Settings node, selecting Reporting. Prior to doing this, make sure that you have installed SSRS on the new reporting server.

Each report that you run will most likely contain numerous parameters. Most reports will allow you to aggregate data by day, week, or month. You should also be able to specify a beginning and ending date for the data that is displayed.

To view a report, complete the following steps:

1. From the Operations Console, open the **Reporting** workspace.
2. Navigate to the report that you want to view.
3. Double-click on the report.
4. Select the values that you want to use for each parameter.
5. Click **Run**.

You can also schedule reports to run and deliver the report to a file share at a specific time. To schedule a report, execute the following steps:

1. From the Operations Console, open the **Reporting** workspace.
2. Navigate to the report that you want to schedule to run.
3. From the Tasks pane, click **Schedule**.
4. On the **Delivery Settings** page, specify a description for the schedule, and then select **Windows File Share** as the delivery method.
5. In the **File name** text box, type a name for the file to be saved in the share.
6. In the **Path** text box, type the path of the share to be used.
7. In the **Render Format** drop-down list box, click the format for the report.
8. In the **Write mode** drop-down list box, click one of the following:
 - **None** to stop the report from running if the file already exists.
 - **Autoincrement** to create a new file each time the scheduled report runs.
 - **Overwrite** to overwrite the previous file if one exists.
9. In the **User name** and **Password** text boxes, type the credentials used to run the report, and then click **Next**.
10. In the **Subscription Schedule** page, specify a frequency for the report, and then click **Next**.
11. In the **Report Parameters** page, specify the values for each parameter, and then click **Finish**.

Considerations for Implementing Notifications and Reports

Before implementing notifications and reports, consider the following:

- **Subscribers.** Use groups instead of individuals for subscribers. Make sure you have a manageable number of subscribers and document what alerts they must receive. In addition, if a user or group of users is constantly looking at the Operations console for alerts, they may not need to be set up as subscribers.
- **Subscriptions.** Make sure you monitor subscriptions and ask subscribers how frequently they receive notifications. It is common for an untuned environment to fill subscribers' mailboxes with more notifications than they can manage. In situations like this, subscribers will most likely ignore the messages. Therefore, make sure only relevant alerts are sent as notifications.
- **Reporting.** We recommend that you use a separate instance of SQL Server to host the reporting data warehouse database, the operational database, and the SSRS database. Make sure you monitor these components.
- **Database grooming.** Review your database grooming settings in the Settings node on the Administration workspace. You use these settings to reduce the amount of space used by the database while maintaining the amount of historical data needed by your organization. Consider any regulations that your company might have to comply with, including monitored data storage.

When you implement notifications, consider:

- **Subscribers**
 - Use groups
 - What alerts they should receive?
- **Subscriptions**
 - Frequency
 - Keep received email down to manageable size

When you implement reporting, consider:

- Database server instance
- Database grooming

Discussion: Planning Monitoring with Operations Manager

Review the questions below with the instructor and the other students. Be prepared to discuss your answers.

Question: The CIO wants you to design a monitoring solution for all aspects of the company's IT infrastructure, client devices, and various networking devices, even those not running a Microsoft operating system. How can you use Operations Manager in these situations?

Question: The CIO wants to know if there is a way that he can use Operations Manager data to gauge the state of the overall IT infrastructure quickly. Is this possible, and how can you present high-level data views?

Question: Contoso, Ltd. has a number of wireless LANs within their headquarters, all within the corporate firewall. Bandwidth is limited on the wireless LANs and many of the devices you wish to monitor over the wireless LANs have limited storage space and memory. Given these limitations, what are your options for monitoring the devices on this network?

- How can you use Operations Manager to monitor all aspects of a company's information technology infrastructure?
- How can you use Operations Manager to quickly view a high-level summary of the infrastructure's health?
- What solution exists to monitor devices with limited resources, even over slow WAN links?

Lesson 4

Configuring Integration with VMM

You face a challenge when you use Operations Manager to monitor virtual environments because hardware resources, such as memory and processor, are virtualized for virtual machines. Lesson 1 details these issues, and how you can address them. If VMM manages your virtual environment, you can integrate that environment with Operations Manager, and use the PRO-enabled management packs provided by Microsoft and other vendors to monitor your virtual machine environment correctly. You can use a technology called Performance and Resource Optimization (PRO) tips in VMM to manage Operations Manager PRO-enabled management packs.

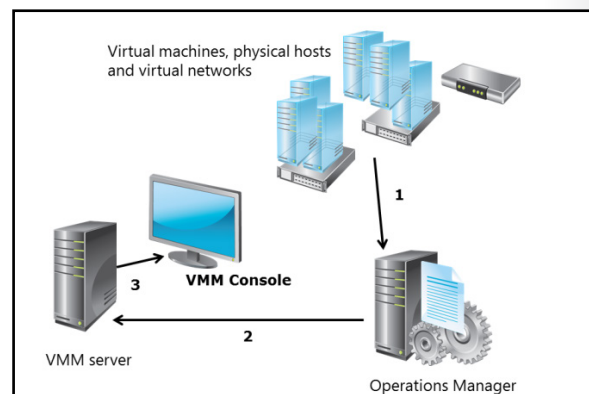
Lesson Objectives

After completing this lesson, you will be able to:

- Integrate Operations Manager with VMM.
- Configure Operations Manager and VMM integration.
- Describe PRO.
- Compare PRO and advanced placement options in VMM.
- Configure advanced monitoring for virtualization components.

Operations Manager Integration with VMM

You can provide many of the monitoring and reporting features that your VMM deployment requires by integrating VMM with Operations Manager. When you configure the integration, you can use Operations Manager to monitor both physical and virtual servers, and then make the data that Operations Manager collects available to VMM administrators as reports in the VMM management Console. Operations Manager can provide automated actions or suggestions on a specific action to take in the virtualized environment. For example, a PRO-enabled management pack can detect when there is a problem with a virtualization host and, if appropriate, can automatically initiate an action in VMM to migrate virtual machines to another virtualization host.



In the Monitoring pane of the Operations Console, the Virtual Machine Manager monitors views that you can find under the Virtual Machine Manager folder. There are views for Active Alerts and Health State. Use these views to gain a quick view into the health and availability of the VMM environment.

The Managed Resources folder contains several health state views that you can use to view the health of individual VMM components, such as the Cloud health and Host health views. Use these views to display the configuration of VMM components. For example, when you select a cloud in the Cloud health view, the cloud capability details are displayed in the Detail View. This view includes the Maximum VM Count, Maximum Storage, and Maximum Memory settings for the cloud.

The Performance folder contains performance views to monitor a component's current performance and performance over a selected period. For example, the Host Performance view displays performance counters that relate to bytes that are sent and received for the virtual network adapter. You can select these counters to display a graph that represents the bytes sent and received over the virtual network adapter. By using the Select Time Range option, you can set the time range of displayed data. This permits you to see over time how much data was sent and received. This can also be useful in capacity planning.

Other performance views include the following:

- Cloud Performance
- Service Performance
- Host Cluster Performance

Use the Diagram View that is located under the Virtual Machine Manager Views folder to view the overall health of the VMM environment. When you use a distributed application diagram, you can expand the components of VMM and instantly view the health of each. For example, by expanding the VMM Infrastructure component group, you can view the health of the VMM Server. By expanding the VMM Server component group, you can then view the health of the virtual machine guests, the VMM database, and any clouds that VMM manages.

You can also open the Health Explorer from any selected component in this view. Use this to troubleshoot failed monitors. The Summary, Causes, and Resolution section is displayed in the Health Explorer. It displays State Change Events that you can use to match problems with other activities that have occurred in the VMM environment.

When you select a virtual machine in this view, several Virtual Machine Tasks become available in the Tasks pane. Use these tasks to perform tasks against the selected virtual machine. For example, the Create Checkpoint task creates a checkpoint on the selected virtual machine. You can also perform other tasks, such as Start, Stop, and Pause.

Configuring Operations Manager and VMM Integration

Before configuring the integration between VMM and Operations Manager, ensure that your system meets the following requirements:

- Correct Windows PowerShell version. You must have Windows PowerShell 2.0 for System Center 2012, Windows PowerShell 3.0 for System Center 2012 SP1, or Windows PowerShell 4.0 for System Center 2012 R2 installed.
- Install the Operations console on the VMM management server.
- Import the following management packs to Operations Manager:
 - Windows Server 2003 Internet Information Services (IIS 6.0).
 - Windows Server 2008 Internet Information Services (IIS 7.0) .
 - Windows Server Internet Information Services Library.
 - SQL Server Core Library.

To configure Operations Manager and VMM integration:

- Check the Windows PowerShell version
- Install the Operations console on the VMM management server
- Import the necessary management packs
- Set up integration
- Change the management pack version in the registry; if using 2012 R2 version, this is not necessary



Note: You must install the Operations Manager agent on the VMM management server and all VMM hosts. However, you do not need to install the Operations Manager agent on the virtual machines.

Once you have verified that your system meets the prerequisites, you can set up and configure the integration by using the following steps:

1. In the VMM console, open the **Settings** workspace.
2. In the Settings pane, click **System Center Settings**, and then click **Operations Manager Server**.
3. On the **Home** tab, in the **Properties** group, click **Properties**.
4. Review the information in the **Introduction** page, and then click **Next**.
5. In the **Connection to Operations Manager** page, enter the server name for a management server in the management group, and select an account to use to connect. You can use the VMM server service account or specify a **Run As** account. This account must be a member of the Operations Manager Administrator role.
6. Select **Enable Performance and Resource Optimization (PRO)**, if desired.
7. Select **Enable maintenance mode integration with Operations Manager**, if desired, and then click **Next**.
8. Enter credentials for Operations Manager to connect with the VMM management server, and then click **Next**. This account will be added to the Administrator user role in VMM.
9. Review the information in the **Summary** page, and then click **Finish**.

After you connect to Operations Manager, if you are running System Center 2012 or System Center 2012 Service Pack 1 (SP1), and you update your VMM management packs, you must update the registry on your VMM management server. If you are running System Center 2012 R2, you do not need to update the registry. To make these necessary changes, perform the following steps:

1. Review the System Center Monitoring Pack for System Center 2012 – Virtual Machine Manager guide to get the version number of the VMM management pack. You can see the current version number of the monitoring pack in the Operations console. Make sure the guide and the management pack refer to the same version number. You will need this number later on step 6.
2. Open the **Administration** workspace.
3. In the Administration pane, click **Management Packs**. Scroll down to find the VMM monitoring packs, such as System Center 2012 Virtual Machine Manager Discovery.
4. On the VMM management server, click **Start**, in the **Search programs and files** text box, type **regedit**, and then press Enter. If the **User Account Control** dialog box displays, click **Yes** to continue.
5. In Registry Editor, locate the following registry key:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Microsoft System Center Virtual Machine  
Manager Server/Setup
```

6. Under the key, if necessary, update the string value **CompatibleMPVersion** to the new version number, and then close Registry Editor.

What Is Performance and Resource Optimization?

PRO is a feature that supports resource optimization in a virtualized environment based on the consumption of resources by your virtual machines. PRO provides a mechanism to link specific Operations Manager alerts to remediation actions that can be executed manually or automatically in VMM. These alerts are generated by monitors that you create in PRO-enabled management packs.

PRO-enabled management packs define the different target classes, groups, and monitors that monitor the performance and availability of virtual machines, hosts, applications, and hardware used in a virtualization environment. Microsoft and other vendors provide PRO-enabled management packs for some of their products. You can also create your own PRO-enabled management pack by following the guidelines provided by Microsoft in <http://go.microsoft.com/fwlink/?LinkID=285322>.

When an alert that targets a PRO class in Operations Manager is raised, a PRO tip is created in VMM. A PRO tip can include an action to fix the alert executed from Operations Manager. The action can be executed automatically or it may require approval from a VMM administrator. PRO tips can perform a live migration of a virtual machine from one node to another in a Hyper-V cluster based on performance thresholds, or to move the virtual machine to a completely new host in case of hardware failure.

The VMM monitoring pack includes several PRO-enabled management packs that define the basic classes and groups that support PRO and provide monitors to optimize the performance of hosts and virtual machines based on CPU and memory thresholds.

The VMM monitoring pack provides several PRO-enabled monitors that you can use. The following table shows two of the most common monitors.

Monitor	Threshold calculation	Sampling interval	Calculation
VMM Dynamic Memory VM Pressure	<ul style="list-style-type: none"> Virtual machine current memory / assigned memory Warning Level – 80% Critical Level – 100% 	300 seconds	Consecutive value of past 3 samples
VMM Maximum Dynamic Memory Monitor	<ul style="list-style-type: none"> Sum of virtual machine configured maximum memory values Warning Level – 125% Critical Level – 150% 	900 seconds	Current sample

Monitor	Threshold calculation	Sampling interval	Calculation
Virtual Machine Manager Dynamic Memory VM Pressure	Virtual machine current memory / assigned memory Warning Level – 80% Critical Level – 100%	300 seconds	Consecutive value of past 3 samples
Virtual Machine Manager Maximum Dynamic Memory Monitor	Sum of virtual machine configured maximum memory values Warning Level – 125% Critical Level – 150%	900 seconds	Current sample

MCT USE ONLY. STUDENT USE PROHIBITED

Comparing PRO and Advanced Placement Options in VMM

Advanced Placement Options in VMM without Operations Manager

Although using PRO tips works well, they require integration with Operations Manager. This means that customers without Operations Manager cannot use them. For customers who do not use Operations Manager, VMM offers support for the following resource optimization settings:

- Host reserves. Host reserves provides a way to allocate resources for the host operating system so that, when a virtual machine is created or moved to host, the action only occurs if the resources allocated for the host operating system can be maintained.
- Dynamic optimization. Use dynamic optimization to migrate virtual machines automatically from one host to another based on resource thresholds.
- Power optimization. Use power optimization to migrate virtual machines from one host to another during periods of low utilization so that a host can be turned off to save energy.

Options in VMM

- Host reserves
- Dynamic Optimization
- Power Optimization

Integration with Operations Manager

- PRO tips
- PRO-enabled management packs
- Flexible framework based on monitors

Using PRO With Operations Manager

PRO provides a lot more flexibility than dynamic optimization or power optimization, because it can monitor any object in the VMM environment before taking action. For example, with PRO, you can have a monitor check the performance of the storage that hosts files for a virtual machine. If the virtual machine's performance falls below or exceeds an established value, then PRO can migrate the virtual machine. This cannot be achieved with Dynamic Optimization or Power Optimization.

Before deciding to implement PRO, you must consider the following factors:

- Dynamic optimization and power optimization. If you have no requirements other than monitoring processor and memory usage, dynamic optimization and power optimization are a better option than PRO.
- Customization. Ensure that you understand the capabilities available to you through the VMM monitoring packs. The management pack guide for the VMM monitoring pack lists the rules and monitors available with the management pack. You must create anything else that you might need for your environment.
- Conflicting settings. If you currently use dynamic optimization and power optimization and then integrate with Operations Manager, you must either disable dynamic optimization and power optimization, or change their settings to avoid possible conflicts with PRO-enabled management pack monitors.

Lab: Implementing a Server Monitoring Strategy

Scenario

A. Datum Corporation has completed their initial deployment of the Windows Server 2012 infrastructure. A. Datum Corporation now wants to implement a strategy for monitoring the IT environment. A. Datum is evaluating the use of the monitoring tools in Windows Server 2012 for some of the branch offices. In addition, A. Datum has installed Operations Manager in the London data center to monitor servers in the data center.

You need to implement server monitoring using the Windows Server 2012 tools, and using Operations Manager.

Objectives

- Configure server monitoring using Windows Server 2012 tools.
- Implement Operations Manager monitoring.
- Configure the Operations Manager monitoring components.
- Monitor server performance.

Lab Setup

Estimated Time: 60 minutes

Virtual machines:

20414C-LON-HOST1, 20414C-LON-DC1,

20414C-LON-OM1, 20414C-LON-SVR1,

20414C-LON-SVR2, 20414C-TOR-SVR1,

20414C-TOR-SS1

User name: **Adatum\Administrator**

Password: **Pa\$\$w0rd**

Lab Setup

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

1. On LON-HOST1, start **Hyper-V Manager**.
2. In Microsoft Hyper-V® Manager, click **20414C-LON-DC1**, and, in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Log on using the following credentials:
 - a. User name: **Adatum\Administrator**
 - b. Password: **Pa\$\$w0rd**
5. Repeat steps 2 to 4 for the rest of virtual machines.

Exercise 1: Configuring Server Monitoring Using Windows Server 2012

Scenario

A. Datum has not yet decided whether they will also deploy Operations Manager in the Toronto data center, so they are first evaluating the available Windows Server 2012 tools. A. Datum would like to evaluate the options for monitoring multiple servers using a single interface, and for collecting information from multiple servers in a single location.

The main tasks for this exercise are as follows:

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

► Task 1: Configure Server Manager to monitor multiple servers

1. On TOR-SVR1, in the **Add other servers to manage** hyperlink on the **Server Manager** front page, configure **Server Manager** to manage TOR-SS1, and then create a server group named **Toronto Servers** that contains **TOR-SVR1** and **TOR-SS1**.
2. Run a Best Practices Analyzer scan on **TOR-SS1**.
3. Configure the performance alerts for TOR-SS1 using the following settings:
 - CPU (% usage): 75
 - Memory (MB available): 100
4. Add the Web Server (IIS) role to TOR-SS1.

► Task 2: Configure a data collector set

1. On TOR-SVR1, in Server Manager, go to **Computer Management** and configure a data collector set with the following settings:
 - Name: **Main Resources**
 - Template: **System Performance**
2. Configure the data collector using the following settings:
 - Schedule: **Mon to Fri starting at 8:00 AM**
 - Stop condition: **Overall duration of 10 hours**.
3. Collect data for 10 minutes, then stop the collector and analyze the report.

► Task 3: Configure an event subscription

1. On TOR-SS1, go to **Administrative Tools** and configure the Windows Firewall with Advanced Security to allow DCOM-In and Remote Event Log Management items.
2. On TOR-SVR1, configure an event subscription to TOR-SS1 with the following settings:
 - Name: **TOR-SS1 Events**
 - Type of events: **Critical and Error**
 - Event logs: **Windows Logs**
3. Add the computer account for TOR-SVR1 to the local **Event Log Readers** group in TOR-SS1.
4. View forwarded events. It may take several minutes for events to be forwarded.

Results: After completing this exercise, you should have configured monitoring for the servers by using the tools available within the Windows Server 2012 operating system.

Exercise 2: Implementing Operations Manager Monitoring

Scenario

A. Datum has deployed a single Operations Manager server in the London data center. You now need to deploy the Operations Manager agent for servers in the London data center.

The main tasks for this exercise are as follows:

1. Deploy the Operations Manager agent to virtual machines.
2. Configure agentless monitoring for host machines.

► Task 1: Deploy the Operations Manager agent to virtual machines

1. Deploy the Operations Manager agent to the following computers by using discovery:
 - LON-SVR1
 - LON-SVR2
2. After the install is complete, on, open the Services administrative tool and verify that the **Microsoft Monitoring Agent** service has started.

► Task 2: Configure agentless monitoring for host machines

3. Configure LON-OM1 to place manually-installed agents in a pending state.
4. On LON-DC1, manually install the Operations Manager agent using the following settings:
 - Agent source: **\\LON-OM1\C\$\Program Files\System Center 2012\Operations Manager\Server\AgentManagement\amd64\MOMAgent.msi**
 - Management group: **Adatum**
 - Management server: **LON-OM1**
5. In Server Manager, verify that the System Center Management service is running.
6. On LON-OM1, approve LON-DC1 in the Operations console and verify that LON-DC1 is now agent-managed.

Results: After completing this exercise, you should have installed and verified the Operations Manager agent on computers in the London data center.

Exercise 3: Configuring the Operations Manager Monitoring Components

Scenario

Now that you have configured Operations Manager to monitor the servers in the London data center, the next step is to deploy and configure the monitoring components.

The main tasks for this exercise are as follows:

1. Install and configure management packs.
2. Configure notifications, subscribers, and subscriptions.
3. Configure reports.
4. To prepare for the next module.

► **Task 1: Install and configure management packs**

1. On LON-OM1, import the SQL Server management packs from **C:\Program Files (x86)\System Center Management Packs\System Center Monitoring Pack for SQL Server**.
2. Create a management pack for overrides named **SQL Server 2008 (Monitoring) – Overrides**.
3. Create an override for the **DB File Space** unit monitor in the list of performance monitors for the **SQL Server 2008 DB File** target based on the following parameters:
 - Lower Threshold: **20**
 - Upper Threshold: **30**

► **Task 2: Configure notifications, subscribers, and subscriptions**

Configure Notifications

1. Create a notification channel with the following settings:
 - Channel type: **Email**
 - SMTP Server: **LON-SVR1.adatum.com**
 - Return address: **om@adatum.com**

Configure Subscribers

2. Create a subscriber with the following settings:
 - Name: **ADATUM\Administrator**
 - Schedule: Mon-Sat from 8:00 AM to 8:00 PM
 - Channel type: **Email**
 - Email address: **administrator@adatum.com**

Configure Subscriptions

3. Create a subscription with the following settings:
 - Name: **Critical SQL Alerts**
 - Criteria: **raised by the SQL Server 2008 Computers group and of severity Critical**
 - Channel: **SMTP Notification Channel**
 - Subscriber: **ADATUM\Administrator**

► **Task 3: Configure reports**

- Verify performance data collection from Exercise 1, Task 1.

► **Task 4: To prepare for the next module**

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

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20414C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20414C-LON-SVR1**, **20414C-LON-SVR2**, **20414C-LON-OM1**, **20414C-TOR-SS1**, and **20414C-TOR-SVR1**.

Results: After completing this exercise, you should have imported and configured the Microsoft SQL Server® Management pack in Operations Manager.

Question: What is the purpose of the channel object when creating notifications?

Question: Can you integrate System Center 2012 SP1 Operations Manager with System Center 2012 R2 Virtual Machine Manager?

Module Review and Takeaways

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
Management pack does not exist for a specific function. Difficult to update overrides for alerts.	

Module 7

Planning and Implementing High Availability for File Services and Applications

Contents:

Module Overview	7-1
Lesson 1: Planning and Implementing Storage Spaces	7-2
Lesson 2: Planning and Implementing DFS	7-7
Lesson 3: Planning and Implementing NLB	7-14
Lab: Planning and Implementing High Availability for File Services and Applications	7-19
Module Review and Takeaways	7-27

Module Overview

Highly available file service infrastructures usually require a large investment in hardware and software. With Storage Spaces, the Distributed File System (DFS), Network Load Balancing (NLB), and the use of the Internet SCSI (iSCSI) technologies in Windows Server® 2012, Microsoft is enabling companies of all sizes to build highly available file service infrastructures, without investing in new and expensive hardware.

Objectives

After completing this module, you will be able to:

- Plan and implement Storage Spaces.
- Plan and implement DFS.
- Plan and implement NLB.

Lesson 1

Planning and Implementing Storage Spaces

In this lesson, you will learn about highly available storage infrastructures by using Storage Spaces.

Lesson Objectives

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

- Describe Storage Spaces.
- Explain storage space considerations and the configuration options for storage pools.
- Describe a clustered storage space.
- Describe the requirements for implementing clustered Storage Spaces.
- Describe the process for configuring clustered Storage Spaces.
- Explain considerations for implementing Storage Spaces.

What Are Storage Spaces?

Storage Spaces are a storage virtualization capability included in Windows Server 2012 and Windows® 8 and newer Windows operating systems. It is a feature that is available for both the NTFS file system and the Resilient File System (ReFS) volumes, providing redundancy and pooled storage for internal and external drives of differing sizes and interfaces. You can use Storage Spaces to add physical disks of any type and size to a storage pool, and then create highly available virtual disks from that storage pool. The primary advantage of Storage Spaces is that you manage multiple disks as one unit, not single disks.

Components of Storage Spaces include:

- Physical disks
- Storage pool
- Virtual disks

New features in Windows Server 2012 R2

- Storage tiers
- Write-back cache
- Parity space support for failover clusters
- Dual parity
- Automatic rebuild of storage spaces from storage pool free space

Storage Spaces

Storage Spaces consist of the following components:

- **Physical disk.** *Physical disks* are actual disks, such as Serial ATA (SATA) or serial-attached SCSI disks. If you want to add physical disks to a storage pool, the disks must meet the following requirements:
 - One physical disk is required to create a storage pool, and a minimum of two physical disks are required to create a resilient mirror virtual disk. If your server is a virtual machine, you can add any .vhdx or .vhdx file presented to the virtual machine as a physical disk.
 - A minimum of three physical disks are required to create a virtual disk with resiliency through parity.
 - Three-way mirroring requires at least five physical disks.
 - Disks must be blank and unformatted; no volume can exist on them.
 - You can attach disks by using a variety of bus interfaces including iSCSI, serial-attached SCSI, SATA, SCSI, and universal serial bus (USB).



Note: If you want to use failover clustering with storage pools, you cannot use SATA, USB, or SCSI disks.

- **Storage pool.** A *storage pool* is a collection of one or more physical disks that you can use to create virtual disks. You can add any available physical disk that is not formatted or attached to another storage pool.
- **Virtual disk (or storage space).** A *virtual disk* is a nonphysical disk that is similar to a physical disk from the perspective of users and applications. However, virtual disks are more flexible because they include thin provisioning or just-in-time (JIT) allocations. In addition, they include resiliency to physical disk failures with built-in functionality such as mirroring.

What Is New for Storage Spaces in Windows 2012 R2?

- **Storage tiers.** Storage tiers automatically move frequently accessed data between different storage tiers. For example, you implement one tier using solid-state drives and another using SATA drives. Infrequently accessed data will be moved automatically to SATA.
- **Write-back cache.** Write-back cache provides buffering of small random writes to solid state drives to reduce latency of write operations.
- **Parity space support for failover clusters.** Windows Server 2012 R2 supports parity spaces in failover clusters.
- **Dual parity.** Dual parity allows you to store two copies of the parity across a parity space, improving resilience. With dual parity, a parity space can still work after the failure of two disks, similar to a redundant array of independent disks (RAID) 6 array.
- **Automatically rebuild storage spaces from storage pool free space.** Administrators are no longer required to add hot spare drives to rebuild data from failed disks in a storage pool. Instead of using hot spare disks, Windows Server 2012 R2 rebuilds the lost data on free space available in the pool, reducing the time it takes to recover from physical disk failures.

Storage Pool Configuration Options

You can configure virtual disks with different storage layout and provisioning types. The storage layout and provisioning type that you select during the virtual disk creation determines the virtual disk's reliability and performance.

Storage Layout

The storage layout of a virtual disk specifies how the storage space configures the underlying physical disks. Settings include:

- **Simple.** This setting provides a stripe set without parity, similar to RAID 0. Compared with a single disk, this configuration increases throughput and maximizes capacity. However, it does not provide any redundancy or protect data from a disk failure.
- **Mirror.** This setting provides a mirror set by duplicating data on two or three disks, similarly to RAID 1. It increases reliability with reduced capacity. This configuration requires at least two disks to protect data from a single disk failure or at least five disks to protect from two simultaneous disk failures.

Storage layout settings:

- **Simple:**
 - Better performance, no redundancy, and similar to RAID 0
- **Mirror:**
 - Slower performance, offers redundancy by copying data, and similar to RAID 1
- **Parity:**
 - Better performance, redundancy by parity, and similar to RAID 5

Provisioning type settings:

- **Thin:**
 - Flexible
- **Fixed:**
 - Better performance

- Parity. This setting provides a striped set with distributed parity by striping data and parity information across multiple disks, similar to RAID 5. It increases reliability with reduced capacity. This configuration requires at least three disks to protect data from a single disk failure. Furthermore, you cannot use this configuration in a failover cluster.

Provisioning Type

The virtual disk provisioning type determines how space is allocated to the virtual disk. This decision can affect the performance of the virtual disk. You can choose between thin or fixed provisioning:

- Thin. This provisioning allows the virtual disk to exceed the capacity of the disk pool by allocating a bigger size than the available amount of space in the disk pool at the time of creation. Although this provides flexibility, it does so at the cost of performance, because the Storage Spaces engine must reassign disk space from the pool as the virtual disk increases its consumption.
- Fixed. Fixed provisioning reserves space for the virtual disk at the time of creation, which maximizes performance. However, the disk's size must be less than the pool's amount of available space. As you add physical disks to the pool, you can extend the fixed disk, which will affect performance only during the extension.

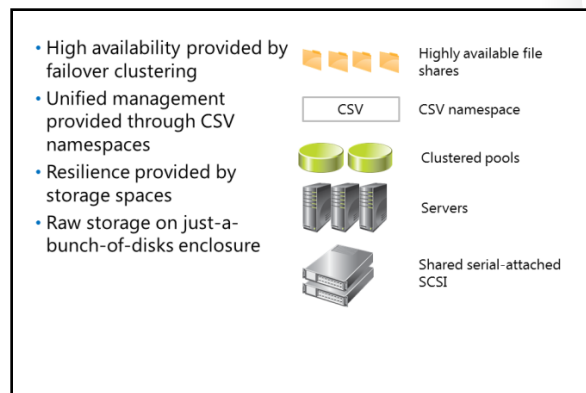
What Is a Clustered Storage Space?

A clustered storage space combines the storage spaces and failover cluster features to provide a highly available, resilient, and cost-efficient storage solution. You can use this concept to build a storage area network (SAN) by using a set of serial-attached SCSI just-a-bunch-of-disks enclosures. To provide for resilience and redundancy, you can connect to the just-a-bunch-of-disks enclosures from multiple servers, and each server should have redundant paths to each disk in the enclosure. Path redundancy can be achieved by using Microsoft® Multipath I/O

(MPIO). Because access to storage is clustered, you can use Cluster Shared Volumes (CSVs) to provide a single namespace for storage access, making it easier to manage a large set of disks as a single unit.

You can use clustered storage spaces to protect against the following risks:

- Server node failures. By using the failover clustering feature, you can provide high availability and fault tolerance services to your storage fabric. If any given server is offline, access to the storage remains available through the remaining nodes in the failover cluster.
- Data access failures. You can use multiple network interface cards in each server node, MPIO, redundant power supplies and even redundant disk enclosures to provide a high level of availability in case a single component in the solution fails.
- Data corruption failures. Both NTFS and ReFS provide features to manage data corruption. Although ReFS is better than NTFS for this purpose, you cannot use ReFS with CSVs.
- Physical disk failures. You can use mirror and parity storage spaces to provide physical disk failure tolerance.



Requirements for Implementing Clustered Storage Spaces

Clustered Storage Spaces use the failover cluster service to provide a highly available storage space environment. However, you must ensure that your environment meets certain hardware and software requirements before it can host Storage Spaces in a failover cluster. To create a clustered storage space, you must:

- Connect all physical disks to the nodes through a serial-attached SCSI interface.
- Have at least three physical disks, with at least 4 gigabytes (GB) each.
- Dedicate all physical disks to the storage pool.
- Ensure that Storage Spaces formatted in ReFS are not added to the CSV. CSVs are optional; you do not need to use them.
- Ensure that all physical disks pass the failover cluster validation tests.

To meet clustered storage space prerequisites, your environment must:

- Use serial-attached SCSI disks only
- Have at least three disks with 4 GB each
- Have dedicated disks
- Not format drives with ReFS
- Pass failover cluster validation tests

Configuring Clustered Storage Space

Once your environment has met the requirements for a clustered storage space, you can create a new storage pool to host your clustered storage space. However, instead of using Server Manager to create this storage pool, you must use the Failover Cluster Manager. Before you create a clustered storage space, you should:

1. Add the Failover Cluster role on all nodes attached to the physical disks that you will use in the storage space.
2. Add the File Services role and File Services Role Administration Tools to all nodes in the failover cluster.
3. Create a Scale-Out File Server cluster.

Once you configure the Failover Cluster role on all nodes that connect to the physical disks in the clustered storage space, you can implement a clustered storage space by performing the following steps:

1. Open the Failover Cluster Manager.
2. Expand **Storage**, right-click **Pools**, and then click **New Storage Pool**.
3. In the New Storage Pool Wizard, type a name for the new storage pool, and then click **Next**.
4. In the **Select physical disks for the storage pool** page, select at least three physical serial-attached SCSI disks with a minimum of 4 GB of available space, and then click **Next**.
5. On the **Confirmation** page, click **Close**. The pool will be added to the cluster and brought online.

To implement a clustered storage space:

1. Create a failover cluster
2. Add the File Services role to all nodes in the cluster
3. Create a new storage pool from the Failover Cluster Manager
4. Create virtual disks in the storage pool



Note: You can also create storage pools before you have a failover cluster, and later add the storage space to the cluster.

Once the pool is created, you can create virtual disks in the same manner you would in a regular storage space pool.

Considerations for Implementing Storage Spaces

Storage Spaces provide flexibility and ease of use when creating volumes. However, this same flexibility might cause performance problems, depending on how you combine different types of physical disks in the same virtual disk. You must consider the underlying physical implementation of a disk subsystem before configuring the different types of virtual disks available in a storage space. Consider the following options when creating Storage Spaces:

- Create separate storage pools based on disk type and speed. For instance, consider a storage pool with three disks: a SATA disk that connects with universal serial bus (USB) 2.0, and two serial-attached SCSI disks that attach directly. A parity virtual disk that you create by using this pool will have subpar performance because the USB disk will be much slower than the serial-attached SCSI disks. In this case, it would be better to use three serial-attached SCSI disks.
- Choose the appropriate storage layout based on the physical disks layout. For example, if you are using hardware RAID, you could have two volumes that present to the Windows Server operating system as two separate physical disks. However, each could be a RAID 5 or RAID 10 set, which provides fault tolerance. Choosing parity as a storage layout would be excessive. Simple storage would increase performance, while the physical level would provide fault tolerance.
- Use fixed provisioning for better performance. Fixed provisioning reserves the amount of space needed for a disk during provisioning, avoiding the cost of just-in-time (JIT) provisioning that thin provisioning requires.
- Use hardware RAID for clustered Storage Spaces. Clustered Storage Spaces do not allow the creation of virtual disks by using a parity layout on Windows Server 2012; however, you can do so in Windows Server 2012 R2.

- Create storage pools based on disk type and speed
- Choose the appropriate storage layout based on the physical disks layout
- Use fixed provisioning
- Use hardware RAID for clustered Storage Spaces

Lesson 2

Planning and Implementing DFS

The DFS feature in Windows Server 2012 allows you to expose shared folders hosted by different servers on your network as if they all belonged to one individual server. DFS eliminates the dependency on a server name when accessing shares and providing data replication among the servers that host shares. As a result, DFS creates a highly available storage system without the need to invest in specific fault tolerant hardware.

Lesson Objectives

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

- Describe DFS components.
- Review scenarios for using DFS.
- Understand guidelines for designing DFS namespace availability.
- Describe considerations for configuring referrals.
- Understand guidelines for optimizing DFS namespaces.
- Describe best practices for deploying DFS namespaces.
- Understand guidelines for designing DFS replication.

Overview of DFS Components

DFS is composed of DFS Namespaces and DFS Replication. Each component is a Windows Server 2012 server role and you can use them separately or together. You can manage DFS by using the DFS Management snap-in, or by using the Windows PowerShell® command-line interface.

DFS Namespaces

DFS namespaces provide the ability to group shared folders that are located on different servers into one or more logically structured namespace. Each namespace appears to users as a single shared folder with a series of subfolders. However, the underlying structure of the namespace can consist of numerous file shares that are located on different servers and in multiple sites.

DFS Replication

DFS Replication provides the ability to replicate folders (including those to which a DFS namespace path refers) across multiple servers and sites. DFS Replication uses a compression algorithm known as remote differential compression (RDC). RDC detects changes to the data in a file and enables DFS Replication to replicate the changed file blocks only, instead of the entire file.

Server roles:

- DFS Namespaces
- DFS Replication

Changes added in Windows Server 2012:

- Windows PowerShell for DFS namespaces
- Site awareness for DirectAccess clients
- WMI provider for DFS namespaces
- Replication for Data Deduplication volumes

Changes added in Windows Server 2012 R2:

- Windows PowerShell for DFS namespaces
- WMI provider for DFS Replication
- Replication changes



Note: You can use DFS Replication to replicate the SYSVOL folder in Active Directory® Domain Services (AD DS) for domains in the Windows Server 2008 domain functional level or newer.

New and Changed Functionality

Windows Server 2012 adds or changes the following features in DFS:

- DFS namespaces:
 - Site awareness for DirectAccess clients. Provides remote clients with optimal referrals to DFS content when the computers connect by using DirectAccess. Only Windows 8 clients support site awareness.
 - When a remote computer accesses a DFS namespace path by using DirectAccess in Windows 7 or Windows Server 2008 R2, remote computers with IP addresses outside the sites that AD DS specifies receive a randomly ordered referral. This referral could include servers in distant sites, even when servers in a nearby site are available.
 - When a remote computer accesses a DFS namespace path by using DirectAccess on computers that are running Windows 8 or Windows Server 2012, the computer provides a site name in the referral request to the namespace server that is running Windows Server 2012. The namespace server uses the site name to provide a referral to the closest site available
 - Management. Provides management methods based on Windows Management Instrumentation (WMI) to manage DFS namespaces.
 - Windows PowerShell module. Provides Windows PowerShell cmdlets to manage DFS namespaces.
 - Command-line tool, Dfscmd. This feature has been deprecated. You should use Windows PowerShell instead.
- DFS Replication. Support for Data Deduplication volumes. Provides support to replicate content stored on volumes that use Data Deduplication.
- File Replication Service (FRS). This feature has been deprecated. You should use DFS Replication instead.



What's New in DFS Namespaces and DFS Replication in Windows Server 2012:

<http://go.microsoft.com/fwlink/?LinkID=392398>

In addition to these new features in Windows Server 2012, Windows Server 2012 R2 includes the following enhancements to DFS Replication:

- Windows PowerShell module for DFS Replication. Windows Server 2012 R2 includes an entire module for managing most of the tasks related to DFS Replication by using Windows PowerShell.
- WMI provider. Allows management of DFS Replication by using WMI-capable tools.
- Database cloning for initial sync. Allows DFS to bypass initial replication when you create new replicated folders, replace servers, or recover from a disaster.
- Database corruption recovery. Allows DFS to rebuild corrupt databases without data loss caused by a nonauthoritative initial sync.
- Cross-file RDC disable. Allows you to disable cross-file RDC between servers.
- File staging tuning. Allows you to configure variable file staging sizes on individual servers.

- Preserved file restoration. Allows DFS to recover automatically after a loss of power or stoppage of the DFS Replication service.
- Membership disabling improvements. Stops DFS Replication private folder cleanup when a server's membership is disabled in a replication folder.

 **For more details on the enhancements to DFS Replication in Windows Server 2012 R2, visit:**

<http://go.microsoft.com/fwlink/?LinkID=392399>

Scenarios for Using DFS

By using DFS Namespace and DFS Replication, your organization can benefit from several implementation scenarios, including:

- Sharing files across branch offices
- Data collection
- Data distribution

Sharing Files Across Branch Offices

Usually, organizations with multiple physical offices share files or collaborate between offices.

You can use DFS Replication to replicate files

between these offices or between branch offices and a hub site. This form of replication reduces wide area network (WAN) traffic and provides high availability should a WAN link or a server fail. DFS Replication ensures that when a user makes changes to a file, delta replication occurs so that the changes replicate to all other sites.

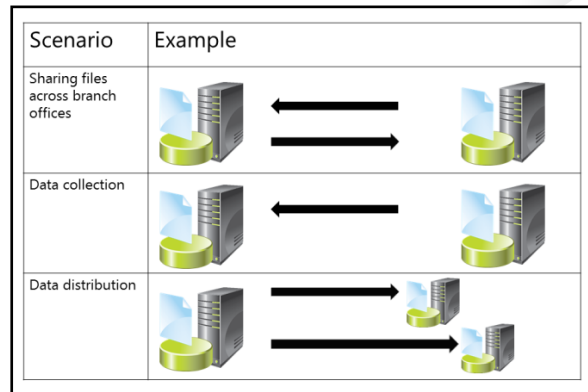
 **Note:**

- Remember that DFS Replication replicates data only after the user closes the file to which he or she made changes. Therefore, we do not recommend DFS Replication for files that remain open for long periods, such as database files and virtual hard drive files that virtual machines use.
- It is important to remember that if a file changes in multiple locations, DFS Replication maintains the changes made by the file with the latest timestamp. This means that changes can be lost in DFS Replication.
- Delta replication copies only the changes committed to a file, therefore using less bandwidth for replication.

Data Collection

You can use DFS Replication to replicate files from all branch offices to a hub site. Once in the hub site, you can back up all data at once, thereby providing a centralized way of recovering data. You can use this centralized backup set as a second backup source if you are already backing up data at the local site level. You can reduce backup hardware, software, and personnel costs by providing the backup at the hub site only.

You can use DFS Namespace to provide fault tolerance to file access. That way, if a branch office server fails, the clients can still access the files from the hub site.



Data Distribution

You can use DFS Replication to replicate data from a hub site to multiple branch sites, which moves the data closer to the user. You can use DFS Namespace to ensure that client computers access the shares closest to them, thereby reducing bandwidth consumption by using site awareness.

Guidelines for Designing DFS Namespace Availability

When users require access to folders in a DFS namespace, DFS redirects them to the folder targets in the namespace. If a namespace is unavailable, users are not able to access files. You should make a namespace highly available to ensure that users have uninterrupted access to files. Consider these guidelines for implementing DFS namespace availability:

- AD DS stores domain-based namespaces, which multiple servers host. Hosting a namespace on multiple servers makes the namespace highly available. All servers that are hosting a namespace must be in the same domain.
- A single server hosts stand-alone namespaces, while configuration data is stored locally. You can use failover clustering to increase the availability of a stand-alone namespace.
- Each folder in a namespace can have one or more targets. To increase availability, use multiple targets for each folder, and then use replication to synchronize data between folder targets.
- When DFS namespaces are in a domain that is running at a domain functional level lower than Windows Server 2008, users can experience performance problems, such as having a limit of 5,000 folders in the namespace. If you must use a lower domain functional level for compatibility with older Windows operating systems, you should use stand-alone namespaces to increase scalability beyond 5,000 folders.
- When DFS namespaces are in a domain that is running at the functional level of Windows Server 2008 or newer, a single namespace does not experience performance degradation until the namespace reaches or exceeds approximately 300,000 folders.

- Increase availability for domain-based namespaces by using multiple namespace servers
- Increase availability for stand-alone namespaces by using failover clustering
- Increase availability for folders by using multiple targets
- In a domain running at a functional level lower than Windows Server 2008, use stand-alone namespaces if there are more than 5,000 folders
- In a domain running at the functional level of Windows Server 2008 or newer, you can include up to 300,000 folders in a domain-based namespace



Windows Server DFS-Namespaces Performance and Scalability:

<http://go.microsoft.com/fwlink/?LinkID=285325>

Considerations for Configuring Referrals

Referral ordering controls the order in which Windows Server 2012 presents folder targets to clients. Each folder uses the configuration of the DFS namespace by default, but you can override this setting.

When a client accesses a folder, Windows Server 2012 presents a list of folder targets to the client. You determine the order of the folder targets in the list by configuring settings for both referrals and target priority.

Considerations for referral options are as follows:

Considerations for referral options include:

- Lowest-Cost Referral
 - Site link costs determine referral
- Random Referral
 - Any available target outside the local site is acceptable
- Exclude Targets Outside Of The Client's Site Referral
 - Clients never use a folder target outside of their site

- Use the Lowest-Cost Referral option when site link costs accurately reflect the order in which you want users to access targets. When you select the Lowest-Cost Referral option, clients use a folder target in the local site first. If there is no target in the local site, clients choose a target in another site based on the lowest site link costs.



Note: If multiple site links have the same cost, Windows Server 2012 selects a site randomly for the client.

- Use the Random Referral option when any available target outside the local site is acceptable. When you select the Random Referral option, clients use a folder target in the local site first. If there is no target in the local site, clients choose a target in another site randomly.
- Use the Exclude targets outside of the client's site referral option if you never want clients to use a folder target outside their site. When you select this option, clients use folder targets in the local site only. If a folder does not have targets in the local site, clients cannot access the folder.

You can use target priority for fine-tuning referrals and for overriding the referral-ordering configuration. Considerations for target priority options include:

- Use the First among targets of equal cost option when multiple targets are in the same site and you want one target to be the primary target. This option is useful when you want to avoid replication conflicts. For example, you might have two copies of data that you store in a site for high availability.
- Use the Last among targets of equal cost option when multiple targets are in the same site and you want clients to use one target only when other targets are unavailable. For example, you might configure one target for use as a back-up archive only.
- Use the First among all targets option to force Windows Server to choose a target as the preferred target, regardless of the cost associated with the site. This is useful when multiple targets exist in multiple sites, but you want one target to be the primary target to avoid replication conflicts. For example, you may keep a primary copy of the data at the head office, which users should always use if it is available.
- Use the Last among all targets option to force users to use a target as the last choice, regardless of the cost associated with the site. This is useful when multiple targets exist in multiple sites, but you want users to use one target only when others are unavailable. For example, you might use one target for a back-up archive only.

Guidelines for Optimizing DFS Namespaces

In addition to referral ordering and target priority, you should consider other options for optimizing DFS namespaces:

- Disable referrals for a folder target when performing maintenance on a server that is hosting that folder target. This ensures that clients do not disconnect from the folder target as you restart the server.
- Enable failback on a folder if you want clients to move back to a preferred target once the preferred target becomes available. This is useful when the preferred target is local and other targets are available only over a WAN link.
- Clients cache referrals to DFS namespace and folders. By default, DFS caches referrals for DFS namespace for five minutes and referrals for folders for 30 minutes. You can shorten these values to ensure that clients find new DFS namespace and folder targets more quickly.
- Namespace servers use namespace polling to retrieve configuration data from AD DS. You have the option to optimize DFS namespace polling for consistency or scalability. When you optimize namespace polling for consistency, all namespace servers retrieve configuration data from the primary domain controller (PDC) emulator. When you optimize namespace polling for scalability, namespace servers retrieve configuration data from a local domain controller to reduce the load on the PDC emulator.

Guidelines for optimizing DFS namespaces include:

- Disabling referrals during server maintenance
- Enabling failback when using multiple servers
- Shortening caching interval for referrals in environments where you add new namespace servers frequently
- Using namespace polling for scalability to reduce the load on the PDC emulator

Best Practices for Deploying DFS Namespaces

Window Server 2012 provides several options for configuring DFS namespaces. Properly configured DFS namespaces make it easier for users to access files and avoid replication conflicts. Best practices for deploying DFS namespaces include:

- Using DFS namespaces to create a unified folder hierarchy. This makes it easier for users to locate files, because they do not need to browse multiple servers.
- Using multiple folder targets to increase availability of individual folders.
- Using the lowest-cost method for ordering target referrals. In most cases, it is preferable if users access files from a target that is within or close to the local Active Directory site.
- Using scalability mode for more than 16 namespace servers. Scalability mode ensures that requests from the namespace servers do not overload the PDC emulator.

- Use DFS namespaces to create a unified folder hierarchy
- Use multiple folder targets to increase availability of individual folders
- Use the lowest cost method for ordering target referrals
- Use scalability mode for more than 16 namespace servers
- Specify a primary server by using target priority to reduce replication conflicts
- Enable access-based enumeration

- Specifying a primary server by using target priority to reduce replication conflicts. If you specify a primary server, all users access files on a single server.
- Enabling access-based enumeration on a namespace to ensure that users see only namespace folders to which they have access. In addition, you should enable access-based enumeration on the folder targets.

Guidelines for Designing DFS Replication

You can control how DFS Replication performs replication. Designing DFS replication appropriately for your environment ensures acceptable performance.

The following are guidelines for designing DFS replication:

- Use a mesh replication topology only with fewer than ten members. In mesh replication, each server replicates its contents to all other servers. This reduces replication complexity and improves performance. With more than ten members, consider a hub and spoke replication topology.
- Use bandwidth throttling to ensure that replication does not overwhelm WAN links, especially when the WAN link speed is slow.
- Use cross-file RDC to reduce replication traffic. Cross-file RDC recognizes patterns in multiple files and uses those patterns to reduce replication. The cross-file feature is available on all editions of Windows Server 2012.
- Use replication filters to prevent replication of unwanted file types. Replication filters reduce replication traffic by restricting replication based on file extensions. For example, if you do not require media files, you could prevent their replication.
- Use read-only replicated folders when you do not intend clients to modify a replica. For example, when a central server is gathering files for backup, using read-only replicated folders eliminates the need to configure Share permissions manually as Read-Only.
- Ensure that you size Staging folders and Conflict and Deleted folders appropriately. The Staging folder with a 4 GB default size should be at least twice the size of the largest replicated file. If the default size of the Conflict and Deleted folder, which is 660 megabytes (MB), is too small, DFS Replication may purge conflicts before addressing them. Once both folders reach 90 percent usage, DFS Replication purges them until they reach 60 percent usage.
- Use multiple smaller replicated folders rather than one large replicated folder. If you must restart replication, it is much faster to restart replication on a folder with less content.

- Use a mesh replication topology with fewer than ten members only
- Use bandwidth throttling to ensure that replication does not overwhelm WAN links
- Use cross-file RDC to reduce replication traffic
- Use replication filters to prevent replication of unwanted files
- Use read-only replicated folders to prevent content modification
- Size Staging folders and Conflict and Deleted folders appropriately
- Use multiple smaller folders to simplify restarting

Lesson 3

Planning and Implementing NLB

NLB provides high availability and scalability for TCP/IP-based applications by sharing the application load among two or more servers. If a server fails in an NLB cluster, other existing servers will assume the failed server's load, providing high availability. If the number of users for the application increases, you can add more servers to the cluster to allow the application to scale-out and provide availability.

Lesson Objectives

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

- Describe NLB scenarios.
- Describe considerations for configuring NLB network settings.
- Describe considerations for configuring port rules.
- Describe considerations for data storage for NLB clusters.
- Provide considerations for deploying NLB on virtual machines.
- Describe managing NLB components in Microsoft System Center 2012 Virtual Machine Manager (VMM).

NLB Scenarios


You can use NLB to provide high availability and scalability for applications and services that use the TCP/IP network protocol on a specific Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) port. Most organizations use hardware load balancers to provide high availability and scalability for critical web applications. These same organizations use the NLB service to support smaller applications without incurring extra hardware costs.

Environments in which you can use NLB include:

- Online Responder service in a Windows public key infrastructure (PKI) environment.
- Regular Internet Information Services (IIS) web applications with fewer than 10 front-end servers.
- File Transfer Protocol (FTP) servers.
- Microsoft System Center 2012 Service Manager Self-Service Portal.
- Microsoft System Center 2012 Configuration Manager management points and software update points.
- Microsoft System Center 2012 Operations Manager management servers.
- System Center 2012 Service Manager management servers.

NLB scenarios:

- Online Responders in a Windows PKI environment
- Regular IIS web applications
- FTP servers
- Service Manager Self-Service Portals
- Configuration Manager management points and software update points
- Operations Manager management servers
- Service Manager management servers

 **Note:** NLB can also be used for other highly available services, such as Microsoft SharePoint® web front-end servers, and web sites hosted over more than 10 servers. However, you may prefer to use hardware load balancers for such scenarios.


Considerations for Configuring NLB Network Settings

NLB network settings define the management of the cluster virtual IP's media access control (MAC) address. Network packets are sent to the cluster virtual IP, because all nodes in the cluster must handle the packets. There are three options for network settings in NLB: unicast, multicast, and Internet Group Management Protocol (IGMP) multicast.

Unicast


Unicast mode is the default network setting for NLB. In this mode, NLB replaces the actual MAC address of each server in the cluster with a common NLB MAC address. We call this operation *MAC address spoofing*. With Mac address spoofing, any packets that are sent to the cluster's virtual IP address will be delivered by using the same MAC address. Therefore, all servers in the cluster will receive the packet. However, if you are using the same switch for your nodes, you might encounter a problem because each port in a switch must have its own unique MAC address. NLB resolves this problem by assigning a separate MAC address for each node. NLB changes the second octet of the common MAC address to match the cluster member ID of the node. Then NLB uses an Address Resolution Protocol (ARP) multicast to send data to the cluster.

- Unicast mode works on all hardware environments
- Unicast mode does not allow you to stop port flooding at the switch level
- Multicast supports a single network adapter
- Multicast allows you to stop port flooding at the switch level by using static ARP entries
- IGMP multicast allows you to use IGMP snooping
- Upstream routers might not support mapping a unicast IP address with a multicast MAC address

 **Note:** Unicast mode requires the use of two network adapters.

Multicast

In multicast mode (multicast or IGMP multicast), NLB assigns a multicast IP address for the nodes in the cluster. You can use this mode with a single network adapter per node.

 **Note:** You can use IGMP multicast in a Windows NLB cluster only if the switch to which the nodes are connected supports IGMP multicast.

Considerations for Configuring NLB Network Settings

The following is a list of considerations for configuring NLB network settings:

- Unicast mode works on all hardware environments.
- Unicast mode does not allow you to stop port flooding at the switch level.
- Multicast supports a single network adapter solution per node.
- Multicast allows you to stop port flooding at the switch level by using static ARP entries.
- IGMP multicast allows you to use IGMP snooping.
- Upstream routers might not support mapping a unicast IP address (the cluster IP address) with a multicast MAC address. In these situations, you must upgrade or replace the router. Otherwise, the multicast method is unusable.

Considerations for Configuring Port Rules

When implementing NLB cluster hosts, consider the following port rule options:

- Create a separate set of port rules per host.
- Avoid adding more ports than necessary for each host. Create a separate port rule per contiguous port range, using either TCP or UDP.
- Avoid using port rules that open both UDP and TCP ports.
- Use None for the affinity option if the application does not require state management at the server level.
- Use Single for the affinity option if the application requires state management at the server level and if all clients are local.
- Use Network for the affinity option if the application requires state management at the server level and the clients are behind a proxy server. Affinity defines what node a client connects to after establishing a first connection with the server. For instance, if no affinity is used, each time the same client connects to the NLB cluster it may receive a response from different nodes. This is fine for applications that do not maintain a session state on the node. If the application maintains a session state, this could be an issue. For example, imagine that a client connects to a node and the node saves the client name to the application state. Then, when the client connects to a different node later, that node will not have the state information for the client. Always check with developers to ensure that you configure the NLB settings for affinity based on the application needs.
- Use Single Host for filtering when executing maintenance on other cluster nodes.

- Create a separate set of port rules per host
- Create a separate port rule per contiguous port range
- Avoid using port rules that open both UDP and TCP ports
- Use Single Host for filtering when executing maintenance on other cluster nodes
- For the affinity options:
 - Use None if the application does not require state management
 - Use Single if the application requires state management and the clients are all local
 - Use Network if the application requires state management and the clients are behind a proxy server

Considerations for Configuring Data Storage for NLB Clusters

The most common use of NLB is to provide high availability and scalability to websites and FTP sites. Typically, IIS hosts these sites. IIS sites provide access to data stored in the file system through the HTTP, HTTPS, or FTP protocols. Therefore, each node in an NLB cluster must have access to the same data.

When designing a storage solution for NLB in IIS, consider the following:

- Use DFS Replication to replicate the data between nodes. This maintains a copy of the data in each node. However, it does not provide fault tolerance if the disk subsystem fails on a node and the NLB node is still active.
- Use a file server in the cluster to provide storage to all cluster nodes. This provides a single storage point for all nodes, but does not scale-out well when adding more nodes.
- Use separate file servers and DFS Replication to provide high availability in case of disk loss, and to reduce the replication between nodes.

- Use DFS Replication to replicate data among all nodes
- Use failover clustering with the file server role to store the data away from the nodes
- Use DFS Replication and DFS namespaces to replicate data to a separate set of file servers

Considerations for Deploying NLB on Virtual Machines

Physical servers are not the only server type that can benefit from the use of NLB. You can use NLB in virtualized servers in the same way you use it in physical servers. However, in a virtual environment, there are a few guidelines to consider:

- Use separate virtual networks for the private and public networks that NLB uses. You can use a private network for the internal traffic, and an external network for public traffic.
- Ensure that you enable spoofing of MAC addresses in the network settings for the adapter that you use for the NLB cluster. By default, in Hyper-V®, virtual machines are not able to spoof their MAC address. The reason behind this is that if a virtual machine is able to spoof its MAC address, it could potentially use the MAC address of an existing virtual machine in the environment and cause a denial of service (DoS) attack or an information disclosure attack. However, NLB uses MAC spoofing to allow the nodes in an NLB cluster to configure their own network interface cards with the same MAC address. This is so all nodes receive traffic from the switch they are connected to. Therefore, you must enable MAC spoofing for virtual machines in an NLB cluster. To do this, follow these steps:
 1. From **Hyper-V Manager**, select the virtual machine that will be part of a Hyper-V cluster.
 2. From the Actions pane, click **Settings**.
 3. Expand either the **Network Adapter** or **Legacy Network Adapter** node.
 4. Click on **Advanced Features**.
 5. Select **Enable MAC address spoofing**, and then click **OK**.
- Consider using VMM to manage NLB for virtual machines. You will learn more about VMM and NLB in the next topic.

- Use separate virtual networks for the private and public NLB networks
- Enable MAC address spoofing
- Use VMM to manage NLB usage through service templates

Managing NLB Components in VMM

VMM includes integration with load balancing to provide virtual machines provisioned by using VMM access to load balancing services. You can integrate VMM with Windows NLB or hardware load balancers. As soon as you install VMM, it already includes NLB with round robin as a load-balancing method. For hardware load balancers, you must download and install the load balancer provider for the specific hardware used.

Once you install all required load balancer providers, you can create Virtual IP (VIP) templates. A VIP template is associated to a load balancer, or NLB, and provides the necessary settings for virtual machines to use network load balancing.

- Create a VIP template
- Create a logical network
- Create static IP address pools
- Configure a physical network adapter on the host to use the logical network

After creating a VIP template, you can use the template on service templates. A service template contains the settings of one or more virtual machines used to provide a given service. For instance, you might have a service template used to deploy a web application that accesses data from a Microsoft SQL Server® database. This service template would contain two tiers: application, and database. The application tier can be composed of one or more web servers in an NLB cluster. Therefore, besides the virtual machine template for the web server, you also need a VIP template for the application tier. The database tier might contain two virtual machines configured with SQL Server and the database setup to be mirrored.

**Configuring Load Balancing in VMM Overview:**

<http://go.microsoft.com/fwlink/?LinkID=392400>

In summary, before you can use NLB with VMM, you must ensure that you meet prerequisites by:

- If a hardware load balancer is being used, install the load balancer provider for the hardware load balancer.

**How to Add Hardware Load Balancers in VMM:**

<http://go.microsoft.com/fwlink/?LinkID=392401>

- Creating a VIP template for NLB.

**How to Create VIP Templates for Network Load Balancing (NLB) in VMM:**

<http://go.microsoft.com/fwlink/?LinkID=392402>

- Creating a logical network and associating it to one or more network sites. Associate the network sites to host groups where you may deploy the service.
- Creating static IP address pools for the associated network sites where you will deploy the service.
- Ensuring that on the host where you will deploy the service, you configure a physical network adapter to use the logical network that the service uses.

Lab: Planning and Implementing High Availability for File Services and Applications

Scenario

One of the key requirements for the Windows Server 2012 deployment at A. Datum Corporation is to provide high availability for all critical network services and applications. A. Datum has identified file services and several web applications as critical services that must be highly available.

A. Datum has identified two requirements for high availability for file services:

- The loss of a single disk that is storing network files should not affect the availability of files stored on the disk.
- The loss of a single server that is providing file shares should not affect the availability of files stored on the server.

In addition, A. Datum requires that the failure of a single web server should not affect the availability of web applications.

Objectives

After completing this lab, students will be able to:

- Plan a high availability strategy for file services.
- Plan a high availability strategy for web applications.
- Implement and validate a high availability strategy for file storage.
- Implement a high availability solution using NLB.

Lab Setup

Estimated Time: 60 minutes

Virtual machines	20414C-LON-HOST1 20414C-LON-DC1 20414C-LON-CL1 20414C-LON-SVR1 20414C-LON-SVR2 20414C-TOR-SVR1 20414C-TOR-SS1
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 LON-HOST1, start **Hyper-V Manager**.
2. In Hyper-V Manager, click **20414C-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 two through four for the remaining virtual machines.

Exercise 1: Planning a High Availability Strategy for File Services

Scenario

You must plan for file services high availability.

Supporting Documentation

Email from Ed Meadows	
Charlotte Weiss	
From:	Ed Meadows [Ed@adatum.com]
Sent:	04 Feb 2013 09:05
To:	Charlotte@adatum.com
Subject:	Toronto office issues
<p>Charlotte,</p> <p>As you are well aware, during the last several months, we have experienced several unplanned outages, which have affected the availability of file services in the organization. These outages have primarily affected the Toronto branch office. To prevent future outages, we need to plan a file services deployment that provides full redundancy against the failure of any single component in the Toronto data center.</p> <p>To implement this solution, we have invested in a new iSCSI SAN in the Toronto office. The iSCSI SAN provides some redundancy, but the server team has also decided to implement both Storage Spaces and DFS to provide a higher level of redundancy.</p> <p>The storage space design must provide for redundancy in the event of the failure of any storage component. This includes servers, network devices, and storage enclosures.</p> <p>We have identified several critical file shares that must be available to users even if the data center in Toronto fails. If the data center fails, users should be able to access the files from a server in London.</p> <p>Regards, Ed</p>	
<p>Proposal</p> <ol style="list-style-type: none"> 1. How should you allow access to the iSCSI SAN in case of network switch failure? 2. How should you allow access to the iSCSI SAN in case of a network interface cards failure at the server level? 3. How should you configure your storage to allow access to data even if a physical disk fails? 4. How should you configure your solution to allow users to access a file share in London when the Toronto servers are offline? 	

The main tasks for this exercise are as follows:

1. Read the supporting documentation
2. Update the proposal document with your planned course of action
3. Examine the suggested proposals in the Lab Answer Key

► **Task 1: Read the supporting documentation**

Read the documentation provided.

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

Answer the questions in the proposals section of the Supporting Documentation.

► **Task 3: Examine the suggested proposals in the Lab Answer Key**

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

Results: After completing this exercise, you should have planned a high availability strategy for file servers.

Students will design the storage space deployment. The design should include storage units and those servers and network components that will be deployed. The design also should contain notes about highly available features that the deployment will include, such as NIC teaming on the servers, MPIO for the network path, and RAID levels on the storage enclosure.

Students will design the DFS deployment. The design will include server locations, the deployed DFS namespaces, and DFS targets. The design should include notes on the configuration for the DFS Replication configuration.

Exercise 2: Planning a High Availability Strategy for Web Applications

Scenario

You must plan for web application high availability based on the following:

Supporting Documentation

Email from Ed Meadows	
Charlotte Weiss	
From:	Ed Meadows [Ed@adatum.com]
Sent:	04 Feb 2013 09:05
To:	Charlotte@adatum.com
Subject:	Web application
<p>Charlotte,</p> <p>We have a new web application that must be accessible to all users in the London office. Our developers tell me that this web application uses session state. I am not sure what that really means, but what I do know is that the application must be available even if a single server fails. Furthermore, we</p>	

Email from Ed Meadows

need to have a copy of the application in Toronto, as backup in case we have disk issues in London. We have two servers in London, LON-SVR1 and LON-SVR2 that you can use to host the application. One more thing, our developers will be making updates to this application frequently. We want to make sure they do not need to copy their updates to each server. They will be updating the application on the TOR-SVR1 server.

Can you please suggest how to configure these three servers to make this happen?

Regards,
Ed

Proposal

1. How can you provide high availability for the web application?
2. How can you manage session maintenance?
3. How can you configure your servers to allow changes to be copied from TOR-SVR1 to the other servers?

The main tasks for this exercise are as follows:

1. Read the supporting documentation
2. Update the proposal document with your planned course of action
3. Examine the suggested proposals in the Lab Answer Key

► Task 1: Read the supporting documentation

Read the documentation provided.

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

Answer the questions in the proposals section of the A. Datum web application high availability plan.

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

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

Results: After completing this exercise, you should have planned a high availability strategy for web applications that meets these criteria:

Your NLB deployment design should include the server deployment and the storage design. You should use the information from the previous exercise to plan the storage design.

Your design should include the NLB port rules and network settings design.

Exercise 3: Implementing a High Availability Solution for File Storage

Scenario

You need to implement the storage space and DFS design that you created. After doing this, you need to validate the deployment, and then test redundancy.

The main tasks for this exercise are as follows:

1. Configure NIC Teaming
2. Configure iSCSI initiators and MPIO
3. Configure Storage Spaces by using iSCSI targets
4. Validate the high availability of the deployment against the loss of a single network adapter

► **Task 1: Configure NIC Teaming**

1. Sign in to TOR-SVR1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. From Server Manager, click Local Computer, then click the listed IP address to open Network Connections, and enable **Ethernet 2**, **Ethernet 3**, and **Ethernet 4**.
3. Create a new NIC team based on the following parameters:
 - Name: **iSCSI Access Team 1**
 - Adapters: **Ethernet** and **Ethernet 2**
4. Configure the TCP/IP settings for iSCSI Access Team 1 by using the following settings:
 - IP address: **172.16.1.20**
 - Subnet mask: **255.255.0.0**
 - Default gateway: **172.16.0.1**
 - Preferred DNS server: **172.16.0.10**
5. Create a new NIC team based on the following parameters:
 - Name: **iSCSI Access Team 2**
 - Adapters: **Ethernet 3** and **Ethernet 4**
6. Configure the TCP/IP settings for iSCSI Access Team 2 by using the following settings:
 - IP address: **131.107.1.10**
 - Subnet mask: **255.255.0.0**
 - Default gateway: blank
 - Preferred DNS server: **172.16.0.10**

► **Task 2: Configure iSCSI initiators and MPIO**

1. On TOR-SVR1, in Server Manager, start the Add Roles and Features Wizard, and then install the **Multipath I/O** feature.
2. In Server Manager, on the Tools menu, open **iSCSI Initiator**, and then configure the following:
 - Enable the **iSCSI Initiator** service
 - Quick Connect to target: **172.16.1.25**
3. In Server Manager, on the **Tools** menu, open **MPIO**, and then configure the following:
 - Enable **Add support for iSCSI devices** on **Discover Multi-paths**
4. After the computer restarts, sign in to TOR-SVR1 with the username **Adatum\Administrator** and the password **Pa\$\$w0rd**.
5. In Server Manager, on the **Tools** menu, click **MPIO**, and verify that **Device Hardware ID MSFT2005iSCSIBusType_0x9** displays on the list.

► Task 3: Configure Storage Spaces by using iSCSI targets

1. On TOR-SVR1, from Server Manager, create a storage pool that uses the following parameters:
 - Name: **iSCSIPool**
 - Disks:
 - **PhysicalDisk1**
 - **PhysicalDisk2**
 - **PhysicalDisk3**
 - **PhysicalDisk4**
2. In the iSCSIPool, create a virtual disk with the following parameters:
 - Name: **DFSDisk**
 - Storage layout: **Parity**
 - Provisioning type: **Fixed**
 - Size: **Maximum**
3. Create a volume with the following settings:
 - Drive letter: **H**
 - File system: **NTFS**
 - Name: **DFS Volume**

► Task 4: Validate the high availability of the deployment against the loss of a single network adapter

1. On TOR-SVR1, use Windows PowerShell to copy the **C:\windows\system32\notepad.exe** file to the **H** drive.
2. On TOR-SVR1, open File Explorer, and then click **DFSVolume (H:)**.
3. Verify that notepad.exe displays in the file list.
4. On the host machine, in Hyper-V Manager, in the Virtual Machines pane, right-click **20414C-TOR-SVR1**, and then click **Settings**.
5. In Settings for 20414C-TOR-SVR1, in the Hardware pane, click the first occurrence of **Network Adapter**. In the **Virtual Switch** drop-down list box, click **External Network**, and then click **OK**.

Results: After completing this exercise, you should have implemented a high availability solution for file storage.

Exercise 4: Implementing a High Availability Solution by Using NLB**Scenario**

You need to implement the NLB design that you created. To do this, you need to deploy an NLB cluster, configure the settings based on the design, and configure a test website to use the highly available storage. Then you need to validate the website's availability by simulating a server failure.

The main tasks for this exercise are as follows:

1. Configure an NLB cluster
2. Configure the Web servers to use highly availability storage
3. Validate the deployment
4. Prepare for the next module

► **Task 1: Configure an NLB cluster**

1. Sign in to LON-DC1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. Open Windows Internet Explorer®, and then verify that the web application is available on the following URLs:
 - **http://LON-SVR1.adatum.com**
 - **http://LON-SVR2.adatum.com**
3. Create a host (A) resource record in DNS for **www.adatum.com** that points to **172.16.0.111**.
4. Install the NLB feature on LON-SVR1 and LON-SVR2.
5. Create an NLB cluster by using the following parameters:
 - Name: **LON-NLB**
 - Operational Mode: **multicast**
 - Cluster IP address: **172.16.0.111**
 - Network interface on LON-SVR1: **Ethernet**
6. Add a second node to the cluster by using the following parameters:
 - Node name: **LON-SVR2**
 - Network interface: **Ethernet**
7. Configure the cluster to use a single port rule with the following settings:
 - Port range: **80 to 80**
 - Protocols: **TCP**
 - Filtering mode: **Multiple Host**
 - Affinity: **Single**

► **Task 2: Configure the Web servers to use highly availability storage**

1. On LON-SVR1, open the DFS Management console, and verify that it is configured with the following three namespace servers:
 - **LON-SVR1**
 - **LON-SVR2**
 - **TOR-SVR1**
2. Verify that replication is configured.
3. Close the DFS Management console.
4. Configure IIS on LON-SVR1 and LON-SVR2 to use the **\\adatum.com\website\wwwroot** share as the source for the default website.

► **Task 3: Validate the deployment**

1. Switch to LON-CL1, and then test the NLB cluster from Internet Explorer by using **http://www.adatum.com** as a URL.
2. Validate the affinity settings by typing your name in the **First name** text box. Click **OK**, and then click **Refresh** on the browser window several times to ensure the data that displays in the browser does not change. This ensures that you are connected to the same server each time you refresh, due to client affinity.
3. Stop the NLB cluster node to which you connected, and then refresh Internet Explorer to test high availability.
4. Start the node again, and then refresh Internet Explorer to test affinity.

► **Task 4: 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 Hyper-V Manager, on the **Virtual Machines** list, right-click **20414C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps two and three for 20414C-LON-SVR1, 20414C-LON-SVR2, 20414C-LON-CL1, 20414C-TOR-SS1, and 20414C-TOR-SVR1.

Results: After completing this exercise, you should have implemented a high availability solution by using NLB.

Module Review and Takeaways

Review Questions

Question: What are the prerequisites for a clustered storage space?

Question: What is one important consideration when deploying NLB by using virtual machines?

Question: What should you do to ensure that a group of shared files is available locally to users from different physical sites, and that all changes to the files are available in all sites?

Question: Your company needs a SAN for a virtualization project. Management is looking at all possible solutions for a SAN due to budget constraints. You are asked to provide a solution using Windows Server 2012 R2. What can you propose to create a SAN by using Windows Server 2012 R2 without buying a Fibre Channel SAN?

Question: You have a server named Server1 that runs Windows Server 2012 R2. Server1 has a hardware RAID array containing 20 drives. The RAID array presents three physical drives to the operating system. Each drive is a RAID 5 array. One drive is used as a startup and system partition, the other drives are currently not used.

You add the two drives not currently used to a storage pool by using Storage Spaces. You need to create two virtual drives and provide availability of data in case up to two physical disks fail. What should you do?

Question: You have a server named Server2 that runs Windows Server 2012 R2. Server2 has several hard drives directly attached to itself. Some of the hard drives are serial-attached SCSI drives, some are SATA drives, and some are solid-state drives. You will add all the attached drives to a storage pool, and you will use the pool to create Storage Spaces used for file sharing. You need to ensure that data that is frequently accessed is stored in physical hard drives that have a better read/write performance. What should you do?

Module 8

Planning and Implementing a High Availability Infrastructure Using Failover Clustering

Contents:

Module Overview	8-1
Lesson 1: Planning an Infrastructure for Failover Clustering	8-2
Lesson 2: Implementing Failover Clustering	8-16
Lesson 3: Planning and Implementing Updates for Failover Clustering	8-22
Lesson 4: Integrating Failover Clustering with Server Virtualization	8-24
Lesson 5: Planning a Multisite Failover Cluster	8-32
Lab: Planning and Implementing a Highly Available Infrastructure by Using Failover Clustering	8-37
Module Review and Takeaways	8-45

Module Overview

Planning for high availability is an essential part of your system design. For most organizations, implementing a disaster recovery plan is too time-consuming or costly. Therefore, these organizations use high availability to avoid having to perform disaster recovery.

Windows Server® 2012 includes failover clustering to provide high availability for applications and services. In Windows Server 2012 R2, failover clustering has improved greatly. In this module, you will learn how to plan and implement failover clustering.

Objectives

After completing this module, you will be able to:

- Plan an infrastructure for failover clustering.
- Implement failover clustering.
- Plan and implement updates for failover clustering.
- Integrate failover clustering with server virtualization.
- Plan a multisite failover clustering.

Lesson 1

Planning an Infrastructure for Failover Clustering

Failover clustering is a highly available feature in Windows Server 2012. When you configure it properly, failover clustering allows an application or service that runs on one cluster node to fail over to another cluster node. However, an improperly configured failover cluster may lead to data loss and performance issues when a cluster node fails. Therefore, to ensure that your failover cluster is successful, you must understand how to plan and design shared storage, hardware capacity, and the quorum.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe failover clustering.
- Describe improvements to failover clustering in Windows Server 2012 and Windows Server 2012 R2.
- Describe the considerations for server workloads on failover clusters.
- Explain how to determine the type and amount of hardware components to utilize for failover clustering.
- Plan network components.
- Plan storage components.
- Describe how to use Cluster Shared Volumes (CSVs) for failover cluster storage.
- Plan cluster quorum.
- Describe application considerations.

Overview of Failover Clustering

A failover cluster is a group of independent computers that work together to increase the availability of applications and services. If a server cluster, or *node*, fails, then another node begins to provide services. This is *failover*, and it results in little or no disruption of service for users.

A failover clustering solution has several components, including:

- Nodes, which are computers that are members of a failover cluster. These computers run cluster service, resources, and applications associated with the cluster.
- A network, across which cluster nodes can communicate with each other and clients. There are three types of networks that you can use in a cluster: public, private, and public-and-private.
- A *resource*, which is an entity that a node hosts. The cluster service manages it, and you can start, stop, or move it to another node.
- Cluster storage, which is a storage system that cluster nodes typically share. Some scenarios, such as clusters of servers that are running Microsoft® Exchange Server, do not require shared storage.

- A failover cluster is a group of independent computers that work together to increase the availability of applications and services
- Components of a failover cluster include:
 - Nodes
 - Network
 - Resource
 - Cluster storage
 - Clients
 - Service or application

- Clients, which are computers (or users) that use the cluster service.
- A service or application, which is a software entity that the cluster presents to clients for their use.

In a failover cluster, each node in the cluster:

- Has full connectivity and communication with the cluster's other nodes.
- Is aware when another node joins or leaves the cluster. Furthermore, each node is aware when a node or resource is failing, and has the ability to take over those services.
- Connects to a network through which client computers can access the cluster.
- Connects through a Fibre Channel, shared serial-attached SCSI bus or Internet Small Computer System Interface (iSCSI) connection to shared storage.
- Is aware of the services or applications that are running locally, and the resources that run on all other cluster nodes.

Usually, cluster storage refers to logical devices, which typically are 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, which the cluster will manage.

Windows Server 2012 includes most of the failover clustering features and administration techniques from Windows Server 2008 R2, in addition to some new features and technologies that increase scalability and cluster storage availability. These features also provide easier management and faster failover.

Failover Clustering Improvements in Windows Server 2012 and Windows Server 2012 R2

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. Server Manager in Windows Server 2012 can discover and manage all clusters that you create in an Active Directory® Domain Services (AD DS) domain.

If you deploy the cluster in a multisite scenario, the administrator can control which nodes in a cluster have votes for establishing quorum. In addition, Windows Server 2012 improves upon the scalability for failover clustering for virtual machines that run on clusters.

- Improved CSVs. Windows Server 2008 R2 introduced this technology for providing virtual machine storage. In Windows Server 2012, CSVs appear as CSV file systems. They support Server Message Block (SMB) version 3.0 storage for the Hyper-V® role in Windows Server 2012 and other applications. Additionally, CSV can use the SMB Multichannel and SMB Direct features that enable traffic to stream across multiple networks in a cluster. You also can implement file servers on CSVs, in scale-out mode. For additional security, you can use 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 CSVs with zero offline time.

Failover clustering improvements in Windows Server 2012	Removed and deprecated failover clustering features in Windows Server 2012
<ul style="list-style-type: none"> • Increased scalability • Improved CSVs • Cluster-aware updating • Active Directory integration improvements • Management improvements 	<ul style="list-style-type: none"> • Cluster.exe command-line tool • Cluster Automation Server (MSCluster) COM interface • Printer cluster

- Cluster-Aware Updating (CAU). In older versions of Windows Server, updating cluster nodes to minimize or avoid downtime required significant preparation and planning. Typically, updating cluster nodes was a manual procedure, which required additional administrative attention. Windows Server 2012 introduces CAU, a new technology that updates cluster nodes automatically with the Windows Update hotfix, while keeping the cluster online and minimizing downtime. You will learn more about this technology later in this module.
- AD DS integration improvements. Windows Server 2008 integrated failover clustering in AD DS. Windows Server 2012 improves on this integration. Now administrators can create cluster computer objects in targeted organizational units (OUs) or, by default, in the same OUs as the cluster nodes. This aligns failover cluster dependencies on AD DS with the delegated Domain Admin model that many Information Technology (IT) organizations use. Additionally, you can deploy failover clusters with access only to read-only domain controllers (RODCs).
- Management improvements. Although failover clustering in Windows Server 2012 still uses almost the same management console and the same administrative techniques, Windows Server 2012 has several important management improvements. These include the improved validation speed in the Validation Wizard and for large failover clusters, in addition to new tests for CSVs, the Hyper-V role, and virtual machines. Furthermore, new Windows PowerShell® cmdlets are available for managing clusters, storage, and LUNs; for monitoring clustered virtual machine applications; and for creating highly available iSCSI targets.

Windows Server 2012 does not include some of the features from older failover clustering versions, while other features have been or deprecated. If you are upgrading from an older version, you should be aware of these changes, which include:

- The Cluster.exe command-line tool is deprecated. However, you have the option to install it with the failover clustering tools, such as Windows PowerShell cmdlets, which provide similar functionality to **cluster.exe** commands.
- The Cluster Automation Server (MSClus) Component Object Model (COM) interface has been deprecated, but you have the option to install it with the failover clustering tools.
- Support for 32-bit cluster-resource dynamic-link libraries (DLLs) is deprecated, but you have the option to install 32-bit DLLs. You should update cluster resource DLLs to 64-bit.
- The High Availability Wizard no longer includes the Print Server role, and you cannot configure it in the Failover Cluster Manager.
- The **Add-ClusterPrintServerRole** cmdlet is deprecated and Windows Server 2012 no longer supports it.

Failover Clustering Improvements in Windows Server 2012 R2

The quorum model in Windows Server 2012 R2 has changed significantly. You now have more options and greater flexibility in maintaining quorum and cluster. In addition, the Failover Cluster Manager console in Windows Server 2012 R2 has a cluster dashboard where you can see the health status of all managed failover clusters. In the console, next to each failover cluster that you manage, there are icons that indicate whether the cluster is running, the number and status of clustered roles, the node status, and the event status. The most important new features in Failover Clustering Quorum in Windows Server 2012 R2 are the following:

- Dynamic quorum. This feature enables a cluster to recalculate quorum in the event of node failure and still maintain working clustered roles, even when the number of voting nodes remaining in the cluster is less than 50 percent.
- Dynamic witness. This feature decides if the witness has a vote to maintain quorum in the cluster.

- Force quorum resiliency. This feature provides additional support and flexibility to manage *split brain syndrome* cluster scenarios. These happen when a cluster breaks into subsets of cluster nodes that are not aware of each other.
- Tie breaker for 50% node split. By using this feature, the cluster can adjust the running node's vote status automatically to keep the total number of votes in the cluster at an odd number.

You will learn more about these new quorum options and modes of work later in this lesson. Besides updating quorum, Microsoft has made other valuable changes to failover clustering. The most important changes in Windows Server 2012 R2 failover clustering are described below.

Global Update Manager Mode

Global Update Manager is responsible for updating the cluster database. In Windows Server 2012, it was not possible to configure these updates, but Windows Server 2012 R2 enables you to configure the mode of work for Global Update Manager. Each time a cluster's state changes, such as when a cluster resource is offline, all nodes in the cluster must receive notification about the event before the Global Update Manager commits the change to the cluster database.

In Windows Server 2012, Global Update Manager works in Majority (read and write) mode. In this mode, when a change happens to a cluster, a majority of the cluster nodes must receive and process the update before it is committed to the database. When the cluster node wants to read the database, the cluster compares the latest time stamp from a majority of the running nodes and uses the data with the latest time stamp.

In Windows Server 2012 R2, Global Update Manager can also work in the All (write) and Local (read) mode. When working in this mode, all nodes in the cluster must receive and process an update before it is committed to the database. However, when it receives the database read request, the cluster will read the data from the database copy that is stored locally. Because all roles receive and process the update, the local cluster database copy is a relevant source of information.

Windows Server 2012 R2 also supports a third mode for Global Update Manager. This mode is Majority (write) and Local (read). In this mode, a majority of the cluster nodes must receive and process an update before it is committed to the database. When it receives the database read request, the cluster reads the data from the database copy that is stored locally.

In Windows Server 2012 R2, the default setting for Hyper-V failover clusters is Majority (read and write). All other workloads in the clusters use All (write) and Local (read) mode. By default, no workloads use Majority (write) and Local (read).

Changing the working mode for Global Update Manager improves cluster database performance and increases the performance of cluster workloads because a cluster database no longer has to perform at the speed of the slowest node.

Cluster Node Health Detection

In Windows Server 2012, the mechanism for node health detection within a cluster declares a node as down if it does not respond to heartbeats for more than five seconds. In Windows Server 2012 R2, specifically for Hyper-V failover clusters, the default threshold value increases from five seconds to 10 seconds if nodes are in the same subnet and to 20 seconds if nodes are in different subnets. This provides increased resiliency for temporary network failures for virtual machines that are running on a Hyper-V cluster, and delays cluster recovery actions in cases of short network interruptions.

AD DS–Detached Cluster

Failover clusters are integrated with AD DS, and you cannot deploy a cluster if nodes are not members of same domain. When a cluster is created, appropriate computer objects for cluster name and clustered role name are created in AD DS.

In Windows Server 2012 R2, you can deploy an AD DS–detached cluster. This cluster does not have dependencies in AD DS for network names. When you deploy clusters in detached mode, the network name and the network names for clustered roles register in local domain name system (DNS), but corresponding computer objects for cluster and clustered roles are not created in AD DS.

Cluster nodes still have to join to the same AD DS domain, but the creator of a cluster does not need to have permission to create new objects in AD DS. In addition, later management of these computer objects is not necessary.

There are also side effects to deploying AD DS–detached clusters. Since computer objects are not created, you cannot use Kerberos authentication when accessing cluster resources. Although Kerberos authentication is used between cluster nodes because their computer accounts and objects are created outside the cluster, AD DS–detached clusters use Windows NTLM authentication. Because of this, we do not recommend that you deploy AD DS–detached clusters for any scenario that requires Kerberos authentication.

To create an AD DS–detached cluster, you must run Windows Server 2012 R2 on all cluster nodes. You cannot configure these features by using the Failover Cluster Manager, so you must use Windows PowerShell.

Considerations for Server Workloads on Failover Clusters

Before you implement failover clustering, you must identify services and applications that you want to make highly available. In other words, you have to identify workloads and their main characteristics before you decide whether to implement a highly available solution. You cannot apply failover clustering to all applications, and sometimes applications have their own redundancy mechanisms that are incompatible with highly available technology that you want to implement.

Additionally, you should be aware that failover clustering does not provide improved scalability by adding nodes. You can obtain scalability by scaling up and by using more-powerful hardware for individual nodes. Therefore, you should use failover clustering when your goal is high availability, not scalability. In Windows Server 2012 and Windows Server 2012 R2, there is one exception to this: if you implement Scale-Out File Services on CSVs, you can achieve a level of scalability while maintaining high availability. In this type of cluster, adding new nodes improves the performance of the whole cluster.

You must analyze workloads that are running on your servers before making them highly available. This becomes even more important if these workloads are running inside virtual machines. If so, you must determine whether you will implement high availability on a virtual machine level or on the service level. Sometimes, you will implement failover clustering inside virtual machines to make some workloads highly available, but you can also choose to implement technologies such as Network Load Balancing (NLB), for workloads like web-based services. You will learn more details about high availability for virtualized workloads in a later lesson.

When designing high availability for workloads that run on your servers, you should:

- Identify workloads and their main characteristics
- Understand that you cannot apply failover clustering to all applications and workloads
- Remember that failover clustering does not provide improved scalability
- Consider high availability approaches for workloads in virtual machines

When analyzing and planning server workloads for high availability, you should consider the following:

- For each specific workload, determine whether it requires scalability, redundancy, or both.
- For each specific workload, determine what hardware resources you require to achieve redundancy.
- When planning for redundancy, always plan for growth. We recommend that you use tools, such as workload and capacity planners, for specific services, including Microsoft SQL Server®, Microsoft SharePoint®, and others.
- Evenly distribute the highly available applications from a failed node. You should spread the highly available services or applications from a failed node across the remaining nodes. This prevents the overloading of a single node.

Selecting Hardware Components for Cluster Nodes

Failover clusters have to satisfy several criteria to meet availability and support requirements. When carefully selecting your cluster-node hardware, consider the following:

- Select hardware for a failover cluster that meets the Certified for Windows Server 2012 logo requirements. Hardware that has this logo has been independently tested to meet the highest technical standards for reliability, availability, stability, security, and platform compatibility. Furthermore, you can access official support options if you experience malfunctions or other issues.
- Install the same or similar hardware on each failover cluster node. For example, if you choose a specific model of network adapter for one node, you should install this adapter on all nodes.
- Ensure that if you use serial-attached SCSI or Fibre Channel storage connections, you use identical mass storage device controllers for cluster storage in all clustered servers. The mass storage device controllers should use the same firmware version.
- Ensure that if you use iSCSI storage connections, each clustered server has one or more network adapters or host bus adapters that are dedicated to the cluster storage. Additionally, whenever possible, do not use the network that you use for iSCSI storage connections for network communication. In all clustered servers, the network adapters that you use to connect to the iSCSI storage target should be identical, and we recommend that you use Gigabit Ethernet or more.
- Ensure that after you configure the servers with the hardware, each clustered server passes all tests in the Validate a Configuration Wizard. Otherwise, Microsoft will not support the cluster's configuration.
- Plan for hardware components by ensuring that nodes have enough capacity to host workloads from other nodes that might fail.

The hardware requirements for a failover implementation include:

- Server hardware components must have the Certified for Windows Server 2012 logo
- Server nodes should all have the same configuration and contain the same or similar components
- Each clustered server passes all tests in the Validate a Configuration Wizard

You should examine all cluster configuration components to identify single points of failure. You can remedy many single points of failure with simple solutions, such as adding storage controllers to separate and stripe disks, teaming network adapters, and using multipath software. These solutions reduce the probability that a single device's failure will cause a failure in the cluster. Typically, server-class computer hardware has power redundancy options for multiple power supplies and for creating a redundant array of independent disks (RAID) for disk data redundancy.

Planning Network Components

Failover cluster network components must have the Certified for Windows Server 2012 logo and pass the tests in the Validate a Configuration Wizard. Additional requirements include the following:

- The network adapters in each node should be identical and have the same IP version, speed, duplex, and flow-control capabilities.
- The networks and network equipment to which you connect the nodes should be redundant. This ensures that the nodes continue communicating even if a single point of failure occurs. You can use Network Adapter Teaming to provide single network redundancy. We recommend that you use multiple networks to provide multiple paths between nodes for internode communication. Nodes in a cluster exchange heartbeats to verify their presence in the cluster. This communication is critical, so we recommend having a separate network for this purpose. If you do not, Windows® will generate a warning during the validation process.
- The network adapters in a cluster network must have the same method for IP address assignment, which means that they all use either static IP addresses or Dynamic Host Configuration Protocol (DHCP).

The network requirements for a failover implementation include:

- The network hardware components must have 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 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 wizard's report will include a warning that the network should not have single points of failure. In addition, Microsoft Product Support will not support this configuration.

Planning Storage Components

Most scenarios for failover clustering require shared storage, which provides consistent data to a highly available service or application after failover. A failover cluster can utilize one of three shared storage options:

- Shared serial attached SCSI is the lowest-cost option. However, shared serial attached SCSI is not very flexible for deployment because the two cluster nodes must be in close physical proximity. Additionally, the shared storage devices that support shared serial attached SCSI have a limited number of connections for cluster nodes.
- iSCSI. iSCSI is a type of storage area network (SAN) that transmits SCSI commands over IP networks. Performance is acceptable for most scenarios when the physical medium for data transmission is between 1 and 10 gigabytes per second (Gbps) Ethernet. This type of SAN is inexpensive to implement, because it requires no specialized networking hardware. In Windows Server 2012, you can implement iSCSI target software on any server and present local storage over iSCSI interface to clients.

Failover clusters require shared storage to provide consistent data to a virtual server after failover

Shared storage options include:

- Serial attached SCSI
- iSCSI
- Fibre Channel
- Shared VHDX



- Fibre Channel. Fibre Channel SANs have better performance than iSCSI SANs, but are more expensive. The implementation of Fibre Channel SANs requires specialized knowledge and hardware.
- VHDX files as shared storage for guest clustering. Guest clusters created in Windows Server 2012 R2 Hyper-V now can use a shared virtual .vhdx drive instead of a SAN or iSCSI-based storage. A shared .vhdx drive is added through the SCSI interface for virtual machines, and it must be stored either on a Scale-Out File Server or on a CSV. You will learn more about shared virtual hard disks later in this course.



Note: The Microsoft iSCSI Software Target is an integrated feature in Windows Server 2012, and can provide storage from a server over a TCP/IP network, including shared storage for applications that a failover cluster hosts. Additionally, 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 a storage solution, you should be aware of the following storage requirements:

- You can use the native disk support that failover clustering includes by using basic disks and not dynamic disks.
- You should format the partitions with New Technology File System (NTFS). For the disk witness, the partition must be NTFS file system, because Windows Server 2012 does not support the File Allocation Table (FAT) file system.
- You can use either master boot record (MBR) or globally unique identifier (GUID) partition table (GPT) for your disk-partition style.
- Your storage must follow the SCSI Primary Commands-3 standard, because improvements in failover clustering require that the storage respond correctly to specific SCSI commands. Specifically, the storage must support Persistent Reservations.
- You must select a miniport driver for storage that works with the Storport storage driver. Storport offers a higher performance architecture and good Fibre 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 that you use on one cluster from being seen on another cluster. Consider using Multipath I/O (MPIO) software. Cluster nodes commonly use multiple host-bus adapters to access storage, which provides additional high availability. However, to use multiple host-bus adapters, you must use multipath software. For Windows Server 2012, you must base your multipath solution on MPIO. Usually, your hardware vendor 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.

Using CSVs for Failover Cluster Storage

When planning for cluster storage, you should be familiar with the concept of CSVs. In a typical deployment of a failover cluster, only a single node controls a LUN on the shared storage. This means that another node cannot see shared storage, until it becomes an active node. Windows Server 2008 R2 introduced the CSV technology, which enables multiple nodes to share a single LUN concurrently. Each node obtains exclusive access to individual files on the LUN rather than 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 is a technology that enables multiple nodes to share a single LUN or volume concurrently
- Windows Server 2008 introduced CSV
- Windows Server 2012 improved CSV
- Windows Server 2012 R2 introduced new features:
 - Optimized CSV placement policies
 - Increased CSV resiliency
 - CSV cache allocation
 - CSV diagnosis
 - CSV interoperability

In the first version of Windows Server 2008 R2, CSV hosted only virtual machines that were run on a Hyper-V server in a failover cluster. This enabled administrators to use a single LUN to host multiple virtual machines in a failover cluster. Multiple cluster nodes could access the LUN, but each virtual machine could run on only one node at a time. If the node on which a virtual machine was running failed, CSV enabled the virtual machine to restart on a different node in the failover cluster. Additionally, this provided simplified disk management for hosting virtual machines, because each virtual machine no longer required a separate LUN.

In Windows Server 2012, CSV has additional improvements. You now can use CSV for other roles and not just for Hyper-V. For example, you can configure the file server role in failover clustering in the Scale-Out File Server scenario. A Scale-Out File Server provides scale-out file shares that are always available for file-based server-application storage. Scale-out file shares enable you 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. Furthermore, a CSV file system can perform Chkdsk.exe without affecting applications that have open handles on the file system.

Windows Server 2012 R2 has further improved CSVs. The following sections describe the important improvements in CSV in Windows Server 2012 R2.

Optimized CSV Placement Policies

In a failover cluster for Windows Server 2012, one node in the cluster is designated as a coordinator for a CSV, and there is no automatic rebalance of this designation. Note that the coordinator for CSV owns the physical disk resource that is associated with a LUN. Also, all I/O operations specific to the file system are performed through the coordinator node. In Windows Server 2012 R2, CSV ownership distributes evenly between cluster nodes based on the number of CSVs that each node owns already. The Failover Cluster service performs a rebalance automatically when a node rejoins a cluster, when you add a new cluster, or when you restart a cluster node.

Increased CSV Resiliency

CSV in Windows Server 2012 uses SMB as a transport for I/O forwarding between nodes in a cluster. SMB uses a Server service on cluster nodes, and if this service becomes unavailable, it can decrease performance or the accessibility of storage. Windows Server 2012 R2 implements multiple instances of Server service, which improves the resilience and scalability of internode SMB traffic. The default instance of Server service now accepts clients that access regular file shares, and a second CSV instance handles only internode CSV traffic. In addition, if Server service becomes unhealthy on one cluster node, CSV ownership can be transitioned to another node automatically to ensure greater resiliency.

CSV Cache Allocation

CSV cache enables the server to improve performance by using random access memory (RAM) as a cache for write-through operations. In Windows Server 2012, CSV cache is disabled by default, but when it is enabled, you can allocate up to 20 percent of the total RAM for cache. In Windows Server 2012 R2, you can allocate up to 80 percent of memory for CSV cache, which improves the performance of the clustered server role. This is especially useful for scale-out file server clusters. In deployments where a Hyper-V cluster runs on a scale-out file server cluster, we recommend that you enable and use the CSV cache, but with greater allocation for a scale-out file server deployment to achieve maximum performance of virtual machines stored on file servers.



Note: In Windows Server 2012 R2, the name of the private property of the cluster physical disk resource has changed from **CsvEnableBlockCache** to **EnableBlockCache**.

CSV Diagnosis

In Windows Server 2012 R2, you can see the state of CSV on a per node basis. For example, you can see whether I/O is direct or redirected or whether the CSV is unavailable. If a CSV is in I/O redirected mode, you also can view the reason. You can retrieve this information by using the Windows PowerShell cmdlet **Get-ClusterSharedVolumeState** with the parameters **StateInfo**, **FileSystemRedirectedIOReason**, or **BlockRedirectedIOReason**. This provides you with a better view of how CSV works across cluster nodes.

CSV Interoperability

CSVs in Windows Server 2012 R2 also support interoperability with the following technologies:

- Resilient File System (ReFS)
- Data Deduplication
- Parity Storage Spaces
- Tiered Storage Spaces
- Storage Spaces write-back caching

This added support expands the scenarios in which you can take advantage of these CSV features.

Planning Cluster Quorum

Quorum is the number of elements that must be online for a cluster to continue to run. Each element can cast one vote to determine whether the cluster continues to run, and each cluster node is an element. If there is an even number of nodes, then Windows assigns an additional element, or *witness*, to the cluster. The witness element can be either a disk or a file share. Each voting element contains a copy of the cluster configuration, and the cluster service works to keep all copies synchronized at all times.

The cluster will stop providing failover protection if most of the nodes fail or if there is a problem with communication between the cluster nodes. Without a quorum mechanism, each set of nodes could continue to operate as a failover cluster, which could cause a partition within the cluster.

In failover clusters, a quorum is a consensus that enough cluster members are available to provide services

- Quorum:
 - Is based on votes in Windows Server 2012
 - 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
- Quorum modes:
 - Node Majority
 - Node and Disk Majority
 - Node and File Share Majority
 - No Majority: Disk Only

MCT USE ONLY. STUDENT USE PROHIBITED

Quorum prevents two or more nodes from operating a failover cluster resource concurrently. If voting does not produce a clear majority among 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 potentially damaging situation might occur in which more than one node tries to control access to a resource. For example, if two or more instances of the same database became available on the network, or if data was accessed and written to a target from more than one source at a time, the application itself could be damaged and, the data could be corrupted.

Because a given cluster has a specific set of nodes and a specific quorum configuration, the cluster can calculate the number of votes that are required for the cluster to continue providing failover protection. If the number of votes drops below the majority, the cluster stops running and it will not provide failover protection if there is a node failure. Nodes will still listen for the presence of other nodes, in case another node appears again on the network, but the nodes will not function as a cluster until they achieve a majority consensus or quorum.



Note: The full functioning of a cluster depends not just on quorum, but also on the capacity of each node to support the services and applications that failover to that node. For example, a cluster that has five nodes could still have quorum after two nodes fail, but each remaining cluster node will continue serving clients only if it has enough capacity, such as processing power, network bandwidth, or RAM, to support the services and applications that failed over to it. An important part of the design process is 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 failover to it.

There are four quorum modes that Windows Server 2012 supports:

- **Node Majority.** Each node that is available and in communication can vote. The cluster functions only with a majority, or more than half of the votes. This model is preferred when the cluster consists of an odd number of server nodes (maintaining or achieving quorum does not require a witness).
- **Node and Disk Majority.** Each node plus a designated disk in the cluster storage can vote. The disk witness can vote when it is available and in communication with the cluster. The cluster functions only with a majority of the votes. Clusters with an even number of server nodes that are able to communicate with one another use this model.
- **Node and File Share Majority.** Each node, plus a designated file share witness that the administrator creates, can vote when they are available and in communication. The cluster functions only with a majority of the votes.
- **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 is 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. The quorum-shared disk is a single point of failure, so we do not recommend this mode.

When you configure a failover cluster in Windows Server 2012, the Installation Wizard selects one of two default configurations automatically. By default, failover clustering selects:

- Node Majority, if there is an odd number of nodes in the cluster.
- Node and Disk Majority, if there is an even number of nodes in the cluster.

Modify this setting only if you determine that a change is appropriate for your cluster and you understand the implications of making the change. In addition to planning your quorum mode, you should also consider the capacity of your cluster's nodes, 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 deploy several applications or services on the cluster, each remaining cluster node may not have the capacity to provide services.

Quorum Changes in Windows Server 2012 R2

In Windows Server 2012 R2, old quorum modes such as Node Majority, Node and Disk Majority, and Node and File Share Witness Majority, are no longer used. Instead, Windows Server 2012 R2 introduces the concept of *Dynamic Quorum*. This feature provides the ability for a cluster to recalculate quorum in the event of node failure and still maintain working clustered roles, even when the number of voting nodes remaining in the cluster is less than 50 percent.

Windows Server 2012 R2 enhances this feature by introducing the concept of *Dynamic Witness*. When you configure a cluster in Windows Server 2012 R2, dynamic quorum is selected by default, but the witness vote is also adjusted dynamically based on the number of voting nodes in the current cluster membership. For example, if a cluster has an odd number of votes, a quorum witness does not have a vote in the cluster. If the number of nodes is even, a quorum witness has a vote. If a witness resource has for some reason failed or is offline, the cluster will automatically set the witness vote to a value of 0. This approach greatly reduces the risk of a malfunctioned cluster because of a failing witness. If you want to see if a witness has a vote, you can use Windows PowerShell and a new cluster property in the following cmdlet:

```
(Get-Cluster).WitnessDynamicWeight
```

A value of 0 indicates that the witness does not have a vote. A value of 1 indicates that the witness has a vote. The cluster can now decide whether to use the witness vote based on the number of voting nodes that are available in the cluster. A much simpler quorum configuration when you create a cluster is an additional benefit. Windows Server 2012 R2 will configure quorum witness automatically when you create a cluster. In addition, when you add or evict cluster nodes, you no longer have to adjust the quorum configuration manually. The cluster now automatically determines quorum management options and quorum witness.

Force Quorum Resiliency

This feature provides additional support and flexibility to *split brain syndrome* cluster scenarios. This scenario happens when a cluster breaks into subsets of cluster nodes that are not aware of each other. The cluster node subset that has a majority of votes will run while others are turned off. This scenario usually happens in multisite cluster deployments. If you want to start cluster nodes that do not have a majority, you can force quorum to start manually by using the **/fq** switch.

In Windows Server 2012 R2, in such scenarios, the cluster will detect partitions in the cluster automatically as soon as connectivity between nodes is restored. The partition that was started by forcing a quorum is considered authoritative, and other nodes rejoin the cluster. When this happens, the cluster returns to a single view of membership. In Windows Server 2012, partitioned nodes without quorum were not started automatically, and administrator had to start them manually with the **/pq** switch. In Windows Server 2012 R2, both sides of the split cluster have a view of cluster membership, and they will reconcile automatically when connectivity is restored.

Tie Breaker for 50% Node Split

Windows Server 2012 R2 enhances dynamic quorum with additional functionality. The cluster can now adjust the running node's vote status automatically to keep the total number of votes in the cluster at an odd number. This is called *Tie breaker for 50% node split*, and it works with dynamic witness functionality. You can use dynamic witness functionality to adjust the value of a quorum witness vote. For example, if you have a cluster with an even number of nodes and a file share witness, if the file share witness fails, the cluster uses dynamic witness functionality to remove the vote from file share witness automatically. However, because the cluster now has an even number of votes, the cluster tie breaker selects a node randomly and removes it from the quorum vote to maintain an odd number of votes. If the nodes are distributed evenly in two sites, this helps to maintain cluster functionality in one site. In previous Windows Server versions, if both sites have an equal number of nodes and a file share witness fails, both sites stop the cluster.

If you want to avoid the node being selected randomly, you can use the **LowerQuorumPriorityNodeID** property to predetermine which node has its vote removed. You can set this property by using the following Windows PowerShell command, where *1* is the example node ID for a node in the site that you consider less critical:

```
(Get-Cluster).LowerQuorumPriorityNodeID = 1
```

Application Considerations

Failover clustering cannot protect all applications. Therefore, it is very important to analyze applications and services that you need to make highly available before you implement any highly available technology. You can make some applications highly available simply by deploying more than one server that is hosting that application, without implementing any specific highly available technology. On the other hand, some applications cannot benefit from a failover cluster, as they do not use any shared data.

When you plan a failover clustering implementation for applications, you should consider:

- Web-based services cannot be made highly available by using failover clustering
- AD DS does not support clustering
- Exchange Server uses failover clustering but not directly
- DHCP servers support failover clustering
- SQL Server is a good candidate for failover clustering
- In a Hyper-V environment you can use host-based or guest-based clustering

When planning for failover clustering implementation for applications, you should consider following:

- If you want to make web-based services highly available, you should consider using NLB instead of failover clustering.
- If you want to make AD DS highly available, you should not put domain controllers in the cluster. You can simply deploy more than one AD DS domain controller. A similar scenario applies to DNS servers.
- If you want to implement high availability Exchange Server, you should not use Windows tools to configure high availability. Exchange Server 2010 and newer versions have their own mechanisms for high availability that rely on failover clustering and NLB.

- If you want to make DHCP servers highly available, you can implement a DHCP cluster, or you can implement DHCP failover as an alternative. This technology is specific to Windows Server 2012.
- SQL Server is typically a good candidate for failover clustering.

When planning application high availability, we recommend that you analyze all possibilities and alternatives before implementing any specific high availability technology. In some cases, applications have their own technology for high availability, which is preferable.

Failover clustering is best suited for stateful applications that are restricted to a single set of data. A database is an example of such an application. Data is stored in a single location, and only one database instance can use it. The best results for failover clustering occur when the client can reconnect to the application automatically after failover. If the client cannot reconnect automatically, the user must restart the client application.

In addition, you can use failover clustering to provide highly available virtual machines. You can implement failover clustering in a virtual environment in a host clustering or a guest clustering scenario. 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. You implement failover protection at the host server level. The applications or services that run in the virtual machine do not have to be compatible with failover clustering, and they do not have to be aware that the virtual machine is clustered.

You configure guest failover clustering in a similar way to a physical server failover clustering, except that the cluster nodes are virtual machines. In this scenario, you create two or more virtual machines, and enable failover clustering within the guest operating system. Then the application or service is enabled for high availability between the virtual machines.

Lesson 2

Implementing Failover Clustering

We recommend specific hardware and software configurations for Windows Server 2012 failover clusters. Failover clusters provide a higher level of service than stand-alone servers. Therefore, cluster hardware requirements are often stricter than the requirements for stand-alone servers.

In this lesson, you will learn how to prepare for cluster implementation, and discuss the hardware, network, storage, infrastructure, and software requirements for Windows Server 2012 failover clusters. In addition, this lesson outlines the steps for using the Validate a Configuration Wizard to ensure correct cluster configuration, in addition to the process for migrating failover clusters.

Lesson Objectives

After completing this lesson, you will be able to:

- Deploy a failover cluster.
- Deploy server roles on a failover cluster.
- Configure failover cluster settings.
- Configure application and resource settings.
- Describe how to use a Scale-Out File Server in Windows Server 2012.
- Configure a scale-out file server cluster.

Deploying and Validating a Failover Cluster

To use failover clustering, you first must install a failover clustering feature on all nodes that you will use in a cluster. You can install the failover clustering feature easily on the full installation of Windows Server 2012 by using the Server Manager Add Roles & Features Wizard, or on a Server Core by using the Windows PowerShell interface.

A Server Core installation installs only essential components, so a failover cluster on a Server Core installation offers several key benefits:

- A reduced surface of vulnerability that is open to attack.
- A reduced number of required updates that you must install.
- Fewer components to manage and less required disk space for the installation.
- Lower resource overhead, leaving more resources for the clustered instances.

You can use the Cluster Creation Wizard in the Failover Cluster Manager console to create a cluster. However, before actually creating a cluster, we highly recommend that you validate the current software and hardware configuration of cluster nodes. You can validate your configuration by using the Cluster Validation Wizard in the Failover Cluster Manager console.

- You can deploy failover clustering feature on the full GUI or a Server Core version of Windows Server 2012
- You should consider the benefits from deploying a cluster on Server Core
- Before creating a cluster, you should run the Validate a Configuration Wizard to perform:
 - System configuration tests
 - Network tests
 - Storage tests

The Validate a Configuration Wizard runs tests that confirm that hardware and software settings are compatible with failover clustering. You can run the entire set of configuration tests or a subset of the tests. As a best practice, run the tests on servers and storage devices before you configure a failover cluster, and run it again after you make any major changes to the cluster. You can access the test results after the wizard completes, or by accessing the report file in the %Windir%\cluster\Reports directory.

The tests in the wizard include specific simulations of cluster actions, and fall into several categories:

- **System Configuration Tests.** These tests determine if the selected servers meet specific requirements. For example, you can test whether you meet the requirement that your servers run the same operating system version and software updates.
- **Network tests.** These tests determine if the planned cluster networks meet specific requirements. For example, you can test whether you meet the network redundancy requirements.
- **Storage tests.** These tests determine if your storage meets specific requirements. For example, you can test whether your storage supports the necessary SCSI commands and handles simulated cluster actions correctly.

The report from the Validate a Configuration Wizard reports one of the following responses for each tested item:

- Meets the requirements for a failover cluster.
- Might meet the requirements. However, you receive a warning notice that suggests that you review the hardware and software settings to ensure that they meet best practices.
- Does not meet the requirements.

Additionally, the report may indicate that the test could not be run. For example, you may receive this result if you are testing only a single node, as certain tests do not run in that case.

The Validate a Configuration Wizard report contains details about tests that were run and the results. The report also provides information about reasons for any failures. If you cannot resolve problems after reading the report, contact the vendor for configuration guidance or software version compatibility changes. In some cases, a vendor may recommend a specific driver or firmware version that is stable in a clustered environment.

Deploying Server Roles on a Failover Cluster

Failover clustering supports clustering of several Windows Server roles, such as File Services, DHCP, and Hyper-V. To implement clustering for a server role, or for external applications such as SQL Server or Exchange Server, perform the following procedure:

1. Install the failover clustering feature. Use Server Manager, `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 a single Server Manager console.
2. Verify configuration and create a cluster. With the appropriate nodes, use the Failover Cluster Manager snap-in to validate a configuration, and then create a cluster with the selected nodes.

To deploy and configure a failover cluster, you must:

- Install the failover clustering feature
- Verify configuration
- Install the role on all cluster nodes
- Create a clustered application
- Configure the application
- Test failover

3. Install the role to be clustered 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 options on the application that you are using in the cluster.
6. Test failover by using the Failover Cluster Management snap-in to move the service intentionally from one node to another.

After you create the cluster, you can monitor its status by using the Failover Cluster Management console to manage available options.

Configuring Settings for a Failover Cluster

You can configure many properties of a newly created cluster. When you open Cluster Properties, you can configure a cluster name, change the name, add various types of resources to the cluster, including IP address and network name, and configure cluster permissions. By configuring cluster permissions, you determine who has full control over that specific cluster and who can just read the cluster configuration. Additionally, you can perform standard management tasks on each cluster on demand, including:

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. For each node in a cluster, you can stop the 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 you want to dedicate 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 how the cluster achieves quorum and who can vote in a cluster. Remember that quorum is configured and used differently in Windows Server 2012 and Windows Server 2012 R2.
- 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, Windows Server 2012 no longer supports `Cluster.exe`, which you could use for some of these tasks in previous Windows Server operating system versions, and it is not part of the default installation. Windows PowerShell 2012 introduces Windows PowerShell cmdlets for cluster management.

Configuring Failover and Failback 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 also 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.

- When using Preferred Owners, you should consider the following:
 - Preferred owners are set on the clustered application
 - Multiple preferred owners can be set in an ordered list
- By setting Preferred Owners you can control:
 - The order in which an application will select a node on which to run
 - The applications that can be run on the same nodes in an active/active configuration
- You can modify failover and failback settings in these ways:
 - Setting the number of times the cluster service will restart a clustered application in a set period of time
 - Setting or preventing failback of the clustered application to the preferred node once it becomes available

Setting	Result
Example 1: General tab, Preferred Owner: Node1 Failover tab, Failback setting: Allow failback (Immediately)	If the service or application fails over from Node1 to Node2, once Node1 is available again, the service or application will fail back to Node1.
Example 2: Failover tab, Maximum failures in the specified period: 2 Failover tab, Period (hours): 6	In a six-hour period, if the application or service fails no more than two times, it will be restarted or failed over every time. If the application or service fails a third time in the six-hour period, it will remain in the failed state. The default value for the maximum number of failures is $n-1$, where n is the number of nodes. You can change the value, but we recommend a low value so that if multiple node failures occur, the application or service will not move between nodes indefinitely.

It is not mandatory to configure these settings if your cluster nodes are equal and you do not have any preference regarding which node will host clustered services most often. However, if you prefer that one node should be active most of the time, and that another node (or nodes) act as a spare node, you should configure failover and failback settings.

MCT USE ONLY. STUDENT USE PROHIBITED

Using Scale-Out File Server in Windows Server 2012

Unlike other roles from previous Windows Server versions, there are some significant changes in the File Server role cluster in Windows Server 2012.

In Windows Server 2008 R2 and older versions, you could make folder shares highly available by deploying the File Server role in the cluster. All highly available shares were accessible on one cluster node, and if that node failed, another node took ownership and began hosting the folder share. Because only one node was serving clients at a time, it was not recommended to store files like virtual machines or databases on file servers in the cluster. However, in Windows Server 2012 and Windows Server 2012 R2 that has changed.

The File Server role cluster can work in two modes:

- Scale-out file server cluster
- File server cluster for general use

The key benefits of scale-out file server cluster:

- Active-active file shares
- Increased bandwidth
- CSV cache
- Simpler management

In Windows Server 2012, you can deploy the clustered File Server role in two modes:

- **Scale-Out File Server for application data (Scale-Out File Server).** Windows Server 2012 introduces the clustered file server mode, which provides the ability to store server application data, such as Hyper-V virtual machine files, on file shares. At the same time, it provides a level of reliability, availability, manageability, and high performance that you would expect from a SAN. These benefits are due to the fact that file shares are online on all nodes simultaneously, which provides better scalability than failover cluster scenarios. Sometimes this is known as an active/active cluster. You should not implement a Scale-Out File Server for general file or folder shares, but rather for SQL Server databases or virtual machine storage.
- **File server for general use.** This is the same file-server clustering that was in the previous Windows Server versions. All shares associated with the clustered file server are online on one node at a time, as an *active-passive cluster*. File shares associated with this type of clustered file server are *clustered file shares*.

The key benefits of using a scale-out file server cluster are:

- **Active/active clustering.** While other failover clusters work in an active-passive mode, in a scale-out file server cluster, all nodes can accept and serve SMB client requests. In Windows Server 2012 R2, SMB 3.0 is upgraded to SMB 3.0.2. This version improves scalability and manageability for Scale-Out File Servers. SMB client connections, in Windows Server 2012 R2, are tracked per file share (instead of per server), and clients are redirected to the cluster node with the best access to volume used by the file share.
- **Increased bandwidth.** In previous version of Windows Server, bandwidth of the file server cluster was constrained to the bandwidth of a single cluster node. Because of the active/active mode in the scale-out file server cluster, you have much higher bandwidth that you can further increase by adding cluster nodes.
- **CSV Cache.** Because the scale-out file server clusters use CSVs, they also benefit from the use of CSV Cache. CSV Cache is a feature that you can use to allocate RAM as a write-through cache. The CSV Cache provides caching of read-only unbuffered I/O. This can improve performance for applications such as Hyper-V, which conducts unbuffered I/O when accessing a virtual hard disk (VHD) file. You can allocate up to 20 percent with Windows Server 2012 and 80 percent with Windows Server 2012 R2 of the total physical RAM for CSV write-through cache, which will be consumed from nonpaged pool memory.
- **Simpler management.** When using a scale-out file server cluster, you can add CSV storage and shares at any time after the cluster is created.

Demonstration: Configuring the Scale-Out File Server Cluster

In this demonstration, you will see how to configure a scale-out file server cluster.

Demonstration Steps

1. On LON-SVR1, verify that the **File Server Role Service** is installed, and then install the **Failover Clustering** feature by using Server Manager.
2. On LON-SVR2, verify that the **File Server Role Service** is installed, and then install the **Failover Clustering** feature by using Server Manager.
3. On LON-SVR1, start **iSCSI Initiator**.
4. Use the **172.16.0.10** address to discover and connect to the **iSCSI Target**.
5. On LON-SVR2, start **iSCSI Initiator**.
6. Use the **172.16.0.10** address to discover and connect to the **iSCSI Target**.
7. On LON-SVR2, open **Disk Management**, and then initialize and bring online **Disk 2** and **Disk 3**. Format these drives.
8. On LON-SVR1, open **Disk Management**, and then bring online **Disk 3** and **Disk 4**.
9. Open the Failover Cluster Manager console on LON-SVR1.
10. Create a cluster with nodes LON-SVR1 and LON-SVR2. Do not run validation at this point.
11. Provide **FSCluster** as the value for **Access Point for Administering the Cluster**.
12. Provide **172.16.0.127** as the IP address for the cluster.
13. Add **Cluster Disk 1** and **Cluster Disk 2** to the **Cluster Storage**, and then assign **Cluster Disk 1** to a CSV.
14. Configure **Quorum Settings** as **Typical (recommended)**, and then start the Configure Cluster Role Wizard to configure a clustered role.
15. Add a new cluster role, select **File Server**, and then select **Configure Scale-Out File Server for Application Data**.
16. Provide **AdatumFS** as the **Client Access Point**, and then add a file share to the existing ADatumFS cluster role.
17. In the New Share Wizard, select **SMB Share-Applications**, and then click **select by volume**.
18. Name the share **TestShare**.

Lesson 3

Planning and Implementing Updates for Failover Clustering

In Windows Server 2012, a new technology enables you to update cluster nodes without downtime, making the process safer and faster. This greatly reduces administrative overhead on cluster nodes updating. In this lesson, you will learn about Cluster-Aware Updating (CAU).

Lesson Objectives

After completing this lesson, you will be able to:

- Describe CAU.
- Describe how CAU works.
- Configure CAU.

What Is CAU?

Special attention is necessary when applying operating system updates to nodes in a cluster. In older versions of Windows Server, such as Windows Server 2008 R2 or older, if you wanted to provide zero downtime for a clustered role, you had to update cluster nodes manually, one after another, and you had to move resources manually from the node being updated to another node. This procedure was very time-consuming. In Windows Server 2012, Microsoft has implemented a new feature for automatic updating of cluster nodes called Cluster-Aware Updating (CAU).

CAU is an automated feature specific to Windows Server 2012 that updates nodes in a cluster with minimal or no downtime

Benefits include:

- Automatic cluster updating
- Can be scheduled
- No downtime

CAU is a feature that enables administrators to update cluster nodes automatically with minimal or no downtime 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 updates 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 a file server with SMB Transparent Failover, CAU can orchestrate cluster updates with no effect on the service availability.

Question: How do you update cluster nodes in your environment?

How Does CAU Work?

CAU is based on orchestrating a process of cluster node updating. CAU can orchestrate the complete cluster updating operation in two modes:

- Remote-updating mode. In this mode, a computer that runs Windows Server 2012 or Windows 8 is configured as an orchestrator. To configure a computer as a CAU orchestrator, you must install failover clustering administrative tools. The orchestrator computer is not a member of the cluster that updates during the procedure. From the orchestrator computer, the administrator triggers on-demand updating by using a default or a 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 performs updates on each cluster node in sequence. 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.

CAU can work in two modes:

- Remote-updating mode
 - Separate computer is configured as an orchestrator
 - Failover Clustering Administrative Tools must be installed
 - Orchestrator computer is not a cluster member
- Self-updating mode
 - CAU clustered role is configured as a workload
 - There is no dedicated orchestrator computer
 - Cluster updates itself

To use CAU, you must install the failover clustering feature in Windows Server 2012 and create a failover cluster. The components that support CAU functionality install automatically on each cluster node.

You must also install the CAU tools, which are included in the Failover Clustering Tools. The CAU tools consist of the CAU UI and the CAU Windows PowerShell cmdlets. The Failover Clustering Tools install by default on each cluster node when you install the failover clustering feature. You can also install these tools on a local or remote computer that runs Windows Server 2012 or Windows 8 and has network connectivity to the failover cluster.

Demonstration: Configuring CAU

In this demonstration, your instructor will show you how to configure CAU.

Demonstration Steps

1. Make sure that the cluster is configured and running on LON-SVR1 and LON-SVR2.
2. Add the **Failover Clustering Feature** to LON-DC1.
3. Run **Cluster-Aware Updating** on LON-DC1, and configure it to connect to **FSCLUSTER**.
4. Preview updates that are available for nodes LON-SVR1 and LON-SVR2.
5. Review available options for the Updating Run Profile.
6. Apply available updates to **FSCLUSTER** from LON-DC1.
7. After updates are applied, configure **Cluster self-updating options** on LON-SVR1.

Lesson 4

Integrating Failover Clustering with Server Virtualization

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 and several tools to manage highly available virtual machines. In this lesson, you will learn how to integrate failover clustering on the Hyper-V server virtualization platform.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe options for making virtual machine workloads highly available.
- Describe how a failover cluster works with Hyper-V.
- Describe infrastructure requirements for implementing failover clustering for Hyper-V.
- Describe how to implement Storage Migration.
- Describe how to implement Live Migration.
- Describe guidelines for implementing virtual machine availability.

Options for Making Virtual Machine Workloads Highly Available

Most organizations have certain critical applications that 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. To make virtual machines (and services that run within virtual machines) highly available, you can choose between several options, including:

- 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), which you implement inside virtual machines.

High availability options	Description
Host clustering	<ul style="list-style-type: none"> • Makes virtual machines highly available • Does not require virtual machine operating system or application to be cluster-aware
Guest clustering	<ul style="list-style-type: none"> • Makes virtual machines 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"> • Makes virtual machines NLB cluster nodes • Use for web-based applications

Host Clustering

With host clustering, you can 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, which implements failover protection at the host server level. This means that the guest operating system and applications that run within the virtual machine do not have to be cluster-aware. However, the virtual machine remains highly available. Some examples of applications that are not cluster-aware are print servers or proprietary network-based applications, such as an accounting application. Should the host node that controls the virtual machine become unavailable, the secondary host node assumes control and restarts the virtual machine as quickly as possible. You can move the virtual machine from one node in the cluster to another in a controlled manner. For example, you could move the virtual machine to another

node while patching the host operating system. The applications or services that run 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. The failover is at the virtual machine level, so there are no dependencies on software that you install on 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 then enable failover clustering within the guest operating system. You then enable the application or service for high availability between the virtual machines by using failover clustering in each virtual machine. You implement failover clustering within each virtual machine node's guest operating system so that 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 protect the application or service more thoroughly if you deploy the virtual machines and configure failover clustering on separate Hyper-V host computers. When you implement failover clustering 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, or *Guest Cluster Across Hosts*, is optimal for virtual machines that run 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, including any Windows Server 2012 services and any applications, such as clustered SQL Server and Exchange Server.
- Hyper-V virtual machines can use Fibre Channel-based connections to share storage, which is specific only to Microsoft Hyper-V Server 2012. You also can implement iSCSI connections from the virtual machines to the shared storage.


You should deploy multiple network adapters on the host computers and the virtual machines. Ideally, 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 that the client computers use.

NLB

NLB works with virtual machines in the same manner as 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 distributes client requests among the hosts transparently, and it enables the clients to access the cluster by using a virtual host name or a virtual IP address. From the client computer's perspective, the cluster 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, and Exchange Server Client Access servers.

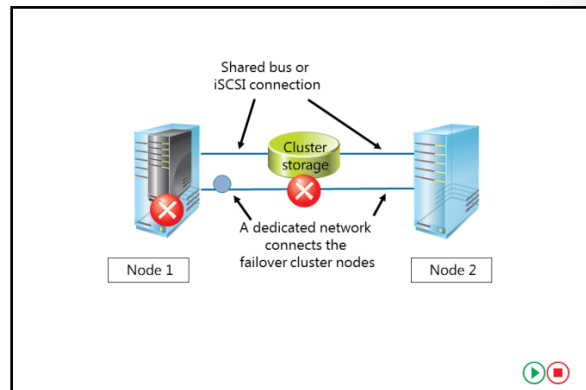
When you configure an NLB cluster, you must install and configure the application on all virtual machines. After you configure the application, you install the NLB feature in Windows Server 2012 within each virtual machine's guest operating system (not on the Hyper-V hosts). Then you must configure an NLB cluster for the application. Older versions of Windows Server also support NLB, so the guest operating system is not limited to Windows Server 2012 only.

 **Note:** As with older Windows Server versions, you should not implement NLB and failover clustering within the same guest operating system. The two technologies conflict with one another.

How Does a Failover Cluster Work with Hyper-V?

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. One at a time, the nodes run the virtual machine. However, you can 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 for other reasons, or when unplanned downtime of one node occurs because of hardware failure, power outage, or similar reasons.



The failover process consists of the following steps:

1. The node where the virtual machine runs owns the clustered instance of the virtual machine and controls access to the shared bus or iSCSI connection to the cluster storage. The node also has ownership of any disks, or 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 indicate that a node is functioning and communicating on the network. The default heartbeat configuration specifies that each node send a heartbeat over Transmission Control Protocol/User Datagram Protocol (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, failover starts after 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 can take over by configuring the Preferred and Possible Owners properties. The Preferred Owner specifies the ownership hierarchy 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 excludes it absolutely from taking over the resource in a failure situation. For example, you may implement a failover cluster by using three nodes, but you configure only two nodes as Preferred Owners. During a failover event, a third node could take over the resource if neither of the Preferred Owners is online. Although you do not configure the third node as a Preferred Owner, as long as it is a Possible Owner, the failover cluster can use it to restore access to the resource, if necessary.

You should bring resources 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 configure the cluster service to fail back to the offline node after it becomes active again. 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 of the resources associated with that instance, moves the instance, and then brings all of the instance's resources back online.

Infrastructure Requirements for Implementing Failover Clustering for Hyper-V

To deploy Hyper-V on a failover cluster, you must ensure that you meet the hardware, software, account, and network infrastructure requirements, which the following sections describe.

Hardware and Software Requirements for Failover Clustering with Hyper-V

Hardware requirements for Hyper-V clusters are the same as any other type of cluster. Lesson 1, “Planning a Failover Clustering Infrastructure” details these requirements. When you are planning hardware for Hyper-V clusters, the most important hardware components to consider are server hardware, network adapters, and HBA.

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"> • Windows Server 2012 Standard or Datacenter Edition • Same software update and service packs • Full installation or Server Core installation
Network infrastructure requirements	<ul style="list-style-type: none"> • Virtual machines are failover cluster nodes • Virtual machine applications must be cluster-aware • Network settings and IP addresses • DNS • Domain role • Domain Admin account

Furthermore, with respect to software requirements, it is mandatory that all servers in a failover cluster must run Windows Server 2012 Standard Edition or Windows Server 2012 Datacenter Edition, and they must have the same software updates and service packs installed. However, using Failover Clustering with Hyper-V has some specifics. When you plan an infrastructure for this type of clustering, you should also consider CSV design as well as virtual network switches design. These components are very important parts of an infrastructure that participates in each Hyper-V failover cluster.

Network Infrastructure Requirements

To implement a failover clustering infrastructure, you need an account with administrative permissions. Also, a failover cluster requires the following network infrastructure specifications:

- Network settings and IP addresses must 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 you do not route to your whole network infrastructure for communication between cluster nodes, ensure that each of these private networks uses a unique subnet.
- The servers in the cluster must use DNS for name resolution. You should use the DNS dynamic update protocol.
- All servers in the cluster must be in the same AD DS domain. As a best practice, all clustered servers should have the same domain role. You should avoid installing cluster nodes on domain controllers because AD DS has its own highly available mechanism. In addition, when implementing clustering on software such as Exchange Server, you should use its own mechanism to build a cluster.
- When you first create a cluster or add servers to a cluster, you must sign in to the domain with an administrator’s account on all the clusters’ servers. Additionally, if the account is not a Domain Admins account, the account must have the Create Computer Objects permission in the domain.

Implementing Storage Migration

Many scenarios may require that you move your 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 between hosts that are not in the failover cluster, resulted in downtime because you had to turn off the virtual machine. If you moved a virtual machine between two hosts, 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.

In Windows Server 2012, virtual machine migration and Storage Migration enable 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. Even more, you can now move virtual machines within the same host or between hosts without downtime, and without the need for establishing cluster infrastructure.

The process of virtual machine migration and Storage Migration is as follows:

1. You can use the Hyper-V console to start live storage migration. 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 write to both the source and destination locations. Read operations occur from only 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. Additionally, if you are moving the virtual machine to another host, Windows copies the computer configuration and associates the virtual machine with the host. If a failure occurs on the destination side, there is a failback 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 several factors, including 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 then retrieves a token that represents the data from the storage device. Then it requests movement from the source disk to the destination disk by using an offload write command with the token.

When you move a virtual machine's virtual hard disks to another location, the Virtual Machine Move Wizard provides you with three options:


- Move all the Virtual Machine's Data to a Single Location. You can specify a single destination location, such as disk file, configuration, snapshot, or smart paging.
- Move the Virtual Machine's Data to a Different Location. You can specify individual locations for each virtual machine item.
- Move Only the Virtual Machine's Virtual Hard Disk. You specify that you want to move only the virtual hard disk file.

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
- You can place storage and virtual machine configuration in different locations

Implementing Live Migration

Live Migration enables you to move running virtual machines between independent Hyper-V nodes or from one Hyper-V node to another Hyper-V node in the same cluster. With Live Migration, users who connect to the virtual machine should experience almost no server outage, and the virtual machine is not turned off. In Windows Server 2012, you can perform Live Migration between hosts in the same cluster, but also between two independent Hyper-V hosts. Migration of virtual machines from two independent hosts is called Shared Nothing Live Migration, and it works in a similar way to Storage Migration, which the previous topic described. To use the Live Migration feature, first you should enable it in the Hyper-V host properties on each host that will participate in Live Migration.



Note: You can perform Live Migration of virtual machines by using the virtual machine migration and Storage Migration method that the previous topic described, but you should be aware that Live Migration is based on failover clustering. Unlike the Storage Migration scenario, you can perform Live Migration only if the virtual machines are highly available. Windows Server 2012 introduces shared-nothing Live Migration that does not use clustering, but it works differently within a cluster than Live Migration.

You can initiate Live Migration by using:

- The Failover Cluster Manager console.
- The Microsoft System Center 2012 Virtual Machine Manager (VMM) administrator console, if you use VMM to manage your physical hosts.
- A Windows Management Instrumentation (WMI) or Windows PowerShell script.

Note: Live Migration enables you to reduce a virtual machine's perceived outage significantly during a planned failover, when you start the failover manually. Live Migration does not apply during an unplanned failover, such as when the node that is hosting the virtual machine fails.

The Live Migration Process

The Live Migration process consists of four steps:

1. **Migration setup.** When you start the failover of the virtual machine, the source node creates a TCP connection with the target physical host. The source node uses this connection 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.** You transfer the guest memory iteratively to the target host, while the virtual machine is 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 modified. During this phase of the migration, the migrating virtual machine continues to run. Hyper-V iterates the memory copy process several times, copying a smaller number of modified pages to the destination

physical computer each time. A final memory copy process copies the remaining modified memory pages to the destination physical host. Copying stops once the number of pages that the system has modified in physical memory but not rewritten to disk (*dirty pages*) drops below a certain threshold, or after 10 iterations are complete.

3. State transfer. To 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 clean-up stage finishes the migration by dismantling the virtual machine on the source host, terminating the worker threads, and signaling the completion of the migration.

When implementing Live Migration, you should:

- 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 that 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 a network adapter with a capacity of 1 Gbps or higher for this network.



Note: When you configure a dedicated network adapter for a private network, you must enable the Client for Microsoft Networks component and the File and Printer Sharing for Microsoft Networks component for that specific network adapter.

- 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 with the same architecture. Similar server hardware provides more consistency when configuring processor compatibility settings, the failover experience, and performance.
- 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. Additionally, 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. However, in Windows Server 2012, you can run multiple Live Migrations simultaneously. By default, you can perform two Live Migrations simultaneously, but you can increase this number depending on bandwidth and available storage.



Note: In Windows Server 2012 R2, you can perform Live Migration of virtual machines by using SMB 3.0 as a transport. This means that you can take advantage of key SMB features, such as SMB Direct and SMB Multichannel, which provide high-speed migration with low central processing unit (CPU) utilization.

Guidelines for Virtual Machine High Availability

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. This is because you must invest in additional server hardware to provide redundancy, and you must implement or have access to a shared storage infrastructure.

When you design your strategy for failover clustering, follow these guidelines to ensure that your design meets your organization's requirements:

When implementing Hyper-V clusters, you should:

- Identify the applications that require high availability
- Identify the application components that must be highly available
- Identify the application characteristics
- Identify the total capacity requirements
- Create the Hyper-V design
 - 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 to help determine 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 is necessary. Other applications may require that you make several servers and components, such as storage or the network, highly available. Additionally, 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 answer several questions about the application, including:
 - Is it an option to virtualize the server that is running the application? Some applications are not supported to run in a virtualized environment.
 - What options are available for making the application highly available? You can make some applications highly available by using options other than host clustering. If other options are available, evaluate the advantages and disadvantages of each.
 - What are the performance requirements for each application? Collect performance information on the servers running the applications currently to gain an understanding of the hardware requirements that are necessary for virtualizing the server.
 - What capacity is required for making the Hyper-V virtual machines highly available? As soon as you identify all of the applications that you must make highly available, you can start to design the actual Hyper-V deployment. By identifying the applications' performance, network, and storage requirements, you can define the hardware that you must implement for all of 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. However, in Windows Server 2012, you can also migrate virtual machines without using failover clustering in a process called *shared-nothing Live Migration*.

Lesson 5

Planning a Multisite 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 multisite failover clusters and their prerequisites. You will also learn about synchronous and asynchronous replication and the process of selecting a quorum mode for multisite clusters.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe challenges for planning multisite clusters.
- Describe considerations for data replication.
- Describe network infrastructure considerations.
- Describe quorum mode considerations.
- Describe the general guidelines for planning multisite clusters.

Challenges for Planning Multisite Clusters

Multisite clusters are more complex to implement than single site clusters, presenting challenges to the administrator. Storage and network issues are the most difficult aspects of implementing multisite clusters.

In a multisite cluster, there is no shared storage. 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 options for replicating data: block level hardware-based replication, software-based file replication installed on the host, or application-based replication.

Storage challenge	Description
Requires a separate storage or a data replication solution other than Microsoft	<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 Database Availability Groups
Can be 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 in 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"> • Inter-node communications are time-sensitive, so 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 host name • 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 	

Multisite data replication can be synchronous or asynchronous. Synchronous replication does not acknowledge data changes made in Site A until the data writes to Site B successfully. With asynchronous replication, data changes made in Site A write to Site B eventually.

When you deploy a multisite cluster and run the Validate a Configuration Wizard, the disk tests will not find any shared storage. Therefore, they will not run. However, you can 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 allows 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. Clients rely on DNS to find a

service or application after a failover, so you might have to adjust the DNS record's 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 internode communication (heartbeat) delay and time-out thresholds.

Designing a quorum properly for multisite clusters is also a very important part of global high availability design. In Windows Server 2012 R2, you can benefit from new dynamic quorum and dynamic witness that enhance the stability of the cluster.

Planning Data Replication in a Multisite Failover Cluster

Usually, a geographically dispersed failover cluster does not share storage between physical locations. WAN links are too slow and have too much latency to support shared storage. This means that you must have separate data instances. To have exact copies of data on both sides, geographically dispersed failover clusters must synchronize data between locations by using specialized hardware and software. Multisite data replication can be synchronous or asynchronous.

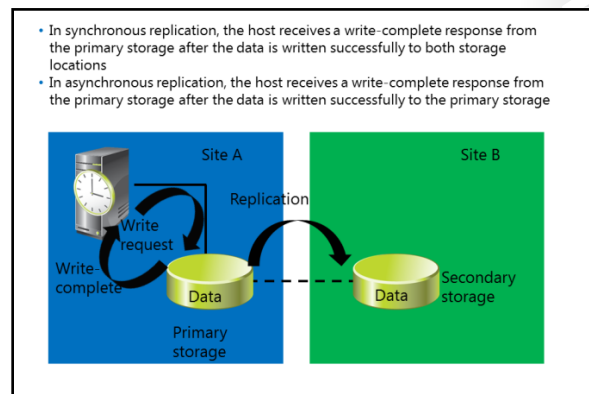
When you use synchronous replication, the host receives a write-complete response from the primary storage after the data writes successfully to both storage systems. If the data does not write successfully to both storage systems, the application must attempt to write to the disk again. With synchronous replication, both storage systems are identical.

When you use asynchronous replication, the node receives a write-complete response from the storage after the data writes successfully to the primary storage. The data writes to the secondary storage on a different schedule, depending on the hardware or software vendor's implementation. Asynchronous replication can be storage-based, host-based, or even application-based. However, not all forms of asynchronous replication are sufficient for a multisite cluster. For example, Distributed File System (DFS) Replication provides file-level asynchronous replication. However, it does not support multisite replication of failover clustering, because it was not designed for high-speed, open-file replication. Instead, it replicates smaller documents that are not open continuously.

When to Use Synchronous or Asynchronous Replication

Use synchronous replication when data loss is not acceptable. Synchronous replication solutions require low-latency disk write, 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. High disk latency can affect the application's performance and stability.

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. Asynchronous replication writes to the remote storage system after it writes to the local storage system, so the possibility of data loss during a failure increases.



Network Infrastructure Considerations

When planning multisite failover clusters, it is very important to plan the necessary network infrastructure properly. In ordinary failover clusters, you need to design only an internal network between cluster nodes, storage, and clients. However, in multisite clusters, you must plan for inter-site network links. When you are planning network infrastructure for multisite clusters, consider the following requirements:

- You must provide at least one reliable, low-latency network connection between sites. This is important for cluster heartbeats. By default, regardless of subnet configuration, heartbeat frequency (also known as *subnet delay*) is once every second (1,000 milliseconds). The range for heartbeat frequency is once every 250 to 2000 milliseconds on a common subnet and 250 to 4,000 milliseconds across subnets. By default, when a node misses a series of five heartbeats, another node will initiate failover. The range for this value, or *subnet threshold*, is three to ten missed heartbeats.
- You must provide a storage replication mechanism. Failover clustering does not provide any storage replication mechanisms, so you must provide another solution. You must also have multiple storage solutions, one for each cluster you create.
- You must ensure that all other necessary services for the cluster, such as AD DS and DNS, also are available on a second site.
- You must ensure that client connections redirect to a new cluster node when failover occurs.

When planning network infrastructure for multisite clusters, you should ensure that:

- You have enough hardware to meet the need for nodes on each site
- Each node has the same operating systems and service packs
- You have a reliable low-latency inter-cluster network
- Your network supports your storage replication mechanism
- You have sufficient infrastructure services on each site

Quorum Mode Considerations

Each failover cluster must have a defined quorum mode, so that it can determine a majority vote easily 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 use these two quorum modes only if the hardware vendor specifically recommends and supports them.

To use the Node and Disk Majority and No Majority: Disk Only modes in a multisite cluster, the shared disk require 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.

When designing automatic failover for geographically dispersed clusters, you must:

- 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 for a file-share witness


Multisite clusters can experience WAN failures in addition to node and local network failures, so Node Majority and Node and File Share Majority are better solutions for multisite clusters. If there is a WAN failure that causes the primary and secondary sites to lose communication, a majority must still be available to continue operations. If there are an odd number of nodes, you can 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 use Node Majority and the sites lose communication, you will 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, you might need to have a file share witness at a third site. In a multisite cluster, a single server can host the file share witness. However, you must create a separate file share for each cluster.

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. 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 continue to communicate and have enough nodes for a quorum.

A new feature in Windows Server 2012 R2 provides additional support and flexibility in scenarios where a cluster breaks into subsets of cluster nodes that are not aware of each other, known as *split brain syndrome*. The cluster node subset that has a majority of votes will run while others are shut down. Usually, this scenario occurs in multisite cluster deployments. If you want to start cluster nodes that do not have a majority, you can force quorum to start manually by using the **/fq** switch.

In Windows Server 2012 R2, in such scenarios, the cluster will detect partitions in the cluster automatically as soon as connectivity between nodes restores. The partition that starts by forcing a quorum is considered authoritative, and other nodes rejoin the cluster. When this happens, the cluster returns to a single view of membership. In Windows Server 2012, partitioned nodes without quorum were not started automatically, and an administrator had to start them manually with the **/pq** switch. In Windows Server 2012 R2, both sides of the split cluster have a view of the cluster membership, and they will reconcile automatically when connectivity restores.

 **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. We recommend that you use an even number of nodes for the quorum.

General Guidelines for Planning Multisite Clusters

Multisite clusters are not appropriate for every application or every business. When you design a multisite solution with a hardware vendor, clearly identify the business requirements and expectations. Not every scenario that involves more than one location is appropriate for multisite clusters.

Multisite clustering is a high-availability strategy that focuses on hardware platform availability primarily. However, specific multisite cluster configurations and deployment have availability ramifications, ranging from the ability of users to connect to the application to the quality of the application's performance. Multisite clustering can be a powerful way of managing planned and unplanned downtime, but you must examine its benefits in the context of application availability.

Multisite clusters do require more overhead than local clusters. In a local cluster, each node of the cluster is attached to the mass storage device, whereas each site of a multisite cluster must have comparable storage. Further considerations include finding vendors to set up your data replication schemes between cluster sites, possibly paying for additional network bandwidth between sites, and developing the management resources within your organization to administer your multisite cluster efficiently.

Additionally, carefully consider the quorum mode that you will use and the location of the available cluster votes.

When planning for multisite clusters, follow these guidelines:

- Clearly identify the business requirements and expectations
- Examine the benefits from multisite clusters in the context of application availability
- Consider storage vendors to set up your data replication schemes
- Carefully consider the quorum mode

Lab: Planning and Implementing a Highly Available Infrastructure by Using Failover Clustering

Scenario

One of the key goals for the Windows Server 2012 deployment at the A. Datum Corporation is that all services, applications, and servers should be highly available. A. Datum has implemented highly available storage and web applications. Now the organization wants to provide high availability for those servers and applications that they cannot make highly available through NLB.

A. Datum has decided to use failover clustering to enable high availability for these services. You must plan and implement a failover clustering solution that will address the requirements at A. Datum. Additionally, A. Datum wants to evaluate new highly available features in Windows Server 2012, such as a Scale-Out File Server, and implement automated procedures for cluster nodes updating.

Since monitoring is an important part of the A. Datum infrastructure, they implemented Microsoft System Center Operations Manager 2012. Now they want to integrate this product with VMM 2012, which they use to manage the virtual environment and Hyper-V hosts.

Objectives

After completing this lab, you will be able to:

- Design high availability for server roles.
- Deploy a failover cluster.
- Implement a Scale-Out File Server.
- Configure CAU.
- Implement highly available virtual machines.
- Implement Operations Manager and VMM integration.

Lab Setup

Estimated Time: 100 minutes

Virtual machines	20414C-LON-HOST1 20414C-LON-DC1 20414C-LON-OM1 20414C-LON-SVR1 20414C-LON-SVR2 20414C-LON-VMM1 20414C-LON-WSUS 20414C-LON-HOST2
User name	Adatum\Administrator
Password	Pa\$\$w0rd

In order to complete this lab, students will need to work with a partner. One student will start LON-HOST1, and the second student will start LON-HOST2. The two servers must connect to each other, but not to the rest of the classroom.

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

1. On the host computer, start **Hyper-V Manager**.
2. On LON-HOST1, in Hyper-V Manager, click **20414C-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Sign in using the following credentials:
 - User name: **Adatum\Administrator**
 - Password: **Pa\$\$w0rd**
5. Repeat steps two and three for the rest of the virtual machines.

Exercise 1: Designing Highly Available Server Roles

Scenario

The server team at A. Datum has presented you with the following list for server roles and applications that need to be highly available. You must identify the best option for implementing high availability for each of the server roles.

Server roles that you must make highly available are:

- AD DS domain controllers
- Web servers
- DHCP servers
- File servers
- SQL Server
- Exchange Server

The main tasks for this exercise are as follows:

1. Analyze server roles that need to be highly available
2. Propose a highly available design
3. Examine the suggested proposals in the Lab Answer Key

► Task 1: Analyze server roles that need to be highly available

Analyze the list of server roles that the exercise scenario provides. Based on what you have learned so far, think about highly available solutions for these server roles.

► Task 2: Propose a highly available design

Propose a design solution for making roles from the exercise scenario highly available by answering the following questions:

1. How will you make AD DS domain controllers highly available, and which highly available technology will you use?
2. What technology or technologies will you use to make web servers highly available?
3. How can you make a DHCP server highly available, and are there any alternative solutions?
4. What is the recommended solution for making SQL Server highly available?
5. How can you make your file server highly available?
6. How can you make Exchange Server highly available? Does the same approach apply for all Exchange Server roles?

► **Task 3: Examine the suggested proposals in the Lab Answer Key**

Read the answers in the Lab Answer Key section, and then compare them with your answers. Discuss with the class.

Results: After completing this exercise, you will have completed the design of high availability for various server roles.

Exercise 2: Deploying a Failover Cluster

Scenario

The first step in implementing failover clustering for applications and virtual machines is to deploy the failover cluster. You must configure a two-node failover cluster. A. Datum has iSCSI storage in place that you will use for clusters. Additionally, you have two physical hosts that you can use as Hyper-V hosts to create a cluster.

The main tasks for this exercise are as follows:

1. Connect to iSCSI targets from both host machines
2. Install and configure failover clustering on both host machines
3. Configure the cluster
4. Validate the cluster

► **Task 1: Connect to iSCSI targets from both host machines**

1. On LON-HOST1, start the **iSCSI Initiator**.
2. Use the 172.16.0.10 address to discover and connect to the iSCSI target.
3. On LON-HOST2, start the **iSCSI Initiator**.
4. Use the 172.16.0.10 address to discover and connect to the iSCSI target.
5. On LON-HOST2, open **Disk Management**, and initialize and bring online the first iSCSI drive (Disk 2). Format it, and then name it **ClusterDisk**.
6. On LON-HOST1, open **Disk Management**, and then refresh Disk 2.

► **Task 2: Install and configure failover clustering on both host machines**

1. Install the Failover Clustering feature on LON-HOST1 and LON-HOST2.
2. Open **Failover Cluster Manager** on LON-HOST1, and then create a new cluster with LON-HOST1 and LON-HOST2 as nodes. Name the cluster **VMCluster**, and then give it the IP address **172.16.0.126**.
3. Do not add the storage to the cluster.

► **Task 3: Configure the cluster**

1. In the Failover Cluster Manager on LON-HOST1, add Cluster Disk 1 to the cluster.
2. Verify that the iSCSI disk appears as available for cluster storage.
3. From the **VMCluster.Adatum.com** node, select **More Actions**, and then configure the Cluster Quorum Settings to use default configuration.

► Task 4: Validate the cluster

1. On the LON-HOST2 machine, open the Failover Cluster Manager console.
2. Start the Cluster Validation Wizard.
3. Select to run all tests.
4. Select to test **Cluster Disk 1**.
5. Review the report. (Note: Some warnings are expected, but you should not have any errors).

Results: After completing this exercise, you will have deployed a failover cluster.

Exercise 3: Implementing a Scale-Out File Server

Scenario

A. Datum has completed the task of making storage highly available for the file services role. The organization also wants to enable high availability for the file servers by configuring the servers as members of a Scale-Out File Server, which you can use as storage for Hyper-V virtual machines.

The main tasks for this exercise are as follows:

1. Install the file server role and failover clustering on LON-SVR1 and LON-SVR2
2. Connect to the iSCSI target from both file server cluster nodes
3. Configure the Scale-Out File Server
4. Create a continuously available file share

► Task 1: Install the file server role and failover clustering on LON-SVR1 and LON-SVR2

1. On LON-SVR1, verify that the File Server role service is installed, and then add the **Failover Clustering** feature by using Server Manager.
2. On LON-SVR2, verify that the File Server role service is installed, and then add the **Failover Clustering** feature by using Server Manager.

► Task 2: Connect to the iSCSI target from both file server cluster nodes

1. On LON-SVR1, start the **iSCSI Initiator**.
2. Use the 172.16.0.10 address to discover and connect to the iSCSI target.
3. On LON-SVR2, start the **iSCSI Initiator**.
4. Use the 172.16.0.10 address to discover and connect to the iSCSI target.
5. On LON-SVR2, open **Disk Management**, and then initialize and bring online Disk 2 and Disk 3 (do not select Disk 1). Format these drives.
6. On LON-SVR1, open **Disk Management**, and then bring online Disk 3 and Disk 4 (do not select Disk 2).

► Task 3: Configure the Scale-Out File Server

1. Open the Failover Cluster Manager console on LON-SVR1.
2. Create a cluster with nodes LON SVR1 and LON SVR2.
3. Do not run validation at this point.

4. Provide **FSCluster** as the value for **Access Point for Administering the Cluster**.
5. Provide **172.16.0.127** as **IP Address** for the cluster.
6. Add **Cluster Disk 1** and **Cluster Disk 2** to the **Cluster Storage**.
7. Assign Cluster Disk 1 to a **Cluster Shared Volume**.
8. Configure quorum settings as **Typical (recommended)**.
9. Start the High Availability Wizard to configure the clustered role.
10. Select **File Server**, and then choose to configure **Scale-Out File Server for application data**.
11. Provide **AdatumFS** as the **Client Access Point**.

► **Task 4: Create a continuously available file share**

1. On LON-SVR1, use the Failover Cluster Manager console to add a file share to an existing ADatumFS cluster role. (If you receive a message that the Client Access Point is not ready, perform these steps on LON-SVR2).
2. In the New Share Wizard, choose to use **SMB Share-Applications**.
3. Choose to configure the share by volume, and then name the share **VMachines**.
4. Configure permissions for the **VMachines** share so that the LON-HOST1 and LON-HOST2 machines have full access.

Results: After completing this exercise, you will have configured a highly available file server.

Exercise 4: Configuring CAU

Scenario

A. Datum is considering using the CAU feature to update all hosts used in a failover cluster. You need to configure and validate the CAU feature to update the Windows Server 2012 host machines.

The main tasks for this exercise are as follows:

1. Configure CAU
2. Update the failover cluster

► **Task 1: Configure CAU**

1. On LON-DC1, install the Failover Clustering feature by using the Server Manager console.
2. On LON-SVR1, open the Windows Firewall with Advanced Security window, and enable the following two inbound rules:
 - **Inbound Rule for Remote Shutdown (RPC-EP-In)**
 - **Inbound Rule for Remote Shutdown (TCP-In)**
3. Repeat step two on LON-SVR2.
4. On LON-DC1, from Server Manager, open **Cluster-Aware Updating**.
5. Connect to **FSCluster**.
6. Preview the updates available for nodes in **FSCluster**.

► Task 2: Update the failover cluster

1. On LON-DC1, start the update process for **FSCLUSTER**, by selecting **Apply updates to this cluster**.
2. Accept the default values in the update wizard.
3. Wait until the update process is completed. The process is finished when both nodes have a **Succeeded** value in the **Last Run status** column.
4. On LON-SVR1, open **Cluster- Aware Updating**, and then connect to **FSCLUSTER**.
5. Select the **Configure cluster self-updating options** option.
6. Choose to add the CAU clustered role with the self-updating mode enabled to this cluster.
7. Configure self-updating to perform **weekly**, on **Sundays** at **4:00 AM**.
8. Close the wizard.

Results: After completing this exercise, you will have configured Cluster-Aware Updating (CAU).

Exercise 5: Implementing Highly Available Virtual Machines

Scenario

You can make certain server roles highly available only by deploying highly available virtual machines that run the server role. You need to configure and validate highly available virtual machines on the failover cluster.

The main tasks for this exercise are as follows:

1. Move a .vhd file to the highly available storage
2. Configure the Hyper-V nodes to use the scale-out file server cluster
3. Configure the virtual machine as highly available
4. Perform a Live Migration for the virtual machine
5. Validate high availability in the event of storage failure

► Task 1: Move a .vhd file to the highly available storage

1. Open **File Explorer** on LON-HOST1.
2. Copy the 20414C-LON-CORE.vhd file from the **E:\Program Files\Microsoft Learning\20414\Drives\20414C-LON-CORE\Virtual Hard Disks** to the **\\AdatumFS\VMachines\LON-CORE** folder. (Note: Depending upon your host environment, your drive letter may differ. You also need to create the LON-CORE folder manually).

► Task 2: Configure the Hyper-V nodes to use the scale-out file server cluster

1. On the LON-HOST1, open **Hyper-V Settings** in the Hyper-V Manager console.
2. Configure the location of Virtual Machines and Virtual Hard Disks as **\\adatumfs\VMachines**.
3. Repeat the same procedure on LON-HOST2.

► Task 3: Configure the virtual machine as highly available

1. On LON-HOST1, open the Failover Cluster Manager console.
2. Create new virtual machine and name it **LON-CORE**.
3. Assign **768 MB** for the **Startup memory**.

4. Connect the virtual machine to **External Network**.
5. Use the existing virtual hard drive located at: `\\ADatumFS\VMachines\LON-CORE`.
6. Configure processor compatibility options for migration to a host with a different processor version.
7. Start the LON-CORE virtual machine.
8. Sign in as **Adatum\Administrator** with the password **Pa\$\$w0rd**.

► **Task 4: Perform a Live Migration for the virtual machine**

1. On the LON-HOST2, open **Windows PowerShell**.
2. In the Windows PowerShell window, type **ping lon-core -t**, press Enter, and then verify that you receive replies.
3. In Failover Cluster Manager on LON-HOST2, perform a Live Migration of the 20414C-LON-CORE virtual machine, and then move it to LON-HOST2.
4. Switch to the Windows PowerShell window, and then monitor the Internet Control Message Protocol (ICMP) packets. You should have minimal (one or two) or no lost packets.

► **Task 5: Validate high availability in the event of storage failure**

1. On LON-SVR1, in Failover Cluster Manager, identify the owner of the AdatumFS.
2. In Hyper-V Manager on LON-HOST1, disconnect the network on the machine that is the owner of AdatumFS.
3. Switch back to Failover Cluster Manager on the current owner of AdatumFS, and then ensure that the owner node changes for AdatumFS.
4. Switch to the Windows PowerShell window on LON-HOST2, and then make sure that there are no lost packets.
5. Reconnect the network to the virtual machine that you disconnected.
6. On LON-HOST1, in Failover Cluster Manager, right-click **20414C-LON-Core**, and then click **Shut Down**.
7. On LON-HOST1 and LON-HOST2, close Failover Cluster Manager.

Results: After completing this lab, students will have virtual machines implemented in a highly available infrastructure.

Exercise 6: Implementing Operations Manager and VMM Integration

Scenario

A. Datum has deployed both Operations Manager and VMM. Now the organization wants to configure the integration of Operations Manager and VMM and implement Performance and Resource Optimization (PRO). You will configure this integration and then validate the PRO implementation.

The main tasks for this exercise are as follows:

1. Import management packs
2. Enable VMM integration with Operations Manager
3. Validate PRO integration
4. Prepare for the next module

► Task 1: Import management packs

1. Open the Operations console on **LON-OM1**, and then navigate to the **Management Packs** node.
2. Import all management packs from the location **\\LON-VMM1\C\$\Program Files\Microsoft System Center 2012\Virtual Machine Manager\Management Packs**.

► Task 2: Enable VMM integration with Operations Manager

1. On LON-VMM1, open the Virtual Machine Manager Console, navigate to **Settings->System Center Settings**, and then open **Operations Manager Server**. (Note: If the VMM console does not start, open **Services** from Administrative Tools and check if all VMM services set to start automatically are running).
2. Add the Operations Manager server **lon-om1.adatum.com** by using the **Administrator RunAs** account.
3. Monitor the details in the Jobs window to ensure that the server was added successfully. If the job does not complete successfully, restart LON-OM1 and LON-VMM1 and repeat steps one and two.

► Task 3: Validate PRO integration

1. On LON-OM1, in the **Monitoring** workspace, expand **Monitoring**, expand **PRO**, and then click **PRO Object State**.
2. Ensure that state of LON-VMM1 is healthy.
3. On LON-VMM1, in the VMM Manager console, open **System Center Settings**.
4. In the Operations Manager Settings window, click **Test PRO**.
5. In the Jobs window, look for a job called **PRO Diagnostic**. Select it, and then click **Details**.
6. Make sure that all tasks in this job completed successfully.

► Task 4: 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 **20414C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps two and three for 20414C-LON-SVR1, 20414C-LON-SVR2, 20414C-LON-VMM1, 20414C-LON-WSUS and 20414C-LON-OM1.
5. On LON-HOST1, from Server Manager, remove the Failover Clustering feature by using the Server Manager console. Repeat the same procedure on LON-HOST2.
6. Restart LON-HOST1 and LON-HOST2. When the server starts, select **20414C-LON-HOST1** in the start menu.

Results: After completing this exercise, students will have Performance and Resource Optimization (PRO) implemented.

Question: What is the benefit of using CAU?

Question: What are the main benefits of having Performance and Resource Optimization (PRO) implemented in VMM?

Module Review and Takeaways

Review Questions

Question: What is the most important benefit of Scale-Out File Server?

Question: Does Live Migration require that you implement a cluster?

Question: How do you replicate data between storage in a multisite cluster?

Tools

Tools	Use	Where to find
Failover Cluster Manager	Manages clusters	Administrative Tools
Windows PowerShell	Command-line management of Windows Server	Administrative Tools
VMM	Virtual environment management	Start menu
Hyper-V Manager	Virtual machines management	Administrative Tools

Best Practice:

- Try to avoid using a quorum model that depends on a disk alone.
- Use CSVs for Hyper-V high availability or Scale-Out File Server.
- Ensure that in case of one node failure, other nodes can handle the load.
- Plan multisite clusters carefully.
- Develop standard configurations before you implement highly available virtual machines. You should configure the host computers as identically as possible. To make sure that you have a consistent Hyper-V platform, you should configure standard network names, and use consistent naming standards for CSV volumes.
- Implement Virtual Machine Manager (VMM). VMM provides a management layer on top of Hyper-V and Failover Cluster Management that can prevent you from making mistakes when you manage highly available virtual machines. For example, it blocks you from creating virtual machines on storage that is inaccessible from all nodes in the cluster.

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
Virtual machine failover fails after I implement CSV and migrate the shared storage to CSV.	The CSV home folder is located on the host server system drive. You cannot move it. If the host computers use different system drives, the failovers will fail because the hosts cannot access the same storage location. All failover cluster nodes should use the same hard drive configuration.
Four hours after restarting a Hyper-V host that is a member of a host cluster, there are still no virtual machines running on the host.	By default, virtual machines do not fail back to a host computer after they migrate to another host. You can enable failback on the virtual machine properties in Failover Cluster Management, or you can implement PRO in VMM.

Module 9

Planning and Implementing a Business Continuity Strategy

Contents:

Module Overview	9-1
Lesson 1: Overview of Business Continuity Planning	9-2
Lesson 2: Planning and Implementing Backup Strategies	9-10
Lesson 3: Planning and Implementing Recovery	9-19
Lesson 4: Planning and Implementing Backup and Recovery of Virtual Machines	9-27
Lab: Implementing a Virtual Machine Backup Strategy with DPM	9-31
Module Review and Takeaways	9-37

Module Overview

A business continuity strategy represents a set of standards, procedures, guidelines, and policies that an enterprise creates to ensure that it can maintain its critical business functions. This strategy outlines activities that you should perform daily to maintain service, consistency, and recoverability.

In this module, you will learn how to plan and implement a business continuity strategy for your organization. Also, you will learn how to plan and implement backup and recovery strategies, including virtual machine backup and recovery.

Objectives

After completing this module, you will be able to:

- Understand the importance of business continuity planning.
- Plan and implement backup strategies.
- Plan and implement recovery.
- Plan and implement virtual machine backup and recovery.

Lesson 1

Overview of Business Continuity Planning

The main objectives of business continuity planning are to recognize and identify your organization's critical operations and risks, and provide effective prevention and recovery should an event occur that could impact your organization's day-to-day business activities. Such events include critical data deletion, disk failure, or server failure.

Lesson Objectives

After completing this lesson, you will be able to:

- Identify requirements for business continuity planning.
- Create strategies for implementing business continuity.
- Understand service level agreements (SLAs).
- Describe components of a backup strategy.
- Describe components of a restore strategy.
- Define considerations for using the Windows Server Backup feature in the Windows Server® 2012 operating system.
- Define considerations for using the Microsoft® System Center 2012 Data Protection Manager.
- Define considerations for using Windows Azure™ Backup.

Collecting Requirements for Business Continuity Planning

You collect requirements for business continuity planning during an initial analysis phase and utilize the results to create your entire infrastructure. Business planning requirements differ depending on the structure and functions of an organization.

Here are several requirements that you should consider for your business continuity plan:

- SLAs for your Information Technology (IT) systems, both hardware and software.
- Contact information for and technical background of your backup administrators.
- The availability of a secondary site from which you can access critical applications and application data for critical business functions.
- Possible workaround solutions.
- Maximum outage time allowed for your applications.

Requirements for business continuity planning should include:

- SLAs for the IT systems
- Contact info and technical background of personnel assigned to recovery
- A secondary site
- Workaround solutions
- Maximum outages allowed for applications

IT employees, business managers, and other high-level decision makers should work together to create a list of requirements. Business managers should perform risk analysis, and they should understand how any failure could impact business. Business managers must decide which applications are critical for their business. They should determine the length of recovery times for each critical application to help define the appropriate backup and restore strategy.



Note: Organizations have different requirements based on their business infrastructure and goals. The requirements for business continuity planning should not be static. You should evaluate and update requirements regularly, and you should reevaluate your business continuity planning every few months.

Strategies for Implementing Business Continuity

To plan your strategies for implementing business continuity, you should collect data from:

- Business impact analysis. Business impact analysis determines an organization's critical business processes and the potential damage or loss that can result from their disruption or failure.
- Risk analysis. Risk analysis identifies possible risks and their probability. Also, risk analysis identifies the single points of failure, such as an organization's disk drives, network switches, storage, or power supply.

You can collect business continuity data from:

- Business impact analysis
- Risk analysis

Technologies for business continuity strategy include:

- NLB
- Failover clustering on physical or virtual machines
- Application-aware high availability
- Conventional data backups
- Online backups
- Virtual machine backups

Business continuity strategy will vary from organization to organization based on business requirements. Technologies that organizations will use to achieve business continuity strategy may include:

- Network Load Balancing (NLB).
- Failover clustering on physical or virtual machines.
- Application-aware high availability.
- Conventional data backups.
- Online backups.
- Virtual machine backups.

Organizations that have business-critical and IT infrastructures may implement a full-continuity, high-cost strategy that includes of different technologies. For example, some organizations will use NLB to provide high availability for web servers, use failover clustering to provide high availability for servers running Microsoft SQL Server®, and perform data backups to both tape, disk and cloud backup services such as Windows Azure Backup for protecting business-critical data. Furthermore, organizations might deploy disaster recovery centers where data from the headquarters data center will be replicated, providing site resilience.

Other organizations might decide to deploy a low-cost strategy that provides protection in situations where potential impacts are minimal or the risk is acceptable. For example, organizations might perform a backup of critical data only, accepting the risk that servers might not be available for several hours or even a day.

Service Level Agreements

SLA is a document that describes the responsibilities and specific objectives of a department, organization, or service provider. Specifically, the SLA IT describes the responsibilities of an IT department or IT service provider regarding the availability, performance and protection of the organization's business critical IT solutions and data. Additionally, SLAs often specify how quickly a provider must restore services after a failure.

Some organizations have formal SLAs, while others have general guidelines. Typically, the performance of an IT department is measured against the objectives that an SLA spells out. These metrics form part of the IT department's performance evaluation, and can influence items such as budgets and salaries. SLAs are critical to the billing structure of managed services and cloud service providers. In other types of organizations, SLAs provide less formal guidelines. A successful SLA must be realistic and achievable, and not set an unachievable standard.

An SLA may include the following elements:

- Hours of operation. The hours of operation define how when 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 a percentage of time, generally of a calendar year, that data and services are available to users. For example, a service availability of 99.9 percent per year means that data and services can have no more than 0.1 per cent per year of unplanned downtime, or 8.75 hours per year on a 24 hours a day, seven days a week basis. However, organizations should also define maintenance windows, which represent time scheduled where systems are offline for maintenance procedures such as hardware upgrades or deploying software updates.
- Recovery point objective (RPO). A recovery point objective sets a limit on how much data is lost due to failure. Recovery point objectives are contractually-determined a unit of time. For example, if an organization sets a recovery point objective of six hours, it is necessary to perform a backup every six hours or to create a replication copy on different locations at six-hour intervals. Should a failure occur, an organization would use the most recent backup, which would be no more than six hours old.

You can configure backup software to perform backups every hour and provide a theoretical recovery point objective of 60 minutes. This means if any data loss occurs 60 minutes after the last backup, only the new data created within that hour will not be restored. All other data created before the backup will be restored using the backup media. When calculating a recovery point objective, it is important to consider the time that it takes to perform a 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 recovery point objective will be one hour and 15 minutes. A realistic recovery point objective must balance your desired recovery time with your network infrastructure's realities. You should not aim for an RPO of two hours when a backup takes three hours to complete. A recovery point objective also depends on the backup software technology that you are using. 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 time when the backup began.

SLA components may include:

- Hours of operation
- Service availability
- Recovery point objective
- Recovery time objective
- Retention objectives
- System performance

- Recovery time objective (RTO). A recovery time objective is the amount of time that it takes to recover from failure. The recovery time objective varies depending on the type of failure. The loss of a motherboard on a critical server has a different recovery time objective than the loss of a disk on a critical server, because motherboard replacement takes significantly longer than disk replacement.
- Retention objectives. Retention objectives measure the length of time that you need to store backed-up data. For example, you may need to recover data quickly from the previous month, but you must store data in some form for several years. The speed at which you agree to recover data in your SLA depends on the data's age. You should consider how quickly data is recoverable or whether it must be recovered from your archives.
- System performance. System performance is an important SLA component, although it often does not relate directly to disaster recovery. Applications that an SLA includes should be available and should have acceptable response times to users' requests. If the system performance is slow, then business requirements will not be met.

Components of a Backup Strategy

Creating a thorough and correct backup strategy is one of the most important goals of the backup planning process. When planning a backup strategy, you should:

- List the data to backup. You must identify all data that requires backing up so that you can restore your data and systems if a disaster occurs. You must identify the quantity of data to back up and which Windows Server 2012 operating system volumes or files and folders to back up. This enables you to choose an appropriate storage medium and identify how long a backup or restore operation will require.
- Create a backup schedule. You must plan how frequently and at what times servers perform automated backup tasks. Most organizations perform at least daily backups.
- Choose a backup type. You must choose a backup type based on the frequency of backups and the time that a backup and restore operation takes to complete. Also, you may need to select a backup type. Your backup software may enable you to choose from the following backup types:
 - Full or Normal. A full backup is a block-level replica of all blocks on all the server's volumes. Windows Server Backup performs full backups by default.
 - Incremental. An incremental backup is a copy of only those blocks that have changed since the last full or incremental backup.
 - Differential. A differential backup is a copy of only those blocks that have changed since the last full backup. Windows Server Backup does not support differential backups.
- Choose a backup location. Organizations might choose to store back up data in following locations:
 - On premise. In this scenario, backup is located on different media in the organization data center.
 - Windows Azure Backup. Windows Server 2012 introduced a feature called Windows Azure Backup, where backup data is stored in the Windows Azure cloud platform.

When you plan a backup strategy, you must:

- List the data to backup
- Create a backup schedule
- Choose a backup type
- Choose a backup location
- Choose the backup medium

Also, you must consider the storage location of your backup media. Tapes are susceptible to magnetic fields and heat, so they should be stored away from these environmental factors. Additionally, you should store backup media offsite in case a disaster, such as a fire or flood, occurs at one of your organizations' offices.

- Choose a backup medium, based on your backup software, the size of your backups, and the timeframe in which you need to restore data. Backup media can include:
 - Tape
 - Hard disk: fixed or removable USB
 - DVD
 - Shared folder
 - Online cloud service

Tape is available in various formats, and supports various data rates and storage capacities. If you back up to tape, you should ensure that the tape format that you use is appropriate for the quantity of data that you are backing up.

The Windows Server 2012 backup feature does not support backing up to tape. Volumes and shared folders are the only supported storage media. Consider the length of time that you must retain backups. Also, consider providing support for different storage formats and media which change over time. In order to recover from backup you have to ensure that the media is still readable. Therefore, organizations have to be able to restore data from one month, six months, and 12 months ago, as well as over longer timeframes, such as several years or more. These timeframes will depend on the organization's compliance and archiving regulations and agreements.

Restore Strategies

When planning your enterprise's backup strategy, you need to develop strategies for restoring data, services, servers, and sites. You also need to make provisions for an offsite backup.

Strategies for Restoring Data

Organization's data is the most commonly recovered item in an enterprise environment, because it is more likely that users will delete files accidentally than it is for server hardware to fail or applications to cause data corruption. When considering data restore strategies, backup is not the only technology you can use for data recovery. You can address many file and folder recovery scenarios by implementing previous versions of file functionality on file shares. You can replicate data in different physical locations, to a public or private cloud, or by using DPM.

Restore strategies include:

- Data restore
- Service restore
- Site restore
- Full server restore
- Offsite backup restore

Strategies for Restoring Services

A network's functionality 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 may experience problems. For example, a major fault may replicate, requiring a restore from backup. Additionally, an enterprise backup solution must ensure timely restoration of services such as Dynamic Host Configuration Protocol (DHCP) and Active Directory Certificate Services (AD CS), in addition to important resources, such as file shares.

Strategies for Full Server Restoration

Developing a full server recovery strategy involves determining which servers you need to recover, including the RPO and RTO for critical servers. For example, suppose that you have a site with two computers that are functioning as domain controllers. When developing your backup strategy, ask yourself the following questions:

- Should you aim to have both servers capable of full server recovery with a 15-minute RPO?
- Is it necessary for only one server to recover quickly if it fails, given that recovered server will be able to provide the same network service and ensure business continuity?

When you develop the full server recovery component of your organization's enterprise backup plan, determine which servers you require to ensure business continuity. Ensure that your plan includes regular backups of these servers.

Strategies for Site Restoration

Most large organizations have branch office sites. While it might be desirable to back up all the computers at those locations, it may not be economically feasible. Developing a site recovery strategy involves determining which data, services, and servers at a specific site must be recoverable to ensure business continuity.

Strategies for Offsite Backup Restoration

Many organizations that do not store offsite backups would not be able to recover from a primary site disaster. Your backup strategies will be irrelevant if your organization's head office site, the location where you store backups, experiences a fire, flood, or cyclone.

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 may occur. You do not need to do this every day. The RPO for recovery at the offsite location—often called the disaster recovery site—is usually different from the primary site's RPO.

Benefits of Using Windows Server Backup

The Windows Server Backup feature in Windows Server 2012 consists of a Microsoft Management Console (MMC) snap-in, the **wbadmin** command, 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 to back up:

- A full server (all volumes).
- Selected volumes.
- Selected specific items for backup, such as specific folders or the system state.

Additionally, Windows Server Backup allows you to:

- Perform a bare-metal restore. A bare-metal backup contains at least all critical volumes, and allows you to restore without installing an operating system first. 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 with Windows RE to recover from a hard disk failure, or if you have to recover the whole computer image to new hardware.

Windows Server Backup in Windows Server 2012 allows you to:

- Perform a full server backup and bare-metal restore
- Backup and restore system state
- Backup and restore individual files and folders
- Exclude selected files or file types
- Select from more storage locations
- Use Windows Azure Online Backup

- Use system state. The backup contains all information necessary to roll back a server to a specific point in time. However, you must install an operating system prior to recovering the system state. System State backup does not include data on the server and includes following components:
 - AD DS and SYSVOL content. The server where the backup is performed must be a domain controller.
 - Boot files and system files.
 - Certificate Services. The server where the backup is performed must be a certification authority.
 - Cluster database. The server where the backup is performed must be a cluster node.
 - The registry.
 - Performance counter configuration information.
 - Component Services Class registration database.
- Recover individual files and folders or volumes. The Individual files and folders option enables you to back up and restore specific files, folders, or volumes. Alternatively, you can add specific files, folders, or volumes to your 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 nondedicated volumes.
- Use Windows Azure Backup. Windows Azure Backup is a cloud-based backup solution for Windows Server 2012 that enables you to back up and recover files and folders offsite, and from the public or private cloud.

If a disaster such as a hard disk failure occurs, you can perform system recovery by using a full server backup and Windows RE. This will restore your complete system to a new hard disk.

Benefits of Using System Center 2012 DPM

DPM is part of System Center 2012. DPM software enables enterprise backup and restore functionality for multiple Microsoft products and technologies. Medium- and large-size organizations might prefer to use DPM rather than Windows Server Backup, because DPM enables you to:

- Manage backup and restore jobs from a single dashboard. Also, you can set configuration settings and reporting for multiple servers.
- Support multiple Microsoft technologies, such as Windows Server, AD DS, Microsoft Exchange Server, Microsoft SQL Server, and Microsoft SharePoint® Server, including virtual servers and Microsoft client operating systems.
- Support multiple versions of Microsoft technologies, such as Windows Server 2012, Windows Server 2008 R2, Microsoft Exchange Server 2013, Exchange Server 2010.
- Support tape backup, disk backup, and Windows Azure Backup.
- Replicate backup data to an offsite location in a disaster recovery site.
- Create rich reporting on backup and restore jobs, backup media utilization, and data protection trends.

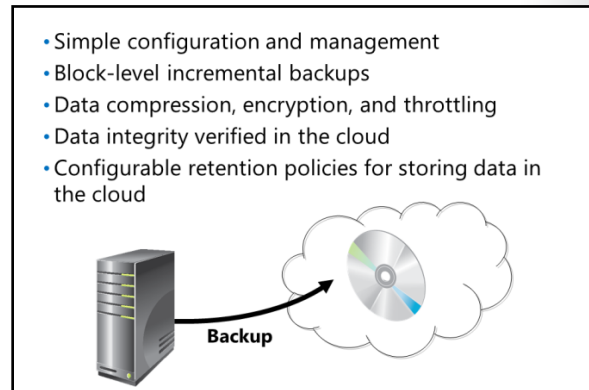
When using DPM, you can:

- Manage backup from a single dashboard
- Support multiple Microsoft technologies
- Support multiple versions of Microsoft technologies
- Support tape backup, disk backup, and Windows Azure Online Backup
- Replicate backup data to an offsite location
- Provide rich reporting

Windows Azure Backup

Windows Azure Backup is a cloud-based backup solution for Windows Server 2012. You can use this service to back up files and folders and to recover them from the public or private cloud. This enables you to provide offsite protection against data loss due to disasters. You can use Windows Azure Backup to back up and protect critical data from any location.

Windows Azure Backup software is located on the Windows Azure platform and uses Windows Azure blob storage for storing customer data. Windows Server 2012 uses the downloadable Windows Azure Backup Agent to transfer file and folder data securely to Windows Azure Backup. After you install the Windows Azure Backup Agent, the agent integrates its functionality through the Windows Server Backup interface.



- Simple configuration and management
- Block-level incremental backups
- Data compression, encryption, and throttling
- Data integrity verified in the cloud
- Configurable retention policies for storing data in the cloud

Key Features

The key features that Windows Server 2012 provides through Windows Azure Backup include:

- Simple configuration and management. Integration with the Windows Server Backup tool seamlessly backs up and recovers data to a local disk or to a cloud platform. Other features include:
 - A simple user interface for configuring and monitoring backups.
 - An integrated recovery experience to recover files and folders from a local disk or from a cloud platform.
 - Easy data recoverability for data that was backed up onto any server.
 - Scripting capability provided by the Windows PowerShell command-line interface.
- Block-level incremental backups. The Windows Azure Backup Agent performs incremental backups by tracking file and block-level changes. It transfers only the changed blocks, which reduces storage and bandwidth usage. Different point-in-time versions of the backups use storage efficiently by storing only the blocks that were changed between these versions.
- Data compression, encryption, and throttling. The Windows Azure Backup Agent ensures that data is compressed and encrypted on the server before sending it to Windows Azure Backup on the network. Therefore, Windows Azure Backup stores only encrypted data in cloud storage. Additionally, users can configure throttling and the way Windows Azure Backup uses network bandwidth when backing up or restoring information.
- Cloud verification of data integrity. The backed-up data is checked automatically for integrity after the backup completes. Therefore, you can identify any corruptions that may arise during data transfer. Windows Azure Backup will fix them automatically during the next backup.
- Configurable retention policies. Organizations can configure retention policies for the data that they store in the cloud. Windows Azure Backup Service supports retention policies that include the deletion of data which has been backed up in the cloud when it exceeds the desired retention range.



Reference Links: For more information about Windows Azure Backup, go to <http://go.microsoft.com/fwlink/?LinkID=288908>

Lesson 2

Planning and Implementing Backup Strategies

This lesson examines the required planning elements that comprise a successful, manageable, and secure backup process. You can apply these considerations when you plan backup strategies for various types of data on your network. Typically, you will distribute backup tasks among various servers and personnel in your environment.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the planning process for server backup.
- Define options for backing up computers running the Windows Server operating system.
- Define options for AD DS backup.
- Define options for backing up Windows Server 2012 roles.
- Define options for backing up file storage.
- Back up servers running Windows Server with DPM.
- Understand the DPM deployment process
- Describe options for DMP configuration.
- Understand considerations for planning a back up strategy.

Options for Planning Server Backups

When planning backups across your organization, you must ensure that you protect your mission-critical resources by:

- Identifying critical resources.
- Performing scheduled backup verification.
- Storing backups in an offsite location.
- Creating a backup logistics strategy.
- Ensuring backup security.
- Ensuring compliance and regulatory requirements.

When planning and implementing backups, you should:

- Identify critical resources
- Perform scheduled backup verification
- Store at least one copy in an offsite location
- Create a backup logistics strategy
- Ensure backup security
- Ensure compliance with regulatory requirements

Determine Critical Resources to Back Up

In an ideal scenario, you would back up all of your organization's data, and you could then restore the data instantly, as it existed at any particular point in time in the last several years. However, that type of backup strategy would be extremely expensive and would consume extensive storage. Therefore, the first step in planning a backup across your enterprise is to determine what data you need to back up.

For example, you should determine whether you need to back up every domain controller, given that AD DS information will replicate to a replacement domain controller as soon as it is promoted. Also, you should determine if it is necessary to back up every file server in all file shares if every file replicates to multiple servers through a distributed file system.

As you begin to plan your backup strategy, you also need to distinguish between technical and regulatory reasons for backing up data. Due to legal requirements, you may need to provide your business with business-critical data for the past 10 years or even longer.

When you determine what data to back up, you should consider whether:

- You are storing the data in only one place. If so, ensure that you back up that location.
- Data replicates. If so, it may not be necessary to back up each replica. However, you must back up at least one location to ensure that you can restore the data from backup.
- The server or data is a mission-critical component. If so, you should back it up.
- Each product such as Microsoft SQL Server or Microsoft Exchange Server has specific steps that you must follow to recover a server or disk that failed or to recover corrupt data.

Many organizations ensure the availability of critical services and data through redundancy. For example, Exchange Server 2010 and Exchange Server 2013 provide continuous replication of mailbox databases to other servers by using database availability groups (DAGs). Some organizations might choose not to perform backups and to use Exchange Server native data protection features instead. However, using DAGs does not mean that an organization can neglect to back up its Exchange Server Mailbox servers. It does change how an organization should back up its Mailbox servers and plan its backup strategies.

Verify Your Backups

You should have a method for verifying that each backup completes successfully. Certainly, you need to know when a backup fails. At a minimum, this means checking the logs on each server to determine whether a failure has occurred. For example, if you have configured backups to occur on each server every six hours, you need to determine how often you should check the logs. You can use an alert mechanism that is available in System Center 2012 Operations Manager to alert you if a backup fails. As a best practice, you should strive to be aware of backup failures as soon as they occur. That way, you do not discover that your backups for a particular server failed only when you need those backups to perform a recovery.

One way to verify backups is to perform regular testing of your recovery procedures by simulating a particular failure. This allows you to verify the integrity of the data that you are using to perform a recovery. Also, you can ensure that your recovery procedures resolve a failure effectively. It is better to discover that you need to add steps to your recovery procedure during a test, rather than during an actual failure.

Store Backups in an Offsite Location

Organizations should retain at least two copies of backup media, storing one copy in its data center and the other copy in an offsite location. You can determine your offsite storage location according to your organizational type, such as banks or government institutions, or by the legal obligations in the country where your organization is located. You may decide to have employees transport copies of your backup media to your offsite location, or you may copy them to the offsite location through your organization's private network. Some organizations replicate their critical data and then perform backups in both the central and offsite locations.

Create Backup Logistics

You can perform backups by using different types of media. You should store those media in a storage location that is waterproof and fireproof. You should label every backup media with information such as the date of backup, type of backup—whether full or incremental—and what data the media is storing. For example, you should include a descriptor such as "Exchange Server Managers database" or "Microsoft SQL Server Finance database". Organizations should decide if they will use a new backup media for every new backup they perform, or implement a rotation scheme where older backup media will be deleted and reused. For example, an organization might perform daily backup and use new media every day from

Monday to Friday. The administrator will store the media with the Friday backup and replace it with a new one, while all other backup media from the previous week will be overwritten from Monday to Thursday. Again, the administrator will store the backup media from next Friday and replace it with a new one. At the end of the year, the organization will have copies of data from each week of the year. When the organization no longer needs those backup media, it can reuse them.

Confirm That Backups Are Secure

A good set of backups contains all of your organization's critical data, which you must protect from unauthorized access. Although you may protect your data with 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 unauthorized users do not access critical data.

When developing an enterprise backup strategy, ensure that you store your backup data in a secure location.

You should 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.

Ensure That Compliance and Regulatory Responsibilities Are Met

System administrators should be aware of their organization's regulatory and compliance responsibilities with respect to data archival. For example, some organizations' policies require that they retain business-relevant email message data for seven years. Regulatory requirements can vary from country to country, and even within countries. When developing your organization's data protection strategy, you should schedule a meeting with your organization's legal team to determine precisely which data you need to store and for how long.

Options for Backing Up Servers Running Windows Server

When planning a Windows Server backup strategy, you should first determine the function of each of the servers, such as a file server, a database server running Microsoft SQL Server, or a mail server running Exchange Server. Backup procedures for Windows Server depend on functionality and the role of the server in the IT infrastructure.

You will perform different backup procedures for Windows Server, depending on the:

- Installed Windows Server roles, such as AD DS, DNS, DHCP, or Microsoft Internet Information Services (IIS). Each server role has its own backup procedure, with which administrators should be familiar. For example, administrators will back up AD DS by backing up the system state of the server where the AD DS server role is installed. However, administrators will back up DHCP by backing up the DHCP database located on the DHCP server.

Backup procedures for Windows Server are different depending on:

- Windows Server roles installed (AD DS, DNS, DHCP, IIS)
- Data stored on the server
- Microsoft applications installed on the server (Microsoft SQL Server, Microsoft Exchange Server)
- Third-party applications installed on the server
- Backup location (onsite or in a cloud)

- Data stored on the server. In some organizations, servers running Windows Server provide file services by storing users' data. Users access their data by connecting to servers running Windows Server through the network. Administrators should develop a backup strategy that details how to back up users' data and how often they should perform a backup to provide the most recent data restore in case of server failure.
- Microsoft applications installed on the server. Servers running Windows Server might have different Microsoft application products installed, such as Microsoft SQL Server, Exchange Server, or SharePoint Server. Each product has its own specifications on how and where it stores and protects the data. Therefore, administrators should learn how to create efficient backup strategies for different Microsoft applications.
- Third-party applications installed on the server. Organizations might use third-party applications with their own instructions on how to perform a backup. You configure backup procedures according to how those third-party applications store data and user settings.

For administrators, the most efficient way of backing up different servers running Windows Server is to use a centralized backup and restore software solution, such as Data Protection Manager (DPM). Compared to configuring different types of backup manually for each server running Windows Server, DPM allows you to perform different types of backup from one centralized console. In this console, DPM contains information about how to back up different applications.

Options for AD DS Backup

Backing up the AD DS role is an important procedure that should be part of any backup and recovery process or strategy. We perform backup of the AD DS role in order to restore data in different data loss scenarios, such as a deleted user object, a deleted group object, a deleted computer object, a deleted organizational unit (OU), or a corrupted AD DS database.

When backing up AD DS, you should consider:

- Your backup schedule. It is very important that you plan your AD DS backup schedule properly, because you cannot restore from a backup that is older than the deleted object lifetime that is 180 days. This is because when a user deletes an object from AD DS, information about that deletion is kept for 180 days. If you have a backup that is newer than 180 days, you will be able to restore the deleted object successfully. If your backup is older than 180 days, the restore procedure will not replicate the restored object to other domain controllers, which means the state of AD DS data will be inconsistent.
- The four backup tools available for AD DS:
 - Windows Server Backup (Wbadmin.msc).
 - The command-line interface (Wbaadmin.exe).
 - Windows Server backup cmdlets in Windows PowerShell.
 - Backup software, such as DPM.

- AD DS backups should not be older than 180 days
- Backup software:
 - Windows Server Backup (Wbadmin.msc)
 - Command-line tools (Wbaadmin.exe)
 - Windows PowerShell
 - Other backup software, such as DPM
- Backup types:
 - System state backup
 - Critical volumes
 - Full server

- Backup types. There are different backup types, including:
 - System state backup. This backs up the system data only. The data that is backed up depends on the roles that you install on the server.
 - Critical volumes. This backs up the volumes where you save system state files.
 - Full server. This backs up the complete server, including all volumes on the targeted server.

Options for Backing Up Windows Server 2012 Roles


You can back up almost all services on computers that are running Windows Server 2012 by performing a system state backup. Some services also allow configuration and data backup from their respective management consoles.

Backing Up Roles

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

Role	Method
DHCP	<ul style="list-style-type: none"> • System state backup backs up scopes and options • DHCP console backup backs up individual scopes or all scopes
Active Directory Certificate Services	<ul style="list-style-type: none"> • System state backup backs up certificate services database • CA console backup backs up certificate services database
IIS	<ul style="list-style-type: none"> • System state backup backs up Microsoft IIS data and settings • Appcmd.exe allows backup of IIS components • Ensure that website files and folders are also backed up. System state backups do not back up these items • Export and backup certificates
Network Policy and Access Services	<ul style="list-style-type: none"> • System state data backs up Network Policy and Access Services Configuration
DNS	<ul style="list-style-type: none"> • System state backup backs up configurations and zones stored on the server • You can use dnscmd.exe to export and import zones
File and Print Services	<ul style="list-style-type: none"> • System state backs up shared folder permissions and settings • File and folder backup backs up content of shared folders

Role	Method
DHCP	<ul style="list-style-type: none"> • System state backup that backs up scopes and options. • DHCP console backup that backs up individual scopes or all scopes.
Active Directory Certificate Services	<ul style="list-style-type: none"> • System state backup backs up certificate services database. • Certification authority (CA) console backup backs up certificate services database.
IIS	<ul style="list-style-type: none"> • System state backup backs up IIS data and settings. • Appcmd.exe allows backup of IIS components. • Ensure that website files and folders are backed up also. System state backups do not back up these items. • Export and back up certificates.
Network Policy and Access Services	<ul style="list-style-type: none"> • System state data backs up Network Policy and Access Services configuration.
DNS	<ul style="list-style-type: none"> • System state backup backs up configurations and zones stored on the server. • Active Directory-integrated DNS zones are already included in the System State back up • You can use dnscmd.exe to export and import zones.
File and Print Services	<ul style="list-style-type: none"> • System state backs up shared folder permissions and settings. • File and folder backup backs up the content of shared folders.

 **Note:** Some server applications, such as Exchange Server and Microsoft SQL Server, register themselves with Windows Server Backup. This allows Windows Server Backup to back up or restore these applications as necessary. In general, you should check a vendor's documentation for specific recommendations regarding successful backup of a specific application.

Options for Backing Up File Storage

In many organizations, file storage technologies and solutions are business-critical. In these situations, you must plan backup and restore scenarios carefully.

File storage solutions include:

- File servers. Depending on how important data is, you should develop different backup and restore strategies for every file server. Additionally, you should ensure that your file servers are protected by backup procedures on shared folders where data is located.
- Distributed File System (DFS). Before backing up data that is located on DFS, refer to your backup software's documentation. Many backup software solutions recognize DFS, so you can configure backup by selecting the appropriate DFS namespace. Furthermore, you should back up the system state on the DFS root server so that you can back up DFS configuration data. Depending on backup strategy, some organizations might choose to back up the replicated servers to provide an additional copy of data in a remote location.
- Storage spaces. If you organize your file servers to use storage spaces, you should back up data on the logical volumes that storage spaces will create. Some organizations might consider storage spaces as a solution for data protection, because they may provide data redundancy. However, you should be aware that, if you delete files, the only way to restore them is from your backup.
- Cluster Shared Volumes (CSVs). Windows Server 2012 provides improved interoperability with different backup software solutions when implementing your backup and restore procedures on CSVs. This is because backup software does not have to be CSV-aware. During the backup, Windows creates a distributed snapshot that is application-consistent. In addition, CSV backup procedure does not require CSV volume ownership, which means that you can perform parallel backups on the same or different CSV volumes or clustered nodes. Windows Server 2012 improves I/O performance during the backup process, compared to previous Windows Server versions.

File storage solutions include:

- File servers
- DFS
- Storage spaces
- CSV

Performing Windows Server Backups with DPM

When planning to backup Windows Server computers, consider the following benefits of DPM:

- Backup centralization. DPM uses a client/server architecture, in which the client software is installed on all the computers that you are backing up. Those clients stream backup data to the DPM server, which allows each DPM server to support an entire small to medium-sized organization. You can manage multiple DPM servers from one centralized DPM console.
- 15-minute recovery point objective. DPM allows 15-minute snapshots of protected data, such as files, folders, Exchange Server databases, Microsoft SQL Server databases, or Hyper-V® virtual machines.
- Supports Microsoft workloads. Microsoft designed DPM specifically to support Microsoft applications and products such as Exchange Server, Microsoft SQL Server, and Hyper-V. However, DPM was not designed to support non-Microsoft server applications that do not have consistent states on disk or that do not support VSS. System Center 2012 R2 added functionality to DPM that supports file-consistent backups for virtual machines running non-Microsoft server operating systems and applications.
- Disk-based backup. You can configure DPM to perform scheduled backups to disk arrays and storage area networks, and to export specific backup data to tape for retention and compliance-related tasks.
- Remote site backup. DPM's architecture allows it to back up clients in remote sites. This means that a DPM server in a head office site can perform backups of servers and clients across wide area network (WAN) links.
- Supports backup to cloud strategies. DPM supports the backup of DPM servers, which means that you can use a DPM server at a cloud-based hosting facility to back up the contents of another office's DPM server. For disaster redundancy, you can configure DPM servers to back up each other.

DPM has the following features:

- Backup centralization
- 15-minute RPO
- Support for Microsoft technologies
- Disk and tape backups
- Remote site backup
- Backup to cloud support

DPM Deployment Process

To deploy DPM, your planning must take into consideration the various requirements for data protection and recovery in your organization. Business requirements may determine the different technologies that you must protect, which DPM features you should use, how many DPM servers you want to deploy, and where you will deploy them.

If you are considering deploying DPM, you should:

- Design an appropriate DPM solution. Organizations should start by analyzing their business-critical data and performing a risk assessment, and then produce a strategy for backups and restoration of data, as well as a DPM deployment design.

If you are deploying DPM, you should:

- Design an appropriate DPM solution
- Install the DPM prerequisites
- Install SQL Server and DPM
- Verify the installation
- Deploy agents
- Perform a test backup and restore of the protected servers

- Install system requirements that are prerequisites for DPM. According to the DPM deployment design that you create, IT administrators should prepare the appropriate server infrastructure and a backup media solution, and then install the required system components.
- Install Microsoft SQL Server and DPM. A DPM deployment requires a Microsoft SQL Server 2012 or Microsoft SQL Server 2008 R2 database. When you have determined your system requirements, you will need to install DPM and Microsoft SQL Server on the same server or on different servers, depending on your deployment design.
- Verify the installation. After installing the DPM server, you should verify whether the installation completed successfully by analyzing the event and installation logs.
- Deploy agents. After you complete the verification, you should deploy agents on your computers that require protection, per your deployment design.
- Perform a test backup and restore of the protected servers. Before putting a DPM server into production, you should perform a test backup and restore of the servers that you are protecting. This enables you to ensure that the backup and restore procedures are fully functional. Also, you need to ensure that your DPM components, such as your backup media hardware, disk drives, and tape libraries, are operational.

DPM Configuration Options

DPM configuration options enable you to design and implement a solution that can help protect your organization's critical business data and system infrastructure data. Configuration settings in DPM are located in different console windows. You can switch between them by choosing the appropriate tabs on the lower-left side of the console windows.

You can configure DPM options on the following workspaces in DPM console:

- **Management.** The Management workspace enables you to install agents on computers that you want to protect. Also, you can configure different media for storing backed-up data, such as configuring disks in storage pools, configuring tape devices and libraries, and configuring Windows Azure Backup.
- **Protection.** The Protection workspace enables you to create protection groups, which each contain data from different technologies that you want to protect, such as file-server folders, Exchange Server mailboxes, Microsoft SQL Server databases, and virtual machines.
- **Recovery.** The Recovery workspace provides you with options to perform recovery of the data that you protected. In this workspace, you can choose to restore from different recovery points.
- **Reporting.** The Reporting workspace provides reports on the status of the backups in process, disk and tape usage, and usage trends. You can schedule reports to run at different time intervals and to be sent through email by using different formats, such as HTML, Microsoft Excel® 2010, or PDF.
- **Monitoring.** The Monitoring workspace provides alerts about errors and different types of backup and recovery events. Also, you can monitor the status of the scheduled, completed, and failed protection jobs.

DPM configuration tabs:

- Management
- Protection
- Recovery
- Reporting
- Monitoring

Considerations for Planning a Backup Strategy

When planning a backup strategy for your organization, you should consider:

- The maximum amount of data you can afford to lose. What is the theoretical RPO of the product? Products that offer restoration closer to the time of the failure are likely to cost more than products that offer 15-minute or 30-minute RPOs. You should determine your organization's needs. Does your organization need to be able to recover up to the last SQL Server transaction, or is a 15-minute recovery window an acceptable compromise?
- Recovery time. You need to consider how long it takes to go from failure to restored functionality. Restoring to the last SQL Server transaction is the optimal solution. If it takes two days to recover to that point, the solution does not work in your enterprise.
- Centralized backup solutions. You need to know how the product allows you to centralize your backup solution on one server. Alternatively, you must consider how backups occur directly on each server in your organization.
- Vendor support. You should not 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.
- Application compatibility. At some point, you may want to deploy a product update, but you may find that it is incompatible with the backup solution. Check with the application vendor to determine whether it supports your enterprise's backup solution.
- Recovery point capacity. You should determine the capacity of your recovery point. You also should consider how many recovery points your enterprise's data protection solution offers, and whether they are adequate for your organization's needs.

When planning a backup strategy, analyze the following parameters:

- Maximum amount of data loss and RPO
- RTO
- Centralized backup
- Vendor and application support
- Backup software compatibility
- Recovery point capacity

As part of your backup strategy, your documentation should contain information about the key parameters in the list of backup strategies above. Furthermore, you should assign administrative roles and responsibilities to backup administrators. Backup administrators will be responsible for performing backups and assuring that backup procedures execute correctly.

Lesson 3

Planning and Implementing Recovery

Disaster recovery is a methodology that prescribes the steps that you should take once a disaster has occurred. You perform disaster recovery 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. When creating a disaster recovery plan, you must balance the cost of a disaster to the organization with the cost of protection from that disaster.

Lesson Objectives

After completing this lesson, you will be able to:

- Plan server recovery.
- Plan and restore AD DS.
- Restore Windows Server 2012 roles.
- Restore files and data.
- Restore servers running Windows operating systems.
- Perform Windows Server restores with DPM.
- Consider and plan a recovery strategy.
- Implement a disaster recovery site.

Options for Planning Server Recovery

Windows Server Backup in Windows Server 2012 provides the following options for recovering data:

- Files and folders. You can back up individual files or folders as long as the backup is on a separate volume or in a remote shared folder.
- Applications and data. You can recover applications and data if the application has a VSS writer and is registered with Windows Server Backup.
- Volumes. 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.

Windows Server Backup in Windows Server 2012 recovery options include:

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

Wizard-driven restore options provide:

- Recovery destination
- Conflict resolution
- Security settings

The Recovery Wizard in Windows Server Backup provides several options for managing file and folder recovery, including:

- **Recovery Destination.** Under Recovery Destination, you can select any one of the following options:
 - Original location. Restores the data to the location to which you backed it up originally.
 - Another location. Restores the data to a different location.
- **Conflict Resolution.** Restoring data from a backup might conflict with existing versions of the data. Conflict resolution helps 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 you are recovering.

Options for AD DS Restore

AD DS is the most important service in any organization. Each domain controller stores a copy of AD DS data, so planning your AD DS backup and restore strategy carefully is a priority for any IT department. You must remember that you can restore only from backups within the past 180 days.

When planning your AD DS restore strategies, you have the following options:

- **Active Directory Recycle Bin.** You use Active Directory Recycle Bin in a scenario where an object is deleted from AD DS. Then you can use the Active Directory Administrative Center GUI console to restore the object. If your organization has an AD DS forest functional level of Windows Server 2008 R2 or newer, you must enable the Active Directory Recycle Bin in Windows PowerShell or by using the ldp.exe tool before you use this functionality. Once you enable the Active Directory Recycle Bin, you cannot disable it.
- **Nonauthoritative restore.** When you restore and bring online a failed domain controller, it will establish communication with other domain controllers. The time at which you performed the backup is before the restore date, so other domain controllers will replicate their current AD DS data to the restored domain controller. Therefore, when you have completed the restore process, the restored domain controller data will be up to date. This is because AD DS replication works by using a timestamp, so the newer AD DS object attributes always replicate to older AD DS object attributes. You should use nonauthoritative restore in a scenario when a domain controller is destroyed due to a hardware failure. To perform a nonauthoritative restore, you must start the domain controller in Directory Services Restore Mode.

Options for AD DS restore:

- AD DS Recycle Bin
- Nonauthoritative restore
- Authoritative restore
- AD DS checkpoints
- Domain controller cloning



Note: This type of restore is not convenient when a user account is deleted. If you try to restore deleted user accounts with nonauthoritative restore, the replication process from other domain controllers will delete the user account again. This happens because the restored data from backup has an older timestamp than the process that deleted the user account.

- **Authoritative restore.** Authoritative restore addresses the scenario in which the restore procedure occurs. The restored deleted object or attribute is deleted again during the replication process. Authoritative restore modifies the selected object or attribute once you perform a restore. This means that the timestamp changes to a greater value than the replicas' timestamps on other domain controllers. This change occurs before a domain controller can communicate with other network domain controllers. Therefore, in this scenario, the restore process restores the AD DS object or attribute, marks it as authoritative, and replicates it to other domain controllers, due to a greater timestamp value. In order to perform an authoritative restore, you must start the domain controller in Directory Services Repair Mode and use the ntdsutil tool.
- **AD DS snapshots.** If Windows Server 2012 is running in a virtual environment, you can restore a domain controller by using virtual machine checkpoints. Once you restore a domain controller, it will perform a nonauthoritative restore. You may use this procedure if a domain controller malfunctions. Then you can use the checkpoint to revert the domain controller state to the moment you took the checkpoint.
- **Domain controller cloning.** If Windows Server 2012 is running in a virtual environment, you can use a clone domain controller. This type of restore is useful if a domain controller has a major hardware failure, as you can deploy a new copy of the domain controller in a very short time.

Restoring Windows Server 2012 Roles

You can recover the configuration of most Windows Server 2012 components by performing a system state recovery. With some components, it is possible to perform recovery by using the tools that are available within the role's associated management console.

The following table lists the methods that you can use to recover specific Windows Server 2012 roles and associated data.

Role	Method
DHCP	<ul style="list-style-type: none"> • Restore system state data • Restore manual backup of DHCP database by using DHCP console
AD CS	<ul style="list-style-type: none"> • Restore system state data • Manually restore AD CS database by using Certification Authority console • Certificate templates are stored in AD DS.
IIS	<ul style="list-style-type: none"> • Restore system state data • Perform file and folder recovery to recover web app and site data • Use appcmd.exe to recover backups taken with appcmd.exe • Restore and bind certificates
Network Policy and Access Services	<ul style="list-style-type: none"> • Recover Network Policy and Access Services configuration by restoring system state data
DNS	<ul style="list-style-type: none"> • Restore system state • AD DS-integrated zones replicate back from AD DS • Import data by using dnscmd.exe
File and folders	<ul style="list-style-type: none"> • Ensure that permissions are restored when performing file and folder recovery • You may need to re-create shares • You may need to re-create quotas and File Server Resource Manager settings

Role	Method
DHCP	<ul style="list-style-type: none"> • Restore system state data. • Restore a manual backup of a DHCP database by using the DHCP console.
AD CS	<ul style="list-style-type: none"> • Restore system state data. • Restore an AD CS database manually by using the Certification Authority console. • Restore certificate templates in AD DS.
IIS	<ul style="list-style-type: none"> • Restore system state data. • Perform file and folder recovery to recover web applications and site data. • Use appcmd.exe to recover backups performed with appcmd.exe. • Restore and bind certificates.

MCT USE ONLY STUDENT USE PROHIBITED

Role	Method
Network Policy and Access Services	<ul style="list-style-type: none"> Recover the configuration of Network Policy and Access Services by restoring system state data.
DNS	<ul style="list-style-type: none"> Restore the system state. Replicate Active Directory–integrated zones from AD DS. Import data by using dnscmd.exe.
Files and folders	<ul style="list-style-type: none"> Ensure that you restore permissions when performing file and folder recovery. Recreate shares, if necessary. Recreate quotas and File Server Resource Manager (FSRM) settings, if necessary.



Note: The drawback of performing a system state recovery is that it recovers all system state settings. You cannot choose to restore some settings and not others. For example, you may alter your organization’s DHCP scope settings, and then determine that you need to do a recovery on the Enterprise CA database hosted on the same server. Recovering the system state will reset both the DHCP settings and the Enterprise CA database back to the state they were in when you performed the backup, thereby erasing the changes that you made to your DHCP scope settings.

Restoring Files and Data

Organizations should have procedures on how to restore files and data in their IT infrastructures. You can use the following strategies when developing a restoration procedure for files and data:

- Allow users to restore their own data.
- Perform a restore to an alternative location.
- Perform a restore to the original location.
- Perform a full volume restore.

Files and data recovery options include:

- Allow users to restore their own data
- Perform a restore to an alternative location
- Perform a restore to the original location
- Perform a full volume restore
- End-user Recovery in System Center 2012 R2 DPM

Allowing Users to Restore Their Files and Data

The most common form of data restore that IT departments perform is the restore of files and folders that users delete, lose, or corrupt in some way. Windows Server 2003 introduced the Volume Shadow Copy functionality, which you can also enable on all computers that are running Windows Server 2012. This functionality lets users restore their own files by using the file or folder properties on their workstation. Once you train end users how to do this, your IT department will spend less time recovering user data, which allows your IT personnel to focus on more valuable tasks.

From a planning perspective, you should consider increasing the generation frequency for previous file versions. This gives users more options when they try to restore their files.

Restoring Files and Data to an Alternative Location

A common restore problem is the unintentional replacement of important data when restoring from backup. This can occur when you perform a restore to a location with live data, instead of a separate location where you can retrieve the necessary data and discard unnecessary data.

When you perform a restore to an alternative location, always ensure that permissions are restored also. A common problem is that administrators recover data that includes restricted material to a location where permissions are not applied. Then unauthorized users can access that data.

Restoring Files and Data to Their Original Location

After some failures, such as data corruption or deletion, you must restore data to its original location. This is necessary when applications that access the data are preconfigured with information about the data's location.


Restoring a Volume

If a disk fails, the quickest way to restore the data could be to perform a volume recovery, instead of a selective restore 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.

End-User Recovery

End-user recovery is a feature in System Center 2012 R2 Data Protection Manager that enables users to recover data by retrieving recovery points of their files. You can enable the end-user recovery feature by performing following steps:

1. Configure the AD DS for end-user recovery feature.
2. Enable the end-user recovery feature on the DPM server.
3. Install the shadow copy client software on the computers.

 **Note:** You should copy event logs before you start the restore process. If you overwrite the event log files, such as with a system recovery, you will not be able to read event log information from before the restore began. However, that event log could be very important in helping you understand what caused the problem initially.

Restoring Windows Servers

You can perform a 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 Windows RE.

When you perform a full server restore, you have several options:

- Conduct a bare-metal restore. During a bare-metal restore, you restore an existing server in its entirety to new or replacement hardware, and then the server becomes operational. In some cases, you may have to reset the computer's AD DS account, because these accounts can become desynchronized.

When you perform a full server restore, consider:

- Performing a bare-metal restore
- Performing a System State and data restore
- Importing to Hyper-V

- Perform system state and data restore. In this scenario, you install Windows Server 2012 operating system on a new hardware and then restore the System State and the data from backup.
- Import to Hyper-V. Server backup data is written to the .vhd format, which is the format for virtual machine hard disks also. You can use full server backup data as the basis for creating a virtual machine, which ensures business continuity while sourcing the appropriate replacement hardware.

Performing Windows Server Restores with DPM

You can use DPM to restore the following:

- Versions of the Windows Server operating system from Windows Server 2003 to Windows Server 2012 on a system state, volume, share, folder, and file level.
- Hyper-V virtualized servers on a volume, folder, file, and .vhd and .vhdx level.
- SQL Server versions from Microsoft SQL Server 2000 to Microsoft SQL Server 2012 on a database level.
- Exchange Server 2003 and Exchange Server 2007 on a storage group, database, and mailbox level, and Exchange Server 2010 and Exchange Server 2013 on a database and mailbox level, including databases protected with DAG.
- Microsoft Windows SharePoint Services 3.0 and Microsoft Office SharePoint Server 2007, Microsoft SharePoint Server 2010, and Microsoft SharePoint Server 2013 on a farm, database, web application, file or list item, SharePoint search, and SharePoint front-end Web server.
- Microsoft® System Center 2012 - Virtual Machine Manager in System Center 2012 and newer versions of Virtual Machine Manager (VMM) on a database level.

When you use DPM, you can perform the following types of Windows Server restore:

- Windows Server versions from 2003 to 2012
- Hyper-V virtualized servers
- SQL Server versions from SQL Server 2000 to SQL Server 2012
- Exchange Server 2003, Exchange Server 2007, Exchange Server 2010 and Exchange Server 2013
- SharePoint Services 3.0, SharePoint Server 2007, SharePoint Server 2010 and SharePoint Server 2013
- Virtual Machine Manager

Considerations for Planning a Recovery Strategy

To ensure that your organization meets the RTOs outlined in your SLAs, you should ensure that you document all recovery procedures and that they are easy to follow. You do not want your organization's personnel to have to figure out how to perform a recovery only after a failure occurs.

Developing a Recovery Plan

When developing your recovery plan, consider the following questions:

- Where should the recovered data be located?
- When should the recovery occur?
- What data should be recovered?

When developing a recovery plan, ask yourself:

- Where should you recover?
- When should you recover?
- What should you recover?
- What are the procedures?

Choosing a Data Recovery Location

Choosing where to locate recovered data is simpler if your organization has replacement hardware, such as a replacement hard disk drive or a full server chassis. With the increased popularity of and support for virtualization, it is increasingly unnecessary to wait for specific hardware to become available when you need to perform a full server recovery. You can perform a temporary recovery to a Hyper-V host, and enable it to host the recovered server virtually until replacement hardware arrives. This gives you time to migrate from the virtual machine to a physical server.

Determining When to Recover Data

If a failure occurs, and your organization does not have an agreement with a hardware vendor for 24-hour replacement of components, you may have to wait longer for the components arrive. This could impact your RTO. Alternatively, you could perform a partial recovery to an alternate location. For example, if a remote branch office file server fails, and replacement components will arrive in 72 hours, you might choose to host the file share temporarily on another file server. Alternatively, if you are using DFS with the shared files, you might create a new replica at the site, removing it after the original file server is back in operation.

Selecting What Data to Recover

Administrators will decide what data to recover depending on the failure that occurred. The documentation should define which recovery steps you should perform. Documentation will contain different failure scenarios and necessary actions. For example, if only one database on the server is corrupted, you will perform recovery of that database only, not the entire server hosting the databases.

Testing the Recovery Procedures

You must test your backup and restore strategies before you deploy them in a production environment. Also, you should test your backup and restore strategies in production on a regular basis. You must perform testing procedures in a way that does not affect your production environment. During the testing process, you should use an isolated, nonproduction environment with a copy of the production data. Then you will ensure that if any unexpected failures occur, you will be able to recover your data, services, and servers in the time defined in your organization's SLA. Because the technology and organization IT infrastructure are dynamic, you should constantly evaluate and update your backup, restore, and testing procedures to ensure business continuity for your organization.

Considerations for Implementing a Disaster Recovery Site

You must provision disaster recovery sites with adequate hardware if you want them to be useful when a disaster affects a site.

Provisioning a Disaster Recovery Site

When provisioning a disaster recovery site, you may wish to consider virtualization as an alternative to purchasing multiple server chassis. A computer running Windows Server 2012 Datacenter with Hyper-V, provisioned with enough disk resources and RAM, can serve as a solution for disaster recovery. The Datacenter edition licenses you for an unlimited number of virtual machines. Many organizations have consolidated their existing physical infrastructure onto virtual hosts. In these situations, using a virtual host as a disaster recovery site host may be necessary.

When planning disaster recovery site deployment, consider these guidelines:

- Provision a disaster recovery site
- Ensure business continuity
- Consider using cloud-based disaster-recovery sites

Ensuring Business Continuity

The aim of a disaster recovery site is to ensure business continuity if a disaster destroys an organization's facility. Business continuity extends beyond server availability to include client computers. If a site is lost to a natural disaster, your organization may lose more than just critical servers or the client computers that you use to access those servers. If you use an enterprise solution to back up clients and servers, you can use Virtual Desktop Infrastructure (VDI) and allow thin client access to those servers until you can provision replacement hardware.

Ensuring business continuity means auditing your organization's IT infrastructure to determine which aspects of that structure must remain available for you to maintain business continuity. This requires substantial planning and preparation. You should perform drills in which you simulate the complete loss of sites to determine whether your disaster recovery preparations meet your goals. Additionally, you must keep your disaster recovery site as up-to-date as possible. This might mean that you keep an AD DS domain controller active at the site; that you have local copies of applications, such as Microsoft SQL Server and Exchange Server; and that you configure replication so that failover can occur automatically if a catastrophe occurs.

Cloud-Based Disaster Recovery Sites

Cloud-based disaster recovery sites are becoming increasingly popular. If a failure occurs, a cloud-based provider can provision an organization with the temporary infrastructure necessary to ensure business continuity. After the organization restores its infrastructure, the services of the cloud-based provider are no longer necessary.

Lesson 4

Planning and Implementing Backup and Recovery of Virtual Machines

Organizations that are planning to deploy their servers in a virtual environment must have a detailed plan for backup and recovery. This is true whether the organization has deployed a few virtual servers only, or has a highly virtualized environment with a private cloud. A backup and restore plan based on an organization's business requirements is one of the most important documents that an IT department will create.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the options for implementing a backup plan for virtual machines.
- Describe DPM and Hyper-V integration.
- Plan the DPM deployment process.
- Configure DPM options.
- Use DPM to back up virtual machines.
- Use DPM to restore virtual machines.
- Describe the considerations for implementing backup and recovery of virtual machines.

Options for Implementing Virtual Machine Backup

When planning your backup strategy for virtual machines, you should follow the technical documentation and supported backup options for each service or product that is running on your virtual machines. These services and products may include AD DS, DHCP, files and folders, Exchange Server, and Microsoft SQL Server.

You should consider the following options for implementing virtual machine backup in your organization:

- Perform backups on the physical server where virtual machines are located. This type of backup is thorough because you are backing up all of your virtual machine files, configurations, and data. However, you must ensure that the services or products running on your virtual machines support this type of backup. For example, we do not recommend this type of backup for Exchange Server if it is running in a virtual environment.
- Perform data backups on the virtualized server. With this option, you perform data backup as you would on a server installed on a physical machine. This means that you are backing up the data inside the virtual machine only. We recommend this type of backup when Exchange Server is running in a virtual environment.

Virtual machine backup options include:

- Performing backups on the physical server where virtual machines are located
- Performing backups of the data on the virtualized server
- Online backup
- Offline backup

- Perform online backup. By using this type of backup, you can back up data without interrupting a production environment. If the products installed in a virtualized environment support this type of backup, we recommend that organizations use it so that their servers can continue to work during the backup process.
- Perform offline backup. This type of backup requires that you stop the virtual machines until the backup is complete. Then the virtual machines can resume working. We do not recommend this type of backup, because it will cause downtime for services that are running on the virtual machine. You may consider performing an offline backup if no other type of backup is possible in your organization.

Using DPM to Back Up Virtual Machines

Organizations can use DPM as a backup and restore solution for their Hyper-V virtualized environment. DPM protects physical and virtual machines by installing an agent, which backs up and restores jobs. DPM integrates with Hyper-V so that it can provide virtual machines with the following:

- Protection of a stand-alone host. In this scenario, you install the DPM protection agent on the Hyper-V host where multiple virtual machines are running. A backup job will occur at the host level, which means that it will back up all virtual machine files and configuration settings.
- Protection of the virtual machine. In this scenario, you install the DPM protection agent on the virtual machine that is running on the Hyper-V host. A backup job will occur within the virtual machine, so it will back up specific application data only.
- Protection of the virtual machine that is running on the clustered host. In this scenario, virtual machines are running on a clustered Hyper-V host, and you install the DPM protection agents on all cluster nodes. The backup occurs at the host level.
- A Hyper-V host with virtual machine storage located on different servers. In this scenario, the Hyper-V host is located on one server and the virtual machine drives are located on a different server that is hosting the storage solution. You should install a DPM protection agent on both servers. The backup will occur at the host level.
- Integration of DPM with VMM. In this scenario, virtual machines are on multiple physical hosts, which you manage by using VMM. DPM integrates with VMM to protect virtual machines on every physical host. It even protects virtual machines in the process of live migration from one physical host to another. The following migration scenarios support live migration protection: intracluster, from stand-alone to clustered computers, from clustered to stand-alone computers, and between stand-alone computers. You should install DPM protection agents on every physical host where virtual machines are running. Backup occurs at the host level. If you want to integrate DPM and VMM, you must install the VMM console on the server that is running DPM.

Supported DPM virtual machine backup options include:

- Protection of a standalone host
- Protection of the virtual machine
- Protection of the virtual machine on a clustered host
- Support for location of Hyper-V host and storage on different servers
- Integration of DPM with VMM

Using DPM to Restore Virtual Machines

To recover virtual machine data or virtual machines by using DPM, you should perform the following actions:

- Recover data within your virtual machines, which requires you to install a protection agent on the virtual machine before performing the backup procedures. During the restore process, you will be recovering the data located within the virtual machine, such as files, folders, or databases. This type of restore procedure is almost identical to restoring data from physical computers.
- Recover the complete content and configuration settings of your virtual machines, which requires you to install a protection agent on the physical server on which the virtual machines are located before you start the backup procedures. For this type of restore, you have several options:
 - Recover a virtual machine to the original location. The original VHD and configuration files are deleted and replaced from the backup.
 - Recover a virtual machine to an alternate location. The VHD file and configuration files are restored to another host.
 - Recover specific items. You can choose to restore specific files, folders, volumes, or VHD files.

- Recover data within virtual machines
- Recover complete content and configuration settings of virtual machines:
 - Recover a virtual machine to the original location
 - Recover a virtual machine to the alternative location
 - Recover specific items, such as files, folders, volumes, or VHD files

Considerations for Implementing Virtual Machine Backup and Recovery

By choosing to deploy a virtualized environment and private cloud, organizations gain an optimized and manageable infrastructure and the ability to adapt to changes in business requirements. You can perform the backup and recovery process within this manageable infrastructure.

When planning backup and restore strategies for virtual environments, consider the following best practices:

- Create backup and restore strategies that will address business requirements.
- Use a software solution for managing, monitoring, and protecting virtual environments, such as System Center 2012.
- Consider using Hyper-V Replica. Hyper-V Replica allows organizations to have a replica of their virtual machines on another Hyper-V host for the purposes of restoration or addressing disaster recovery scenarios. Hyper-V replica is not suitable for recovering deleted or corrupted data that is already replicated. Instead, Hyper-V is suitable to mitigate hardware failure in the host or site failure.

Virtual machine backup and recovery best practices:

- Create the backup and restore strategies for your virtual machines, according to your business requirements
- Deploy System Center 2012
- Consider using Hyper-V Replica
- Create strategies according to specific product instructions for running in a virtual environment
- Test a virtual environment's backup and restore strategies before deploying it in a production environment
- Test a virtual environment's backup and restore strategies regularly

- Read the documentation about each application that is running on virtual environments, such as Exchange Server, Microsoft SQL Server, and SharePoint Server. Familiarize yourself with the specific recommendations for backup and recovery processes.
- Test your virtual machines' backup and restore strategies before you deploy them in a production environment.
- Test your virtual machines' backup and restore strategies when they are in production on a regular basis. Be careful not to affect the production environment when testing is ongoing. You should use an isolated, nonproduction environment with a copy of the production data for testing.

Lab: Implementing a Virtual Machine Backup Strategy with DPM

Scenario

As part of the virtualization strategy, A. Datum Corporation is converting most of its physical servers into virtual machines and deploying almost all new servers as virtual machines. A. Datum needs to develop a strategy for virtual machine backups to ensure the backup of any data stored in virtual machines. In some cases, the virtual machines need to be backed up also.

You need to configure DPM to back up the data that a virtual machine stores in in addition to the actual virtual machine itself.

Objectives

- Configure DPM.
- Backup and restore virtual machine data.
- Backup and restore virtual machines.

Lab Setup

Estimated Time: 60 minutes

Virtual machines:

20414C-LON-HOST1, 20414C-LON-DC1,

20414C-LON-SVR1, 20414C-LON-DM1

User Name: **Adatum\Administrator**

Password: **Pa\$\$w0rd**

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

1. On LON-HOST1, start **Hyper-V Manager**.
2. In Microsoft Hyper-V® Manager, click **20414C-LON-DC1**, and then, in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Sign in by using the following credentials:
 - User name: **Administrator**
 - Password: **Pa\$\$w0rd**
 - Domain: **Adatum**
5. Repeat steps 2 through 4 for **20414C-LON-SVR1** and **20414C-LON-DM1**.

Additionally, for this lab you should create virtual machine 20414C-LON-TEST, where you will perform backup and recovery. To do this, complete the following steps:

6. In Hyper-V Manager, in the Actions pane, click **New**, click **Virtual Machine**, and then, in the New Virtual Machine Wizard, on the **Before you Begin** page, click **Next**.
7. On the **Specify Name and Location** page, in the **Name** box, type **20414C-LON-TEST**, click the **Store the virtual machine in a different location** check box, and then click **Browse**.

8. On the **Select Folder** page, on the navigation pane, browse to **E:\Program Files\Microsoft Learning\20414**, click **Select Folder**, and then click **Next**. Note that your drive letter may differ based upon your host machine configuration.
9. On the **Specify Generation** page, click **Next**.
10. On the **Assign Memory** page, in the **Startup Memory** box, type **1024**, and then click **Next**.
11. On the **Configure Networking** page, in the **Connection** drop-down list box, click **External Network**, and then click **Next**.
12. On the **Connect Virtual Hard Disk** page, accept the default settings, and then click **Next**.
13. On the **Installation Options** page, ensure that **Install an operating system later** is selected, and then click **Next**.
14. On the **Completing the New Virtual Machine Wizard** page, click **Finish**.
15. Wait until the wizard finishes, and then, in Hyper-V Manager, verify that the **20414C-LON-TEST** virtual machine has been created with the state of **Off**.

Exercise 1: Configuring DPM

Scenario

A. Datum has deployed a DPM server in its London data center. You must configure the DPM components required to back up and restore data and virtual machines.

The main tasks for this exercise are as follows:


1. Configure a storage pool in Microsoft System Center 2012 Data Protection Manager.
2. Deploy DPM protection agents.
3. Configure protection groups.

► Task 1: Configure a storage pool in Microsoft System Center 2012 Data Protection Manager

1. Switch to LON-DM1, and then minimize Server Manager.
2. On the desktop, double-click **Microsoft System Center 2012 R2 Data Protection Manager**.
3. In System Center 2012 R2 DPM Administrator Console, in the **Management** workspace, add a new disk to the storage pool named **Disk1**.
4. You should see **Disk1 (Virtual HD ATA Device)** added in the **DPM Storage Pool**, displayed with a green check box.

► Task 2: Deploy DPM protection agents

1. Switch to LON-SVR1.
2. On the taskbar, open **Server Manager**, and then open **Windows Firewall with Advanced Security**.
3. In the Windows Firewall with Advanced Security window, in the navigation pane, open the **Windows Firewall with Advanced Security on Local Computer Properties**.
4. In the Windows Firewall with Advanced Security on Local Computer Properties window, in the **Inbound connections** drop-down list box, select **Allow**, and then click **OK**.
5. Close the Windows Firewall with Advanced Security window.
6. Switch to LON-HOST1.
7. Repeat steps 1 to 5 on LON-HOST1.

 **Note:** In a production environment, you might customize firewall settings on corporate servers to allow traffic from IP addresses and ports necessary for communication with DPM.

8. Switch to LON-DM1.
9. In the DPM Administrator Console, click the **Management** workspace, and then, on the navigation bar, click the **Agents** link.
10. Deploy the DPM protection agent on LON-SVR1 and LON-HOST1 by using the following settings:
 - User name and password: **Adatum\Administrator** with the password **Pa\$\$wOrd**
 - Restart method: **Yes**
11. In the Summary window, start the installation, and then wait until a **Success** status appears in the **Results** column.

► Task 3: Configure protection groups

Create a data protection group

1. On LON-DM1, in the DPM Administrator Console, click the **Protection** workspace.
2. On the ribbon, click **New** to start the Create New Protection Group Wizard, and then select the following options to create a new protection group:
 - Select Protection Group Type: **Servers**
 - Select Group Members: **LON-SVR1\All Volumes\Drive C:\Financial Data**
 - Protection group name: **Protection Group Financial Data Folder**
 - Protection Method: **Short-term protection using Disk**
 - Specify short-term goals: default settings
 - Review disk allocation: default settings
 - Replica Creation Method: default settings
 - Consistency check options: default settings
3. On the **Status** page, verify that both processes show **Success** in the **Results** column, and then click **Close**.

Create a virtual machine protection group

4. On LON-DM1, in the DPM Administrator Console, click the **Protection** workspace.
5. On the ribbon, click **New** to start the Create New Protection Group Wizard, and then select the following options to create a new protection group:
 - Select Protection Group Type: **Servers**
 - Select Group Members: **LON-HOST1\Hyper-V\Offline\20414C-LON-TEST**
 - Protection group name: **VM Protection Group**
 - Protection Method: **Short-term protection using Disk**
 - Specify short-term goals: default settings
 - Review disk allocation: default settings
 - Replica Creation Method: default settings
 - Consistency check options: default settings
6. On the **Status** page, verify that both processes show **Success** in the **Results** column, and then click **Close**.

Results: After completing these tasks, you will have created a storage pool in the Microsoft® System Center 2012 R2 Data Protection Manager containing Disk1. Next, you will have deployed Data Protection Manager (DPM) protection agents on LON-SVR1 and LON-HOST1. At the end of this exercise, you will have configured two protection groups. You will use the first protection group to protect data located in the Financial Data folder within the virtual machine LON-SVR1. You will use the second protection group to protect the virtual machine LON-TEST located on the physical host LON-HOST1.

Exercise 2: Implementing Backup and Restore for Virtual Machine Data

Scenario

The virtual machines at A. Datum will host much of the organization's corporate data. It is critical that you can back up and restore this data. Therefore, you need to configure DPM to back up and restore file folder data on a virtual machine.

The main tasks for this exercise are as follows:

1. Configure DPM to back up virtual machine data.
2. Delete data.
3. Restore the deleted data.

► Task 1: Configure DPM to back up virtual machine data

1. On LON-DM1, in the DPM Administrator Console, click the **Protection** workspace.
2. In the details pane, ensure that the status of the **Protection Group Financial Data Folder** is displayed with a green check box.
3. In the details pane, create a recovery point for the **C:\Financial Data** folder.

► Task 2: Delete data

1. Switch to LON-SVR1.
2. Open File Explorer, and then delete the **C:\Financial Data** folder.

► Task 3: Restore the deleted data

1. Switch to LON-DM1.
2. In the DPM Administrator Console, click the **Recovery** workspace.
3. Configure a recovery for the **Financial Data** folder located on LON-SVR1 with the following options:
 - Review Recovery Selection: **Financial Data**
 - Select Recovery Type: **Recover to the original location**
 - Specify Recovery Options: default settings
4. On the **Recovery Status** page, verify that the **Recovery status** is **Successful**, and then click **Close**.
5. Click the **Protection** workspace, and then click **All Protection Groups**.
6. Perform a consistency check for **C:\Financial Data**. In a few moments, the **Protection Status** should show a green check mark.
7. Switch to LON-SVR1.
8. On LON-SVR1, ensure that the **C:\Financial Data** folder has been restored.
9. Close File Explorer.

Results: After completing the exercise, you will have configured DPM to back up virtual machine data from the Protection Group Financial Data Folder created in the first exercise. After completing the backup, you will simulate data loss by deleting the Financial Data folder on LON-SVR1. Then you will restore the deleted data by using DPM.

Exercise 3: Implementing Virtual Machine Backup and Recovery by using DPM

Scenario

A. Datum has decided to implement DPM to back up and restore virtual machines. You need to configure DPM to perform a backup of a virtual machine. Then you will recover the virtual machine to a previous state.

The main tasks for this exercise are as follows:

1. Back up a virtual machine by using DPM.
2. Change a configuration in the virtual machine.
3. Restore the virtual machine.
4. To prepare for the next module.

► Task 1: Back up a virtual machine by using DPM

1. On LON-DM1, in the DPM Administrator Console, click the **Protection** workspace, and, then in the details pane, click **Protection Group: VM Protection Group**.
2. In the details pane, ensure that the status of the **VM Protection Group** is marked with a green check mark. It may take as long as 10 minutes for the status to show **OK**.
3. In the details pane, create a recovery point for the **\Offline\20414C-LON-TEST** virtual machine.

► Task 2: Change a configuration in the virtual machine

1. Switch to LON-HOST1.
2. In Hyper-V Manager, change the properties of 20414C-LON-TEST, so that startup RAM for LON-TEST is configured to be **256 MB**.



Note: This change will disrupt LON-TEST from normal operation because of the small amount of memory allocated. In the next task, you will restore the original memory setting for LON-TEST from backup.

► Task 3: Restore the virtual machine

1. Switch to LON-DM1.
2. In the DPM Administrator Console, click the **Recovery** workspace.
3. In the navigation pane, expand **Recoverable Data\Adatum.com\LON-HOST1**, and then click **All Protected Hyper-V Data**.
4. In the results pane, under **Recoverable Item**, configure a recovery for the **Offline\20414C-LON-TEST** recoverable item with the following options:
 - Select Recovery Type: **Recover to original instance**
 - Specify Recovery Options: default settings

5. On the **Recovery Status** page, verify that the **Recovery status** is **Successful**, and then click **Close**.
6. Switch to LON-HOST1.
7. In Hyper-V Manager, verify that the settings for **20414C-LON-TEST** have been restored.

► **Task 4: To prepare for the next module**

When you are finished with 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 **20414C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20414C-LON-SVR1** and **20414C-LON-DM1**.

Results: After completing this exercise, you will have configured DPM to back up the virtual machine LON-TEST located on the LON-HOST1 physical host from the VM Protection Group created in the first exercise. After completing the backup, you will simulate corruption of the virtual machine by changing the configuration of LON-TEST in the Microsoft Hyper-V® console on LON-HOST1. At the end, you will restore the corrupted virtual machine configuration by using DPM.

Question: Why is it important to prepare a detailed backup and restore strategy for your organization?

Question: Has an organization addressed potential risks by simply identifying them?

Question: Why did you create separate protection groups for backup file server data and virtual machines?

Question: Why do you need to install protection agents on a Hyper-V host computer?

Module Review and Takeaways

Best Practice: You must test your backup and restore strategies before you deploy them in a production environment. Also, you should test your backup and restore strategies in production on a regular basis. However, be careful not to affect your production environment when testing. For testing, you should use a copy of the production data on an isolated, nonproduction environment.

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
DPM protection agent installation fails	

Review Question

Question: What is more convenient for organizations: using a centralized protection solution, such as DPM, or using separate products to protect different servers, data, and services?

Real-world Issues and Scenarios

Scenario 1:

Your organization has defined its backup and restore strategy. However, after several months of running successful backup jobs, your IT manager wants to test the restore of data located on a file server. No server failures have been detected so far. When performing restore, a backup administrator finds that no corporate data has been restored because the wrong folder from the file server has been configured for backup. Now your organization will have to conduct regular testing on the restore procedures to verify that the correct data will be restored if a failure occurs.

Scenario 2:

A. Datum has seven file servers. You decide to consolidate to a new virtual clustered file server.

You copy the first of several file server's data to the new virtual clustered file server. The initial backup takes nine hours, and incremental daily backups take around 20 minutes.

Over the following year, A. Datum transfers the data successfully from all the remaining file servers.

During restore testing, administrators notice that restores are slower than backups. During a disaster recovery planning meeting, they calculate the time it would take to recover the virtual file server fully to another site. They discover that it would take more than 24 hours, and determine this to be unacceptable. To resolve this issue, the company must invest in a storage infrastructure immediately to ensure that it can recover the corporate file server. They could have avoided this problem if they had calculated the time for backup and restore of the file servers before migrating the production systems.

Tools

- DPM Administrator Console. A GUI for configuring and managing DPM.
- DPM Management Shell. Windows PowerShell® for configuring and managing DPM.

Module 10

Planning and Implementing a Public Key Infrastructure

Contents:

Module Overview	10-1
Lesson 1: Planning and Implementing Deployment of a Certification Authority	10-2
Lesson 2: Planning and Implementing Certificate Templates	10-16
Lesson 3: Planning and Implementing Certificate Distribution and Revocation	10-22
Lesson 4: Planning and Implementing Key Archival and Recovery	10-32
Lab: Planning and Implementing an Active Directory Certificate Services Infrastructure	10-36
Module Review and Takeaways	10-44

Module Overview

By using certificates, the Microsoft public key infrastructure (PKI) for the Windows Server® 2012 operating system improves the security of your information exchange. In addition, this integrated PKI provides easy administration across the Internet, extranets, intranets, and applications. In Windows Server 2012, you can use built-in services to build an internal PKI for issuing and managing certificates. In this module, you will learn how to plan and implement the various aspects of a PKI, and build an internal PKI by using Active Directory® Certificate Services (AD CS).

Objectives

After completing this module, you will be able to:

- Plan and implement certification authority (CA) deployment.
- Plan and implement certificate templates.
- Plan and implement certificate distribution and revocation.
- Plan and implement key archival and recovery.

Lesson 1

Planning and Implementing Deployment of a Certification Authority

To support PKI-enabled applications in your organization, you must plan, design, and implement a CA hierarchy. In a network, a CA serves as an authority that issues and manages security credentials and public keys for encryption. In this lesson, you will learn about the key components for a successful CA design.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe CAs.
- Gather information for designing a CA hierarchy.
- Describe internal and public CAs.
- Plan a CA hierarchy.
- Describe considerations for stand-alone and enterprise CAs.
- Describe considerations for deploying a root CA.
- Deploy a root CA.
- Describe considerations for deploying a subordinate CA.
- Describe guidelines for designing a CA hierarchy.
- Design a disaster recovery plan for a CA hierarchy.
- Migrate and upgrade CAs.

Overview of Certification Authorities

The CA is a fundamental component of a corporate PKI solution. In a Windows Server 2012 network, a CA is a server that runs Windows Server 2012 and that has AD CS installed. A CA is a mandatory component for organizations that want to manage their own certificates.

A certificate is a small file that contains several pieces of information about its owner. This data can include the owner's email address, the owner's name, the certificate usage type, the validity period, and the URLs for authority information

access (AIA) and certificate revocation list distribution point (CDP) locations. A certificate also contains a key pair that is the private key and its related public key. These keys are components of the process of validating identities, digital signatures, and encryption.

A CA can issue and revoke certificates and publish AIA and certificate revocation list (CRL) information about revoked certificates. This ensures that the CA issues certificates that can be validated to users, services, and computers.

The CA performs the following tasks:

- Verifies the identity of a certificate requestor
- Issues certificates to requestors
- Manages certificate revocation

Common CA roles:

- Root CA
- Subordinate CA
- Policy CA
- Issuing CA

Types of CAs:

- Stand-alone
- Enterprise

A CA performs multiple functions or roles in a PKI. In a large PKI, distributing CA roles among multiple servers is common. A CA performs several management tasks, including:

- Verifying the identity of the certificate requestor.
- Issuing certificates to requesting users, computers, and devices.
- Managing certificate templates.
- Managing certificate revocation.

You assign a role to each CA, depending on where you choose to locate the CA in the CA hierarchy. Common roles in a CA hierarchy include:

- **Root CA.** All other CAs in the hierarchy trust the root CA. The root CA produces and signs its own certificate. Typically, the root CA does not issue certificates to end users or computers, unless it is the only CA in the PKI.
- **Subordinate CA.** A subordinate CA trusts the root or parent CA. This trust occurs when a subordinate server receives a CA certificate from the root CA. Subordinate CAs issue certificates and implement policies.
- **Policy CA.** This is a subordinate CA, which sits directly below the root CA. You can use policy CAs to issue CA certificates to subordinate CAs that are directly below it in the hierarchy. In addition, you can use policy CAs when your organization's different divisions, sectors, or locations require separate issuance policies and procedures. Typically, you deploy a policy CA in more complex PKI environments.
- **Issuing CA.** The issuing CA issues certificates to users and computers, and is online continuously. In many CA hierarchies, an issuing CA is a subordinate CA. It authenticates users who are applying for certificates, initiates revocation requests, and assists in key recovery.

Two types of CAs are used in Active Directory Domain Services (AD DS) operations:

- **Stand-alone.** Typically, you use a stand-alone CA for offline CAs, but you can also use a stand-alone CA for a CA that is consistently available on the network. An offline CA is one that you take off the network to prevent compromising the keys that it uses to sign certificates. Additionally, stand-alone CAs do not integrate with AD DS, so you cannot use them for automatic enrollment or customization of certificate templates.
- **Enterprise.** Typically, you use an enterprise CA to issue certificates to users, computers, and services, or as a policy CA. This requires AD DS, and the enterprise CA always integrates with AD DS. You can use this type of CA as a configuration and registration database, and it provides a publication point for certificates issued to users and computers. Additionally, it supports automatic certificate enrollment and certificate template customization.

The following table lists the key differences between enterprise and stand-alone CAs.

Feature	Enterprise CA	Stand-alone CA
Integrates with AD DS	Yes	No
Autoenrollment	Yes	No
Customization of certificate templates	Yes	No
Propagation on the client's trusted root store	By using Group Policy	Manually, except when you install it in the domain

Gathering Information for the Design of a CA Hierarchy

Before beginning the design process for your CA hierarchy, you should decide how you plan to issue and manage certificates and how you will administer CAs. In this process, you should collect and analyze information about the following:

- Applications that use a PKI. You should collect information about applications that use PKI, including the technologies that those applications use. A Windows Server PKI supports the following types of PKI-enabled applications:
 - Digital signatures. Make Internet transactions more secure by encrypting and decrypting messages, authenticating the account that sends the message, and confirming that the content that is received is identical to the content that was sent.
 - Smart card logon. Implements two-factor authentication and provides a smart card and a PIN for network authentication of credentials.
 - Secure email. Provides confidential communication, data integrity, and nonrepudiation for email messages. You can enhance email security by using certificates to verify a sender's credentials, a message's point of origin, and a message's authenticity.
 - Software code signing. Protects computers from the installation of unauthorized Microsoft® ActiveX® controls or Java applets. Microsoft Authenticode®, a security feature of Windows® Internet Explorer®, technology enables software publishers to sign any form of active content digitally, including archives with multiple files.
 - IP security. Allows communications that are encrypted and digitally signed to pass between two computers, or between a computer and a router over a public network.
 - 802.1x. Allows only authenticated users to access a network, and protects the data that is transmitted across a network. An Institute of Electrical and Electronic Engineers, Inc. (IEEE) standard, 802.1x in PKI provides centralized user identification, authentication, dynamic key management, and accounting to grant authenticated network access to 802.11 wireless networks and wired Ethernet networks.
- Software restriction policy. You should identify the need for implementing a software restriction policy, and integrate it in the PKI design plan. Integration between a PKI and the software restriction policy allows you to sign applications digitally by using a certificate from your CA, and then configure a software restriction policy that will allow only signed applications to run.
- Internet authentication. You must identify how client and server authentication occurs for transactions in a client-server transmission. For example, when you use Secure Sockets Layer (SSL) encryption, a client authenticates the web server by validating the certificates that the server presents. SSL implementation requires a public or private CA, so this information is important for your CA design.
- Remote access authentication. You must identify the certificates that you want to use for mutual authentication. If your users are connecting to an internal network from the Internet, you will use certificates. For example, if you use Windows 8 DirectAccess, you must implement certificates. Furthermore, Remote Desktop (RD) Gateway requires a certificate on the RD Gateway server. If you plan to use these services, you should include them in your CA hierarchy design.

Before designing your CA hierarchy, you must identify and analyze the following information:

- Applications that use a PKI
- Accounts that use PKI-enabled applications
- Business requirements for designing a CA hierarchy
- Technical requirements for designing a CA hierarchy
- Management and security for CAs and certificates

- Encrypting File System (EFS). You must decide whether you want to implement EFS, which encrypts data. To recover EFS-encrypted data, you can implement key or data recovery, or both. Key recovery retrieves the user's private key from a CA database and imports it into any user's certificate store, thereby enabling decryption of encrypted files. To perform data recovery, you implement EFS recovery agents, which cannot access a user's private key. They can access only the randomly generated file encryption key. If you plan to implement EFS, you should consider several aspects of your CA's design, such as data recovery and Key Recovery Agents (KRAs).
- Accounts that use PKI-enabled applications. You should determine which accounts will use certificates. Several types of accounts can obtain digital certificates in a Windows Server 2012 AD DS environment. When designing CA hierarchy, you should remember that the following accounts are supported:
 - Users. When a CA issues a digital certificate to a user, it uniquely identifies the user to a PKI-enabled application. The user may obtain one or more digital certificates for different purposes on the network.
 - Computers. When a CA issues a digital certificate, or a machine certificate, to a computer, it uniquely identifies the computer to a PKI-enabled application. You use a digital certificate to authenticate a computer to other computers or users. A computer may obtain one digital certificate that is enabled for multiple purposes or several digital certificates that have a different network purpose.
 - Services. When a CA issues digital certificates to a service, it uniquely identifies the service when it participates on the network. The digital certificate authenticates the service with computers, users, or other services, and provides encryption services if the service must encrypt transmitted data. CAs do not issue certificates directly to services. A CA issues a certificate either to the computer account that hosts the service, such as Microsoft Internet Information Services (IIS), or to a user account that the service uses, such as the EFS Recovery Agent.
- Business requirements for designing a CA hierarchy. You should identify which business requirements you are addressing with your proposed CA hierarchy. Additionally, if you are implementing a CA hierarchy in a global company, you should verify that your CA hierarchy design complies with local regulations.
- Administrative requirements for CA hierarchy. Before implementing a CA hierarchy, you should define who will manage CAs in your organizations, and what kind of rights and permissions you should delegate.
- Technical requirements for designing a CA hierarchy. You must define hardware requirements, allocate resources, and provide availability for your CA hierarchy. A CA is not a hardware-demanding role. However, if you want to provide high availability, you should consider implementing CAs in a cluster. Although you can deploy AD CS on a single server, many deployments will involve multiple servers that you configure as CAs, including other servers that you configure as Online Responders, and still others that you configure as web enrollment portals. You can configure CAs on servers that run a variety of operating systems, including Windows Server 2012 R2, Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, and Windows Server 2003. However, not all operating systems support all features or design requirements. Therefore, creating an optimal design will require careful planning and testing before you deploy AD CS in a production environment.
- Management and security for CAs and certificates. You must remember that you will need to manage your CA by using the CA management console or command-line tools. In Windows Server 2012 specifically, you also can manage CAs by using Windows PowerShell®. CA servers must be highly protected. When planning a CA hierarchy, you should decide on a method for securing servers and certificates. You can choose to use hardware security modules (HSMs) or another method such as BitLocker® Drive Encryption.

Internal and Public CAs

When you plan a PKI implementation for your organization, one of the first choices that you should make is whether to use private or public CAs. If you decide to use a private CA, you will 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, which the following table specifies.

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, such as web browsers and operating systems
- Have slower certificate procurement

CA type	Advantages	Disadvantages
External public CA	<ul style="list-style-type: none"> • Trusted by many external clients such as web browsers and operating systems • Requires minimal administration 	<ul style="list-style-type: none"> • Higher cost compared to an internal CA • Cost is based per certificate • Certificate procurement is slower
Internal private CA	<ul style="list-style-type: none"> • Provides increased control over certificate management • Lower cost compared to a public CA • Customized templates • Autoenrollment 	<ul style="list-style-type: none"> • By default, not trusted by external clients such as web browsers and operating systems • Requires more administration

Some organizations use a hybrid approach to their PKI architecture. They use an external public CA for the root CA and a hierarchy of internal CAs for certificate distribution. In this way, the organizations enjoy the advantages of an internal CA, and their external clients trust their internally issued certificates. The only disadvantage is cost. Typically, a hybrid approach is the most expensive approach, because public certificates are very expensive.

Alternatively, you can deploy internal PKIs, with which you can use EFS and digital signatures. For external purposes, such as protecting web or mail servers with SSL, you must buy a public certificate. This approach is the most cost-effective solution.

Question: If you have already implemented a CA hierarchy in your environment, do you also use external certificate for some purposes? If yes, for what?

Planning a CA Hierarchy

When you decide to implement a PKI in your organization, you must decide how to design your CA hierarchy. This decision determines the core design of your internal PKI and the purpose of each CA in the hierarchy. Each CA hierarchy includes two or more CAs. Only the root CA is mandatory. Therefore, you deploy the second CA, and all those after it in the hierarchy, with a specific purpose. You cannot have a subordinate CA without having a root CA. However, a root CA can exist without subordinate CAs. If your environment is small and security is not a primary concern, then you could have a root CA only, which issues certificates to users and computers.

When planning your CA hierarchy, you might consider deploying multiple CAs, including to:

- Specialize in generating certain types of certificates
- Meet the needs of several divisions
- Improve performance
- Restrict administrative access
- Deploy a policy CA

To connect two CA hierarchies, you can implement:

- Cross-certification trust
- A bridge CA

You can create a hierarchy of CAs to do the following:

- Create CAs that specialize in generating certain types of certificates or certificates for a specific purpose.
- Meet the needs of several of your organization's divisions, which might require various CA policies or specific administrator access.
- Improve performance by offloading the certificate-issuing process to dedicated CAs.
- Restrict administrative access by delegating administrative permissions to specific CAs and not the entire PKI.

Additional scenarios for implementing a CA hierarchy in more complex environments include:

- Policy CA. In this scenario, you use policy CAs as subordinate CAs, which sit directly below the root CA in the CA hierarchy. You utilize policy CAs to issue CA certificates to other subordinate CAs that are placed directly below the policy CA. 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.
- Two-tier hierarchy. In this scenario, you have two tiers: a root CA and at least one subordinate CA. The subordinate CA is responsible for policies and for issuing requested certificates.

The benefits of creating a CA hierarchy include:

- Enhanced security and scalability that you can achieve by using dedicated CAs for specific types of tasks, such as smart card management. In addition, you can balance certificate issuance across multiple CAs.
- Flexible administration for the CA hierarchy, which enables role-based access control (RBAC) and decentralization of CA management.
- Support for commercial CAs, which allows a hierarchy's root to begin at a commercial CA root.

To connect PKIs in two or more organizations, you can choose to implement a bridge CA design model. One approach for establishing trust between separate CA hierarchies is to implement cross-certification trust. However, implementation of cross-certification can sometimes be problematic for administrative or policy reasons. When you establish bridge CA, you establish peer-to-peer trust relationships between different organizations' PKIs. This approach allows users to keep their natural trust points and enable users from different companies to interact through the bridge CA with a specified level of trust.

Considerations for Stand-Alone and Enterprise CAs

In Windows Server 2012, you can deploy two types of CAs: stand-alone and enterprise. The two types are different in terms of functionality and configuration storage, not hierarchy. The most important difference between these two CA types is Active Directory integration and dependency. A stand-alone CA can work without AD DS and does not depend on it. An enterprise CA requires AD DS, but it provides several benefits, such as autoenrollment.

The following table details the most significant differences between stand-alone and enterprise CAs.

Stand-alone CAs	Enterprise CAs
Must be used if any CA (root, intermediate or policy) is offline, because a stand-alone CA is not joined to an AD DS domain	Requires the use of AD DS and stores information in AD DS Can use Group Policy to propagate certificates to the trusted root CA certificate store
Users must provide identifying information and specify the type of certificate	Publishes user certificates and CRLs to AD DS
Does not support certificate templates	Issues certificates based on a certificate template
All certificate requests are pending until an administrator approves them	Supports autoenrollment for issuing certificates

Characteristic	Stand-alone CA	Enterprise CA
Typical usage	You use a stand-alone CA for offline CAs, but you can also use it for a CA that is available on the network consistently.	You use an enterprise CA to issue certificates to users, computers, and services. Typically, you do not use it as an offline CA.
AD DS dependencies	A stand-alone CA does not depend on AD DS. You can deploy it in environments other than Active Directory environments.	An enterprise CA requires AD DS, which you can use 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 request certificates from a stand-alone CA only by using a manual procedure or CA Web enrollment.	<ul style="list-style-type: none"> • Users can request certificates from an enterprise CA by using the following methods: • Manual enrollment • CA Web enrollment • Autoenrollment • An Enrollment Agent
Certificate issuance methods	A certificate administrator must approve all requests manually.	The CA can issue or deny requests automatically based on the template's discretionary access control list (DACL).

In general, the first CA that you deploy is a root CA, and this is a stand-alone CA. Then, you take it offline after it issues a certificate for itself and for a subordinate CA. Alternatively, you can deploy a subordinate CA as an enterprise CA, and configure it in one of the scenarios that the previous topic detailed.

Generally, you deploy stand-alone CAs when you do not want integration with AD DS and you do not require automation in the certificate issuance process. For example, if you want to issue certificates for users that are not members of the domain, a stand-alone CA is a good solution. On the other hand, we recommend an enterprise CA in scenarios where you want to leverage AD DS and Group Policy to automate most of the processes for certificates.

Considerations for Deploying a Root CA

You must make several decisions before you deploy a root CA. First, decide if you need to deploy an offline root CA. Then, determine if you need to deploy a stand-alone root CA or an enterprise root CA.

It is most common to choose an enterprise root CA if you are deploying a single-layer CA hierarchy, which contains only a single CA. However, if you deploy a two-layer hierarchy, the most common scenario is to deploy a stand-alone root CA, which will be taken offline, and an enterprise subordinate CA.

- Computer name and domain membership cannot change
- When you plan private key configuration, consider the following:
 - CSP
 - Key character length with a default of 2,048
 - The hash algorithm that is used to sign certificates issued by a CA
- When you plan a root CA, consider the following:
 - Name and configuration
 - Certificate database and log location
 - Validity period

Also, you must consider the operating system installation type. AD CS supports both the full installation and the Server Core installation scenarios. Server Core provides fewer vulnerabilities and less administrative overhead. Therefore, you should seriously consider Server Core installation if you maintain an enterprise environment.

You should also be aware that after you deploy a CA on a computer, you cannot change its name, its domain memberships, or the domain name. Therefore, it is important to determine these attributes before installing a CA.

The following table details additional considerations.

Consideration	Description
Whether you will use a cryptographic service provider (CSP) to generate new keys	<ul style="list-style-type: none"> • The default CSP in Windows Server 2012 is the Microsoft Strong Cryptographic Provider. • Any provider whose name starts with a number sign (#) is a Cryptography Next Generation (CNG) provider.
How long the key character can be	The default key length for the Microsoft Strong Cryptographic Provider is 2,048 characters. This is the minimum recommended value for a root CA. For CAs with a longer root CA certificate lifetime, you can select the 4096 bits key length.
What hash algorithm you will use to sign certificates that a CA issues	The default value of the hash algorithm is SHA-1. The largest number of operating systems supports this value. If you use newer operating systems on clients, such as Windows 7 or newer, you can choose a stronger hash algorithm such as SHA256.
How long the validity period will be for certificates that a CA issues	The default value for certificates is five years. You can modify this later to a shorter value on the certificate templates.
What the status of the root server--online or offline--will be	You should deploy the root server as an offline CA. This enhances security by safeguarding the root certificate, because it is not as susceptible to attack over the network.

If you decide to deploy an offline stand-alone root CA, you should consider the following factors:

- Before you issue a subordinate certificate from a root CA, make sure that you provide at least one CDP and AIA location that will be available to all clients. This is because, by default, a stand-alone root CA includes the CDP and AIA. Therefore, when you take the root CA off the network, revocation checks will fail, because the CDP and AIA locations will be inaccessible. When you define these locations, you should copy the CRL and AIA information manually to that location.
- You must set a validity period for CRLs that the root CA publishes for an extended time, such as one year. This means that you will have to turn on the root CA once per year to publish a new CRL, and then copy it to a location that is available to clients. If you do not do this, after the CRL on the root CA expires, revocation checks for all certificates will fail.
- You must use Group Policy to publish the root CA certificate to a trusted root certification authority store on all server and client machines. You must do this manually, because a stand-alone CA cannot do it automatically, unlike an enterprise CA. Also, you can publish the root CA certificate to AD DS by using the **certutil** command-line tool.
- If you decide to deploy an offline stand-alone root CA as a virtual machine, you should ensure that the machine is as secure as possible. We recommend that you restrict access to a CA virtual or physical machine.

Demonstration: Deploying and Configuring a Stand-Alone Root CA

When deploying a two-tier CA hierarchy with an offline root CA, you need to install and configure AD CS on a non-domain joined server. After you deploy a root CA, you can deploy subordinate CAs as domain members. However, it is very important that you configure CDP and AIA locations before taking the root CA offline.

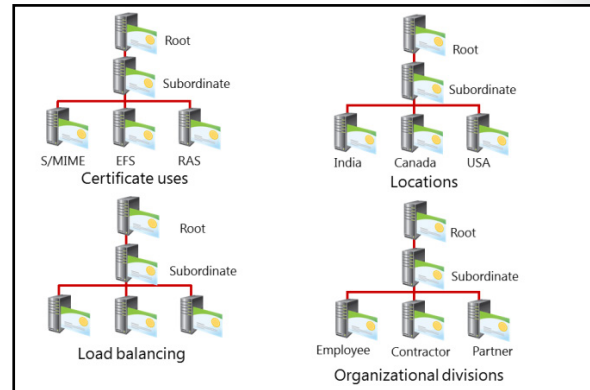
Demonstration Steps

1. Sign in to LON-CA1 as **Administrator** with the password **Pa\$\$w0rd**.
2. Use the Add Roles and Features Wizard to install the **Active Directory Certificate Services** role.
3. After installation completes successfully, click **Configure Active Directory Certificate Services on the destination server**.
4. Configure the AD CS role as a stand-alone root CA. Name it **AdatumRootCA**.
5. Set the key length to **4096**, and then accept all other values as default.
6. On LON-CA1, open the **Certification Authority** console.
7. Open the Properties window for AdatumRootCA.
8. Configure the new locations for CDP to be **http://lon-svr1.adatum.com/CertData/<CaName><CRLNameSuffix><DeltaCRLAllowed>.crl**.
9. Select the **Include in the CDP extension of issued certificates** and **Include in CRLs. Clients use this to find Delta CRL locations** options.
10. Configure the new location for AIA to be **http://lon-svr1.adatum.com/CertData/<ServerDNSName>_<CaName><CertificateName>.crt**.
11. Select the **Include in the AIA extension of issued certificates** check box.
12. Publish the certificate revocation list on LON-CA1.

Considerations for Deploying a Subordinate CA

You can use a subordinate CA to implement policy restrictions for a PKI and to issue certificates to clients. After installing a root CA for the organization, you can install one or more subordinate CAs.

When you use a subordinate CA to issue 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 issue and manage certificates and to publish certificates to AD DS. To complete this procedure, however, you must be a member of the local Administrators group or have equivalent permissions. If the subordinate CA will be an enterprise CA, you must also 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, stand-alone root CA and an enterprise subordinate CA.



Usually, you deploy a subordinate CA to achieve some of the following functionalities:

- Usage. You can issue certificates for a number of purposes, such as Secure/Multipurpose Internet Mail Extensions (S/MIME), EFS, or Remote Access Service (RAS). The issuing policy for these uses might be different, and role separation provides a basis for administering these policies.
- Organizational divisions. You might have different policies for issuing certificates that depend on an entity's role in the organization. You can create subordinate CAs to separate and administer these policies.
- Geographic divisions. Often, organizations have entities at multiple physical sites. Limited network connectivity between these sites might necessitate individual subordinate CAs for many or all sites.
- Load balancing. If you use your PKI to issue and manage a large number of certificates and have only one CA, it can result in considerable network load for that single CA. Using multiple subordinate CAs to issue certificates divides the network load between CAs.
- Backup and fault tolerance. Multiple CAs increase the possibility that your network has operational CAs available to respond to user requests.

Guidelines for Designing a CA Hierarchy

Consider the following guidelines as you design your organization's CA hierarchy:

- Decide how many CAs you require, what types of CA you will use, and where to locate them. Collect the requirements for each CA.
- Select the CA type before you deploy a CA. You cannot convert a root CA to a subordinate CA, or vice versa. Therefore, you must determine the CA type before you begin the deployment. Additionally, remember that you cannot convert stand-alone CAs to enterprise CAs.

When you design your organization's CA hierarchy, you should:

- Determine the number and type of CAs you require and where to locate them
- Select the CA type
- Deploy the root CA first, keeping it offline
- Restrict the CA hierarchy to three or four layers
- Define security levels and appropriate CA policies
- Implement role separation

- Start at the top, and work down the hierarchy. Deploy the root CA first. If you choose to deploy a private root CA, ensure that the root CA is secure by keeping it offline. Deploy the root CA in a location that is physically secure. Do not make the computer a member of any domain.
- Create a CA hierarchy that is a maximum of three to four layers deep. Deploying more than four layers adds complexity to the CA design that can be difficult to manage. Fewer than three layers might not provide enough security if your enterprise has high security requirements. However, in most cases, a two-layer hierarchy is the most common type of deployment. Smaller organizations might also choose to have a single enterprise root CA.
- Define security levels and appropriate CA policies for each CA in your hierarchy, depending on the design requirements.
- Implement role separation so that one person cannot compromise the security of your organization's PKI.

Question: If you have deployed a CA hierarchy in your company, how many layers does it have? Why did you choose this design?

Planning Disaster Recovery for a CA Hierarchy

PKI is a very important in each network infrastructure, so you must have a disaster recovery plan. A nonfunctional CA hierarchy may prevent both the issuance of new certificates and the use of existing certificates.

AD CS can stop working for many reasons. For example, if the CA service cannot start on a CA machine, you will not be able to issue new certificates. Furthermore, you will not be able to process the revocation check on an existing certificate. If the hardware on a CA server fails, that can affect AD CS. Also, failures on network infrastructure can prevent AD CS from working properly.

For any scenario, you should have a disaster recovery plan that will enable you to restore your CA hierarchy as soon as possible. One important part of any disaster recovery plan for PKI is backup. You must identify components and settings of your PKI that you need to back up regularly. Additionally, it is important that you document all changes that you perform on your CA hierarchy, so that you can restore objects easily if necessary, including:

- Certificate template definitions for all certificate templates that you create manually.
- A list of certificate templates that are published in your CA hierarchy.
- Permissions on certificate templates.
- Rights to manage a CA.
- All specific settings in the CA's properties.
- Data paths for a CA database.
- All custom locations defined for AIA and CDP points.
- Content of the CAPolicy.inf file, if you use it during CA deployment.
- The CSP that you used to protect the CA's private key.

A disaster recovery plan for CA hierarchy should include:

- Documentation about CA hierarchy settings
- Documentation about custom certificate templates
- A backup procedure for vital CA components
- A restore procedure for restoring AD CS functionality on the new server
- Solutions for potential problems and issues

By documenting these items, you can restore your CA's functionality, even without a backup. However, please note that you do need to back up the CA's private key to resume functionality after a disaster, and that it is much more reliable to maintain regular backups of the key CA components.

You can perform a backup of your CA by using Windows Server Backup. Also, you can use System State backup to perform the backup of the CA, with all necessary components.

Alternatively, you can use the **certutil** command-line tool and Windows PowerShell cmdlets to manage backups of the CA hierarchy.

For example, if you want to perform a backup of CA, use this command:

```
certutil -backup -p Password filepath
```

For exporting all registry-based settings for a CA, run the following:

```
reg export HKLM\System\CurrentControlSet\Services\CertSvc\Configuration filepath
```

To export all information about certificate templates, run the following:

```
Certutil -catemplates > filepath
```

Restoring AD CS Functionality

If you cannot get your current CA online after a disaster, you will need to enable AD CS functionality on another server. You must perform several steps to restore AD CS on another server. First, you should make sure that you have reliable and up-to-date backups of your previous CA, which this topic detailed previously.

Then, perform these steps to restore AD CS functionality on another server:

- Extend the life of the CRL file, by configuring the appropriate options on the current CA. This will provide you with additional time in case of CA downtime.
- Decommission the old CA.
- Install AD CS at the new server.
- Restore the CA configuration.
- Restore the database and templates to the CA.

To restore AD CS functionality on the new server, you must have an up-to-date backup. Additionally, you must make sure that the old CA never again connects to your network, because the new CA will have the same identity. This could cause a conflict and result in rejections of certificates.

Restoring AD CS on another computer is very similar to CA migration, which the next topic addresses.

Migrating and Upgrading CAs

As discussed earlier in this module, you can design and configure CAs to work for several years. During this time, you may want to upgrade the hardware and operating system that supports the CA. Usually, this would involve moving a CA from one computer to another.

A CA is unlike some other services in that it will continue to work if you install it on a new computer. For example, when you move a CA from one computer to another as part of deploying a new operating system version, you must keep the identity of the CA during this process. Then it can continue to work on the new hardware or operating system with the same identity.

In general, the procedure for moving a CA can be divided into two phases:

- CA backup
- CA restore

Performing a CA Backup Before a Move

You should have a CA backup even if you are not moving a CA to another computer. A CA backup is different from ordinary backup scenarios. To perform a CA backup to move a CA to another computer, you should follow this procedure:

1. If you are backing up an enterprise CA, click the **Certificate Templates** item in the CA console, and then record the names of the certificate templates that are listed. These templates are in AD DS, so you do not have to back them up. You must note which templates publish on the CA that you move because you will have to add them manually after you move the CA.
2. In the CA snap-in, right-click the **CA name**, click **All Tasks**, and then click **Back up CA** to start the Certification Authority Backup Wizard. In the backup wizard, choose to make the backup of the CA a private key, CA certificate, certificate database, or certificate database log. Also, provide an appropriate location for the backup content. For security reasons, protect a CA private key with a password.
3. After the backup is done, open Registry Editor. Locate and export the following registry subkey, which is located at:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\CertSvc\Configuration.



Note: We recommend that you save this registry key to file in the same folder with the CA backup from the previous step.

4. Uninstall the CA from the old server, and then rename the old server or disconnect it from the network permanently.

Before you begin the restore procedure, confirm that the `%SystemRoot%` folder of the target server matches the `%SystemRoot%` folder of the server that is the source of the backup.

In addition, the location of the CA restore must match the location of the CA backup. For example, if you back up the CA from the `D:\Winnt\System32\Certlog` folder, you must restore the backup to the `D:\Winnt\System32\Certlog` folder. After you restore the backup, you can move the CA database files to a different location.

To move a CA from one computer to another, perform backup and restore:

- To back up a computer, follow this procedure:
 1. Record the names of the certificate templates
 2. Back up a CA in the CA console
 3. Export the registry subkey
 4. Uninstall the CA role
 5. Confirm the `%Systemroot%` folder locations
 6. Remove the old CA from the domain
- To restore, follow this procedure:
 1. Install AD CS
 2. Use the existing private key
 3. Restore the registry file
 4. Restore the CA database and settings
 5. Restore the certificate templates

Performing a CA Restore on a New Computer

After you finalize the backup procedure successfully, you must restore the CA on another computer. The new CA should have the same name as the old CA.

To restore the CA, follow this procedure:

1. Install AD CS on the target computer. Install either **Stand-alone** or **Enterprise**, depending on the type of CA that you are moving. When you come to the **Set Up Private Key** page, click **Use existing private key**. Then choose a certificate and use its associated private key. This will provide you with the ability to use an existing certificate from an old CA.
2. On the **Select Existing Certificate** page, click **Import**, type the path of the .p12 file in the backup folder, type the password that you chose in the previous procedure to protect the backup file, and then click **OK**. When you see the prompt for **Public and Private Key Pair**, verify that **Use existing keys** is selected. This is very important, as you want to keep the same root CA certificate.
3. When prompted on the **Certificate Database** page, specify the same location for the certificate database and certificate database log as on the previous CA computer. After you choose all these options, wait for the CA setup to finish.
4. After the setup is done, open the Services snap-in to stop the AD CS service. This will restore settings from the old CA.
5. Locate the registry file that you saved in the backup procedure, and then double-click it to import the registry settings.
6. After you restore the registry settings, open the CA management console, right-click the **CA name**, click **All Tasks**, and then click **Restore CA**. This will start the Certification Authority Restore Wizard. In the wizard, you should select the **Private key and CA certificate** and the **Certificate database and certificate database log** check boxes. This specifies that you want to restore these objects from backup. Next, provide a backup folder location and verify the settings for the restore. The **Issued Log** and **Pending Requests** settings should display.
7. When the restore process finishes, restart the AD CS service.
8. If you have restored an enterprise CA, ensure that the certificate templates from AD DS that you recorded in the previous procedure are present on the new CA.

Lesson 2

Planning and Implementing Certificate Templates

Certificate templates define how certificates are requested and used. You configure templates on the CA and they are stored in the AD DS database. There are different versions of templates: the Windows 2000 Server Enterprise CA supports version 1 certificate templates; the Windows Server 2003 Enterprise Edition supports versions 1 and 2 templates; and Windows Server 2008 Enterprise supports versions 1, 2, and 3 templates. Windows Server 2012 and Windows Server 2012 R2 introduce version 4 templates, but also support all three previous template versions.

In this lesson, you will learn how to plan and implement certificate templates.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe certificate templates.
- Describe certificate template versions.
- Describe the methods for modifying certificate templates.
- Describe the considerations for designing certificate template security.
- Describe the considerations for designing certificate templates.

What Are Certificate Templates?

Certificate templates allow administrators to customize the distribution method of certificates, define certificate purposes, and mandate the type of usage that a certificate allows. Administrators can create templates easily and can deploy them quickly to the enterprise by using the built-in GUI or command-line management tools. CAs use certificate templates to:

- Define the format and content of certificates.
- Define the certificate's purpose and application policy.
- Specify which users and computers can enroll for which types of certificates.
- Specify details about the enrollment process, such as autoenrollment, enrollment only with authorized signatures, and manual enrollment.

Associated with each certificate template is its DACL, which defines the security principals that have permissions to read and configure the template, and to enroll or autoenroll for certificates based on the template. You define the certificate templates and their permissions by using a Certificate Templates snap-in. Certificate Templates snap-ins are stored in AD DS, and they remain valid within the forest. If more than one enterprise CA runs in the AD DS forest, permission changes will affect all CAs.

A certificate template defines the:

- Format and contents of a certificate
- Process for creating and submitting a valid certificate request
- Purpose and application policy
- Security principals 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

When you define a certificate template, the definition must be available to all CAs in the forest. You can ensure this 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. Additionally, the certificate template may not be available to all CAs until replication completes. Storage and replication occur automatically.



Note: Prior to Windows Server 2008 R2, only the Enterprise version of Windows Server supported the management of certificate templates. In Windows Server 2012 and Windows Server 2012 R2, you can also manage certificate templates in the Standard editions.

Certificate Template Versions

Windows Server 2012 CAs support four versions of certificate templates. Certificate templates versions 1, 2, and 3 are legacy templates from previous Windows Server versions, while version 4 is new in 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 together with Windows Server 2012 R2 correspond to version 1, version 2, version 3, and version 4, respectively.

- Version 1:**
- Introduced in Windows Server 2000, provided for backward compatibility in newer 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 you can then modify)
- Version 2:**
- Default template introduced with Windows Server 2003
 - Allows customization of most settings in the template
 - Provides several preconfigured templates when a CA is installed
- Version 3:**
- Supports advanced Suite B cryptographic settings
 - Includes advanced options for encryption, digital signatures, key exchange, and hashing
 - Supports only Windows Server 2008 and Windows Server 2008 R2 servers
 - Supports only Windows Vista and Windows 7 client computers
- Version 4:**
- Available only for Windows Server 2012 and Windows 8 clients
 - Supports both CSPs and KSPs
 - Supports renewal with the same key

Certificate template versions have some functional differences, including the following:

- **Version 1:** The Windows 2000 Advanced Server operating system provides support for version 1 certificate templates. The only modification to version 1 templates that it allows is the changing of permissions to 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, Microsoft no longer supports Windows 2000 Server.
- **Version 2:** The Windows Server 2003 Enterprise Edition operating system provides 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 your organization's requirements. Alternatively, you can duplicate a version 1 certificate template to create a new version 2 of the template. Then you can modify and secure the newly created version 2 certificate template. When you add new templates to a Windows Server 2003 enterprise CA, they are version 2, by default.
- **Version 3:** The Windows Server 2008 Enterprise operating system supports version 3 certificate templates, in addition to version 1 and version 2. 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 the default version 1 and version 2 templates to save them as version 3 templates. Additionally, Windows Server 2008 provides two new certificate templates by default: Kerberos Authentication and Online Certificate Status Protocol (OCSP) responder service. In Windows Server 2008 R2, the Standard version also supports 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.

MCT USE ONLY. STUDENT USE PROHIBITED

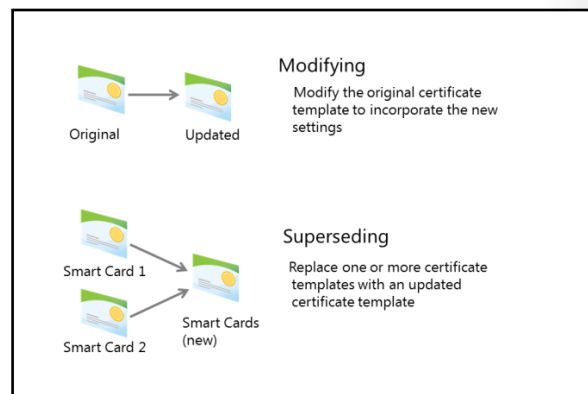
- Version 4: Windows Server 2012 and Windows Server 2012 R2 operating systems provide support for version 4 certificate templates, and for all other, older versions of Windows Server. These certificate templates are available only to the Windows Server 2012 and Windows 8 and newer operating systems. To help administrators remember which operating system versions support which features, the compatibility tab was added to the certificate template properties tab. It marks options as unavailable in the certificate template properties, depending on the selected operating system versions of certificate client and CA. Version 4 certificate templates also support both CSPs and key storage providers, and you can configure them to require renewal with the same key.

You upgrade certificate templates only when you upgrade the CA from Windows Server 2008 or Windows Server 2008 R2 to Windows Server 2012 or Windows Server 2012 R2. After the upgrade, you can upgrade the certificate templates by launching the CA Manager console, and then click Yes to accept the upgrade prompt.

Methods for Modifying Certificate Templates

In most organizations, the CA hierarchy has one certificate template for each function. For example, there may be one certificate template for file encryption and another for code signing. Additionally, there may be templates that cover functions for most of the common subject groups.

As an Information Technology (IT) administrator, you may need to modify an existing certificate template because of incorrect settings or other issues in the original certificate template. Also, you may need to merge multiple existing certificate templates into a single template.



You can update a certificate template by using one of the following methods:

- **Modifying 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.
- **Superseding 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 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. When the user renews the certificate, it will be issued based on a new certificate template that supersedes the old one.

Designing Certificate Template Security

To configure certificate template permissions, you must define the DACL for each certificate template. The permissions assigned to a certificate template will define which users and groups can read, modify, enroll, and autoenroll for that certificate template. You should configure a DACL for each certificate template that you will use, so that you can define the tasks that users or computers can perform on that template.

You can assign the following permissions to certificate templates:

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 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

- **Full Control.** The Full Control permission allows a security principal to modify all attributes of a certificate template, including permissions for the certificate template.
- **Read.** The Read permission allows users and computers to view the certificate template when enrolling for certificates. The certificate server requires the Read permission to find the certificate templates in AD DS.
- **Write.** The Write permission allows users and computers to modify the certificate template's attributes, including the permissions that you assign to the certificate template.
- **Enroll.** The Enroll permission allows users and computers 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 users and computers to receive a certificate through the autoenrollment process. However, the Autoenroll permission also requires that users and computers have both Read and Enroll permissions for a certificate template.

We recommend that you assign certificate template permissions to global or universal groups only, because the certificate-naming context in AD DS stores the certificate template objects. You cannot assign permissions by using domain local groups found within an AD DS domain, and 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 allows all users and computers to view the certificate templates in AD DS, and allows the CA that runs under the system context of a computer account to view the certificate templates when assigning certificates.

Question: To which role will you assign full access rights for one or more certificate templates?

Considerations for Designing Certificate Templates

When you design a certificate template, you set available options on the copy of the existing certificate template. As a result, you get a new certificate template with settings that satisfy your requirements. When designing certificate templates, consider the following:

- Subject name requirements. Certificate Subject Name is a value that normally describes the identity of a user, computer, or service that is enrolling a certificate. Be sure that you enter the proper subject name when enrolling a certificate. This prevents problems with the certificate trust. The Windows Server 2012 CA can issue certificates with subject alternative names (SANs), which allow you to use several subject names on one certificate.
- Certificate life span. Each certificate template defines the life span of certificates that the CA issues based on that template. You cannot issue a certificate with an unlimited life span. Additionally, there is no definite recommendation regarding the duration that you should configure for a template life span. The life span depends on security requirements, certificate usage, and type. Certificate life spans work as a subset of the CA's certificate life span. All certificates, including CA certificates, have expiration dates, after which they are no longer valid. As a result, you cannot issue a certificate with a life span that exceeds the issuing CA's life span.
- Certificate usage. You can issue many certificates, each with a particular purpose, or fewer certificates, each with broad usage. This decision depends on your environment, the desired level of administration, the possible effects on the subjects, and the effects of multiple certificates on the applications that will use them. One strategy for certificate administration is to create a template for each function, such as file encryption or code signing. Then subjects can enroll for each certificate as necessary for the appropriate function. This allows subjects to start with a few certificates, and then obtain only those new certificates that they need. The drawback to this strategy is that the subject may accumulate a large number of certificates and private keys that become difficult to manage.
- The CSP that you will implement. A version 2 certificate template allows you to define one or more CSPs that a template will use. Thus, you can control the types of cryptography that subjects can use within an enterprise. This is useful when security is your priority. Because subjects use the CSP for both portions of any cryptographic service—either encryption and decryption, or signing and confirming signatures—you must ensure that all subjects can use the same CSP. The easiest way to do this is to configure each certificate template to identify one CSP. The administrator should determine the CSP to use for each template, depending on the level of security required, the intended purposes of the certificate, and the presence of security hardware, such as smart cards.
- Key length. You can define a minimum key size allowed for a certificate template. In general, for the same algorithm, longer keys provide more protection than shorter keys. However, larger keys take longer to generate and use. You should select a minimum key size that ensures the necessary amount of protection without affecting performance.

When designing your certificate templates, consider the following:

- Subject name requirements
- Certificate life span
- Certificate usage
- CSP
- Key length
- Deployment methods
- Key archival
- Private key exportability

- Deployment methods. You can configure certificate issuance to subjects in many ways, including manual enrollment, autoenrollment, and CA Web enrollment. Certificate strategies include issuing one all-inclusive certificate to all subjects, or issuing several application-specific certificates to subjects as necessary. Because there are so many options, you should plan your deployment method well in advance of certificate deployment.
- Key archival. When subjects lose their private keys, they will not be able to access any information that was encrypted persistently with the corresponding public key. To prevent this problem, key archival allows you to encrypt and archive keys in the CA database when you issue certificates. If subjects lose their keys, you can retrieve the information from the database and provide the information to them. This enables you to recover the encrypted information rather than lose it.
- Whether users can export the private key. If you want to allow users to export certificates together with their private keys, you should configure this option on the certificate template. When you enable this option, users can use the Certificates Management console to export their certificates, and then move them to another machine.

You can customize certificate templates with several extensions that regulate their use, including:

- Issuance policies. An issuance policy, also known as an enrollment or certificate policy, is a group of administrative rules that you implement when issuing certificates. An object identifier (OID) that you define at the CA represents an issuance policy in a certificate. The issued certificate includes the OID. When a subject presents its certificate, the target can examine it to verify the issuance policy, and then determine if that level of issuance policy is sufficient to perform the requested action.
- Application policies. Application policies allow you to decide which certificates can be used for certain purposes. This enables you to issue certificates widely without worrying that they are being used for unintended purposes. Sometimes application policies are called extended key usage or enhanced key usage. Because some implementations of PKI applications cannot interpret application policies, both application policies and enhanced key usage sections appear in certificates that Windows Server-based CAs issue.
- Key usage. A certificate enables the subject to perform a specific task. To help prevent misuse of certificates, restrictions are placed on certificates automatically. Key usage is a restriction method that administrators use to define certificate usage.
- Basic constraints. Basic constraints ensure that CA certificates are used only in certain applications. For example, a basic constraint may be the path length. A path length defines the number of CAs that are permitted below the current CA. This path length constraint ensures that CAs at the path's end can issue only end entity certificates, not CA certificates.

Lesson 3

Planning and Implementing Certificate Distribution and Revocation

When deploying a PKI in your organization, you must define methods for certificate distribution and revocation. There are several reasons for revoking certificates, such as if a key is compromised or an employee leaves your 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 you can use 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 certificate enrollment options.
- Describe the certificate autoenrollment process.
- Describe considerations for choosing an enrollment method.
- Describe Enrollment Agents.
- Describe options for implementing certificate revocation.
- Describe considerations for designing certificate revocation.
- Describe how to deploy and configure CDP and AIA locations.
- Delegate a CA and certificate management.

Certificate Enrollment Options

In Windows Server 2012, you can use several methods to enroll for a user or computer certificate. The use of these methods depends on the scenario. For example, you will use autoenrollment to mass-deploy certificates to a large number of users or computers, while you will use manual enrollment for certificates dedicated to specific security principals only.

The following list describes the different enrollment methods and when to use them:

- **Autoenrollment.** The administrator defines the permissions and the configuration of a certificate template. These definitions help the requestor to request, retrieve, and renew certificates automatically, without user interaction. Use this method for AD DS domain computers and configure the certificate for autoenrollment through Group Policy.

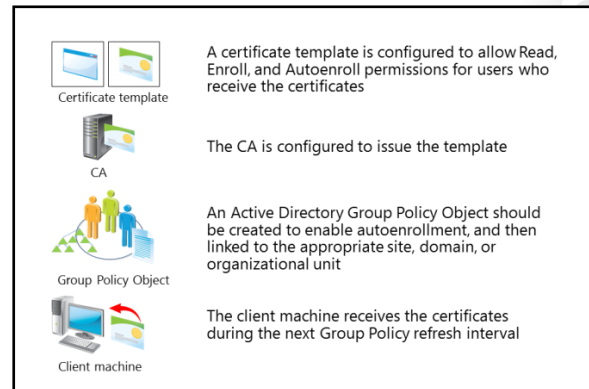
Method	Use
Autoenrollment	<ul style="list-style-type: none"> • To automate the request, retrieval, and storage of certificates for domain-based computers
Manual enrollment	<ul style="list-style-type: none"> • To request certificates by using the Certificates Templates console or Certreq.exe when the requestor cannot communicate directly with the CA
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)

- CA Web enrollment. You can enable a website CA so that users can obtain certificates. To use CA Web enrollment, you must install Internet IIS and the CA 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 CA issues the certificate automatically if the user has the appropriate enrollment permissions. Use this method to issue certificates when you cannot use autoenrollment, such as for an Advanced Certificate request.
- Manual enrollment. A device, such as a web service or a computer, generates the private key and a certificate request. Then the certificate request is transported to the CA to generate the requested certificate. The certificate is 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). A CA administrator creates an Enrollment Agent account for a user, who can then enroll for certificates on behalf of other users. For example, you would use this method if you need to allow a manager to preload new employees' logon certificates on smart cards.

Question: How do you usually enroll for certificates in your environment? Do you use CA Web enrollment? Why or why not?

The Certificate Autoenrollment Process

Autoenrollment is one of the most common methods for deploying certificates in an Active Directory environment. This method automates certificate deployment to users and computers within a PKI. You can use autoenrollment in environments that meet specific requirements, such as the use of certificate templates and Group Policy in AD DS. However, it is important to note that you cannot use autoenrollment with a stand-alone 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 AD CS administrator duplicates a certificate template, and then configures the permissions to allow Read, 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, and periodically refreshes during a user's session. This Group Policy setting is named Certificate Services Client – AutoEnrollment, and it is located under the Computer Configuration node's Security Settings, in the Group Policy Management Editor.

An internal timer triggers autoenrollment of the certificate for the user and computer. 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 signs in.

You can distribute many certificates without the client even being aware that enrollment is taking place. These include most types of certificates that are issued to computers and services, including many certificates issued to users.

To configure automatic enrollment for certificates in a domain environment, you must:

- Have membership in Domain Admins, Enterprise Admins, or an 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 signs in and, if desired, remove them when the user signs off. Additionally, Credential Roaming maintains the integrity of credentials under any conditions, such as when certificates update or when users sign in to more than one computer at a time. This prevents the autoenrollment of the user on each new machine to which he or she signs in.

Credential Roaming triggers whenever 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. Windows 7 and newer operating systems support Credential Roaming.

Considerations for Choosing an Enrollment Method

When choosing the best enrollment method for your clients, consider the following:

- Only enterprise CAs support autoenrollment. It is not possible to configure autoenrollment by using a stand-alone CA, because this functionality requires AD DS.
- Only enterprise CAs support smart card enrollment. If you are going to deploy smart cards in your environment, you should be aware that stand-alone CAs do not support the issuance of smart card certificates.

When planning an enrollment method for certificates, keep the following in mind:

- Only enterprise CAs support autoenrollment
- Autoenrollment is intended for domain clients
- Your clients' operating systems support different enrollment methods
- Policies that you establish to manage certificate distribution

- In Windows XP, Windows Server 2003, and all newer versions, autoenrollment is available for users, computers, and smart cards.
- Autoenrollment is available for domain clients only.

To select the certificate enrollment and renewal processes that are appropriate for your organization, you should consider several factors carefully, including:


- The users, computers, and services for which you intend to provide services. Determine whether they are internal or external to the organization. Identify the operating systems that they are running, and then determine whether they are connected to AD DS.
- The operating system that your clients are using. Clients running Windows Server 2003, Windows XP, Windows Vista®, Windows 7, and Windows 8 can use the Certificate Request Wizard, autoenrollment, or the smart card enrollment station. Windows 2000 Server supports the Certificate Request Wizard but not smart card enrollment. Autoenrollment and the smart card enrollment station also require AD DS. Most other clients can use their web browsers to access web-based enrollment and renewal services.

- The policies that you establish to manage certificate distribution. This includes both the procedural policies that you establish for your PKI and the Group Policy settings that you use to implement those policies. For example, you might want to configure Microsoft Forefront® Identity Manager to establish workflows for certificate enrollment.
- The type of CA that issues the certificates. For example, you must have an enterprise CA to use the smart card enrollment station. Additionally, stand-alone CAs support only web or manual enrollment, whereas enterprise CAs support web and manual enrollment, in addition to autoenrollment.

Enrollment Agent Overview

In the Windows Server 2012 CA, it is possible to configure certificate enrollment on behalf of another user. To do this, you must have a specific certificate issued. This certificate is based on the *Enrollment Agent* template. When a user receives a certificate based on an Enrollment Agent template, he or she has the ability to enroll for a certificate on behalf of another user. Unlike a certificate manager, an Enrollment Agent can only process the enrollment request and cannot approve pending requests or revoke issued certificates.

- An Enrollment Agent is a user who has the appropriate certificate assigned and has the ability to request certificates on behalf of other users or computers
- The restricted Enrollment Agent has limited permissions:
 - For specific group of users
 - For specific certificate templates
- Requires Windows Server 2008 Enterprise edition or Windows Server 2012 CA

 **Note:** An Enrollment Agent is a certificate with a very high security risk. A person that has an Enrollment Agent certificate can impersonate other users, because he or she is able to issue a certificate for other users—for example, smart card certificates. Ensure that you protect this certificate template.

Windows Server 2012 includes three certificate templates that enable different types of Enrollment Agents:

- Enrollment Agent. This template is used to request certificates on behalf of another subject.
- Enrollment Agent (Computer). This template is used to request certificates on behalf of another computer subject.
- Exchange Enrollment Agent (Offline Request). This template is used to request certificates on behalf of another subject and supply the subject name in the request. The Network Device Enrollment Service uses this template for its Enrollment Agent certificate.

When you create an Enrollment Agent, you can further refine the agent's ability to enroll for certificates on behalf of others by a group and by a certificate template. For example, you might want to implement a restriction that the Enrollment Agent can enroll for smart card logon certificates only and just for users in a certain office or organizational unit (OU) that is the basis for a security group.


In older versions of Windows Server CA, it was not possible to permit an Enrollment Agent to enroll only a certain group of users. As a result, every user with an Enrollment Agent certificate was able to enroll on behalf of any user in an organization.

The restricted Enrollment Agent is functionality that was introduced in the Windows Server 2008 Enterprise edition operating system. This functionality allows you to limit the permissions for users who are designated as Enrollment Agents in enrolling smart card certificates on behalf of other users.

Typically, one or more authorized individuals within an organization are designated as Enrollment Agents. The Enrollment Agent needs to be issued an Enrollment Agent certificate, which enables the agent to enroll for certificates on behalf of users. Enrollment Agents are typically members of corporate security,

IT security, or help desk teams, because these individuals already have been entrusted with safeguarding valuable resources. In some organizations, such as banks that have many branches, help desk and security workers might not be conveniently located to perform this task. In this case, designating a branch manager or another trusted employee to act as an Enrollment Agent is required to enable smart card credentials to be issued from multiple locations.

On a Windows Server 2012 CA, the restricted Enrollment Agent features allow an Enrollment Agent to be used for one or many certificate templates. For each certificate template, you can choose which users or security groups the Enrollment Agent can enroll. You cannot constrain an Enrollment Agent based on a certain Active Directory OU or container; instead, you must use security groups.

 **Note:** Using restricted Enrollment Agents will affect the performance of the CA. To optimize performance, you should minimize the number of accounts that are listed as Enrollment Agents. To do this, you minimize the number of accounts in the Enrollment Agent's certificate templates access control list. As a best practice, use group accounts in both lists instead of individual user accounts.

Options for Implementing Certificate Revocation

During the certificate management process, you may need to revoke certificates for several reasons. For example, a key may be compromised or a user may leave the organization. You need to ensure that network clients can determine which certificates are revoked before they accept authentication requests.

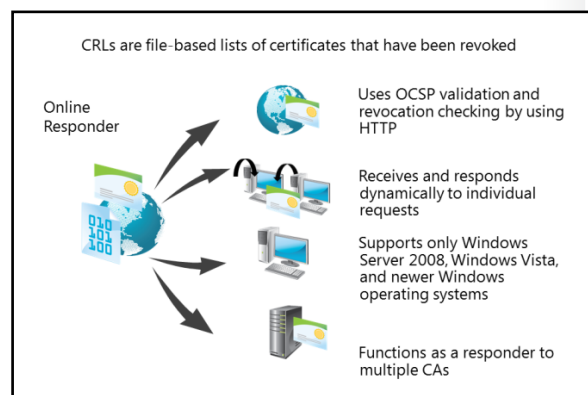
Each time that you revoke a certificate, you should specify a revocation reason. Windows Server 2012 CAs allow you to select a reason code from a list of predefined codes. We recommend that you use these codes every time you revoke a certificate, so that you can track that activity.

When you revoke a certificate, you must announce that the certificate was revoked, so that no services will use or accept that certificate. You can do this by using CRLs or configuring the Online Responder service based on OCSP.

CRLs

CRLs are lists of certificates that have been revoked. The CAs maintain these lists as part of the certificate database. CRLs provide clients with one method of checking certificate revocation before accepting a certificate and proceeding with secure communication.

Clients use CRL distribution points as references to locate up-to-date CRLs. Each CRL distribution point defines one point where clients can access the CRLs, such as on a file share, a Lightweight Directory Access Protocol (LDAP) location, or an HTTP location. However, not all of these locations may be accessible or published. You can choose to configure locations to which the CA publishes CRLs and delta CRLs. You can publish base CRLs periodically.



If the CA issues and revokes many certificates, you might have to publish a large base CRL. A base CRL is a CRL that contains all revoked certificates. To avoid publishing large base CRLs, you can publish delta CRLs. Delta CRLs are smaller, interim CRLs that contain only the certificates revoked since the last base CRL was published. Clients can retrieve the delta CRL, and then quickly build a complete list of revoked certificates. You can transfer delta CRLs faster than base CRLs. Also, you can use the delta CRLs to publish revocation data frequently. Computers that run the Windows operating system can also use delta CRLs.

You can configure a CRL publication setting and a CRL publish period. The CRL publish period defines when a CA must publish an updated CRL automatically. When you first install a CA, the default CRL publish period is one week. You can configure the CRL publish period to suit your needs by opening the properties of the Revoked Certificates container in the CA Management console.

Online Responder

By using OCSP, Online Responders provide clients with an efficient method for determining a certificate's revocation status. OCSP submits certificate status requests by using HTTP.

Clients access CRLs to determine the revocation status of a certificate. If CRLs are large, clients might spend a long time searching through them. An Online Responder can search the CRLs for the clients, and then respond to the requested certificate only. Online Responders receive all the certificate revocation data instead of the relying clients. A relying party submits a status request about an individual certificate to an Online Responder, which returns a definitive, digitally signed response that indicates only that certificate's status. The amount of data retrieved per request is constant, no matter how many revoked certificates exist in the certificate database on the CA.

You can use a single Online Responder to determine revocation status information for certificates that either a single CA issues or that multiple CAs issue. 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 R2 Enterprise, Windows Server 2008 R2 Datacenter, Windows Server 2012, or Windows Server 2012 R2. It can work with a CA that runs Windows Server 2003, or a newer operating system. You should install the Online Responder and CA on different computers.



Note: You must configure the CRL and OCSP settings on the CA before issuing certificates. The location of these services is included in each certificate that the CA issues. Certificates issued before you configure OCSP will not contain the correct information.

For scalability and high availability, you can deploy the Online Responder on a single computer or on a load-balanced group, which contains one or more computers. Additionally, you can configure arrays of multiple linked computers that host Online Responders, and then process 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 Online Responder's URL in the AIA extension of issued certificates. The OCSP client uses this URL to validate the certificate status. Also, you must issue the OCSP Response Signing certificate template so that the Online Responder can enroll that certificate.

Considerations for Designing Certificate Revocation

When designing a certificate revocation infrastructure, first you must determine whether you will use CRLs, Online Responders, or both. Before choosing a revocation technology, you should be aware of differences between the two.

The most important characteristics of CRLs are that:

- A CRL is a file that a CA creates and signs.
- CRLs contain serial numbers of certificates that a CA issues and that you then revoke.
- The CRL also contains the revocation reason for each certificate and the time that the certificate was revoked.
- By default, the CRL is published in systemroot\system32\CertSrv\CertEnroll\.
- If the computer is a domain member and has permission to write to AD DS, the CRL is published to AD DS also.
- The publishing period for a CRL is not the same as the validity period for a CRL. By default, the validity period of a CRL exceeds the publishing period of a CRL by 10 percent, up to a 12-hour maximum, to allow for directory replication.

Potential difficulties with using CRLs include:

- Potentially large file size. This could limit scalability.
- Bandwidth and storage overhead. Large file sizes on the server or client sides may affect this adversely.
- CA processing capacity. High-frequency publication may affect server performance negatively.
- Latency. There is a time lag between when a certificate is revoked and when that information is available.

The most important characteristics of OCSP are:

- OCSP is an HTTP protocol.
- OCSP responders receive certificate status requests.
- The responder's response is signed digitally and indicates the certificate status.
- The amount of data retrieved per request is constant regardless of the number of revoked certificates in the CA.

Most OCSP responders receive their data from published CRLs. Therefore, they rely on the CA's publishing frequency. Some OCSP responders can receive data directly from the CA's certificate status database, so they can provide almost real-time status.

AD DS deploys CRLs on each CA, by default. You can add additional CRL distribution point locations, but each CA will publish CRLs to a share folder and also to AD DS, if you use an enterprise CA.

When designing certificate revocation, follow these guidelines:

- Evaluate the potential benefits of supplementing CRLs with the use of Online Responders
- Identify potential locations where Online Responders would be beneficial
- Identify the installation configuration that best suits your organization
- Identify the locations for every Online Responder and determine how you will manage them
- Test the Online Responder and PKI configuration

However, if you decide to deploy an Online Responder, follow these guidelines:

- Evaluate the potential benefits of supplementing CRLs with the use of Online Responders to manage your organization's revocation checking.
- Identify potential locations where Online Responders would be beneficial.
- Identify the installation configuration that best suits your organization, depending on the number of CAs and locations that you are supporting, the volume of certificate validation requests that you anticipate, and network conditions between your CAs and locations.
- Identify the locations for every Online Responder and determine how you will manage them.
- Test the Online Responder and PKI configuration in a lab environment to validate the PKI design and to identify configuration options for each Online Responder and revocation configuration.

Deploying and Configuring CDP and AIA locations

When you manage and issue certificates, it is important that you properly configure the certificate extensions that verify the certificate of the CA and the certificate that is used by the user, computer, or device. These extensions—AIA and CDP—are part of each certificate. They must point to proper locations, or the PKI might not function correctly.



Note: One common cause for CA hierarchy malfunctions and downtime is improperly configured AIA and CDP extensions. Make sure that you configure these options properly before you put your CA hierarchy in production.

- The AIA specifies where to retrieve the CA's certificate
- The CDP specifies from where the CRL for a CA can be retrieved
- Publication locations for AIA and CDP:
 - AD DS
 - Web servers
 - FTP servers
 - File servers
- Ensure that you properly configure CRL and AIA locations for offline and stand-alone CAs
- Ensure that the CRL for an offline root CA does not expire

What Is AIA?

AIA addresses are the URLs in the certificates that a CA issues. These addresses inform the verifier that a certificate exists and from where to retrieve it. AIA URLs can be HTTP, File Transfer Protocol (FTP), LDAP, or file location addresses.

What Is CDP?

The CDP is a certificate extension that indicates from where the CRL for a CA can be retrieved. It can contain none, one, or many HTTP, FTP, FILE, or LDAP addresses.

Each certificate that you issue from your CA contains information about the CDP and AIA location. Each time a certificate is used, these locations are checked. The AIA location is checked to verify the validity of the CA certificate, while the CDP location is checked to verify content of the CRL for that CA. At least one AIA and one CDP location must be available for each certificate. If they are not available, the system will presume that the certificate is not valid, the revocation check will fail, and you will not be able to use that certificate for any purpose.

AIA and CDP Publishing

If you only use an online CA, these values are configured by default locally on the CA. However, if you want to deploy an offline root CA, or if you want to publish AIA and CDP to an Internet-facing location, you must reconfigure these values so that they apply to all the certificates issued by the root CA. The AIA and CDP extensions define where client applications can locate AIA and CDP information for the root CA.

The formatting and publishing of AIA and CDP extension URLs are generally the same for root CAs and subordinate CAs. You can publish the root CA certificate and the CRL to the following locations:

- AD DS
- Web servers
- FTP servers
- File servers

Publication Points

To ensure accessibility to all computers in the forest, publish the offline root CA certificate and the offline root CA's CRL to AD DS by using the `certutil` command. This places the root CA certificate and the CRL in the Configuration naming context, which is then replicated to all domain controllers in the forest.

For computers that are not members of an AD DS domain, place the CA certificate and the CRL on web servers by using the HTTP protocol. Locate the web servers on the internal network, and on the external network if external client computers—or the internal clients from the external networks—require access. This is very important if you use internally issued certificates outside of your organization.

You also can publish certificates and CRLs to the `ftp://` and `file://` URLs, but we recommend that you use only the LDAP and HTTP URLs because they are the most widely supported URL formats for interoperability purposes. The order in which you list the CDP and AIA extensions is important, because the certificate-chaining engine searches the URLs sequentially. If your certificates mostly are used internally in an AD DS domain, place the LDAP URL first in the list.



Note: Besides configuring CDP and AIA publication points, you also should make sure that the CRL is valid. An online CA will automatically renew the CRL periodically, but an offline root CA will not. If the offline root CA CRL expires, revocation check will fail. To prevent failure, make sure that you configure the validity period for the offline root CA CRL to be long enough, and set a reminder to turn that CA on and issue a new CRL before the old one expires.

Delegating CA and Certificate Management

By configuring security options on the CA properties, you can delegate management of the CA and certificates. As a best practice, configure these options so that you dedicate personnel to managing the CA and the certificates. By default, the Administrators, Domain Admins, and Enterprise Admins security groups can manage CAs and issue and manage certificates, and the Authenticated Users group can request certificates from a CA.

Role-based administration in AD CS provides the ability to delegate predefined permissions to users or groups based on built-in CA roles. Each role can perform a predetermined task or series of tasks. The following table provides details about the roles and groups involved in role-based administration.

- By configuring the security options of the CA, you can implement role separation
- By default, Administrators, Domain Admins, and Enterprise Admins groups are allowed to administer CAs and certificates
- You can configure the following permissions:
 - Read
 - Issue and Manage Certificates
 - Manage CA
 - Request Certificates
- Best practices:
 - Group accounts that you have assigned a CA administrator or certificate manager role should not be members of the local Administrators group
 - Assign CA roles to group accounts only

Role/group	Purpose	Information
CA administrator	Manage the CA	Assigned by using the CA console
Certificate manager	Issue and manage certificates	Assigned by using the CA console
Backup operator	Backup and restore files and directories	Operating system role
Auditor	Manage auditing and Security Event log	Operating system role
Enrollees	Read and enroll	Can request certificates

Role-based administration combines operating system roles and AD CS roles to provide a complete, segmented management solution for your CAs. Instead of assigning local administrative privileges to the IT personnel who manage the CA, you can assign roles, which ensure that administrators have the minimum permissions necessary to perform their jobs.

Role-based administration also reduces the administrative overhead of granting rights to administrators because the process involves adding a user to a group or role.

If you want to implement role separation for the CA and certificate management, you should sign in to the CA as a member of the Administrators, Domain Admins, or Enterprise Admins group, and configure security options for the CA. By implementing role separation, you are delegating the rights for CA and certificate management to other users who are not members of administrative groups.

To configure delegation options for the CA and certificates, right-click your CA server, and then select Properties. On the security tab, you can configure the following permissions:

- **Read.** This permission allows a user or security group to see all CA options, but does not allow any modification of settings. You can assign this permission to a person or a group that performs auditing of CA settings.
- **Issue and Manage Certificates.** This permission allows a user or security group to approve pending certificate enrollment and revocation requests. Sometimes this role is known as a CA officer or certificate manager. A person or a group with this permission cannot modify any CA setting.
- **Manage CA.** This permission allows a user or security group to configure and maintain the CA. This includes the ability to assign all other CA roles and renew the CA certificate. However, a person with this role cannot manage certificates. Typically, this role is known as a CA administrator.
- **Request certificates.** This permission allows users, computers, or groups to request certificates from the CA. It does not allow certificate enrollment, but does allow the sending of certificate requests to the CA. You can configure enroll permissions on a specific certificate template. Unless you want to specify that only certain users or computers can request certificates from a specific CA, we recommend that you designate Authenticated Users as the group with this permission.

As a best practice, group accounts that you have assigned a CA administrator or certificate manager role should not be members of the local Administrators group. Additionally, you should assign CA roles to group accounts only and not to individual user accounts. However, membership in the local Administrators group on the CA is required to renew a CA certificate. Members of this group can assume administrative authority over all other CA roles.

Lesson 4

Planning and Implementing Key Archival and Recovery

Certificate or key recovery is one of the most important management tasks you perform during the certificate life cycle. You use a key archival and recovery agent for data recovery if you lose your public and private keys. Also, you can use automatic or manual key archival and key recovery methods to ensure that you can gain access to data if you lose your keys. In this lesson, you will learn how to plan and implement key archival and recovery in AD CS in Windows Server 2012.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe private key recovery scenarios.
- Describe options for configuring key archival.
- Describe options for configuring key recovery.
- Describe considerations for managing key archival and recovery.

Scenarios for Private Key Recovery

When you deploy PKI in your environment, you must be aware of the importance of the private key in the certificate. Each certificate is associated with a key pair, which consists of a public key and a private key. While the public key is accessible without any limitations, the private key is accessible only to a certificate owner, as it is stored in the user profile by default. A private key is especially important if you use certificates for encryption. When you encrypt data, you need your private key to decrypt it. If you lose your private key, there is no other way to decrypt data.

Therefore, you should plan and define strategies for private key protection and recovery.

Many situations can result in the loss of certificate private keys and the need for key recovery. For example, if a user profile is corrupted or deleted, a user will not be able to access his or her private key. If a hard disk fails and you reinstall the operating system on another drive, the private key will be inaccessible. Additionally, a user can delete the certificate inadvertently from his or her personal store, which deletes the private key. Furthermore, if someone steals your computer or you lose it, you will lose access to your private key. In each case, you need to recover your private key. In some cases, a certificate owner can export the private key with the certificate, and copy it to a backup location. In this scenario, recovery is very simple: the user imports the certificate. However, some certificates do not allow private key exporting, so other technologies are necessary for key backup and recovery.

In an AD DS environment, we recommend that you have a centralized method for private key archival and a precisely defined procedure for key recovery. The next topics detail the methods for key archival and key recovery.

You may lose the key pair when:

- A user profile is deleted or corrupted
- An operating system is reinstalled
- A disk is corrupted
- A computer is stolen

As a best practice, have a centralized method and technology for private key backup and recovery

Options for Configuring Key Archival

Conducting a manual backup is the easiest method of key archival. When you enable certificates with a private key export, each user can export the certificate manually with the private key (in .pfx format) and copy it to a safe location, such as a USB device or cloud-based storage. When exporting a private key, it is mandatory that you protect it with a password, which prevents a user from importing the certificate without providing the password.

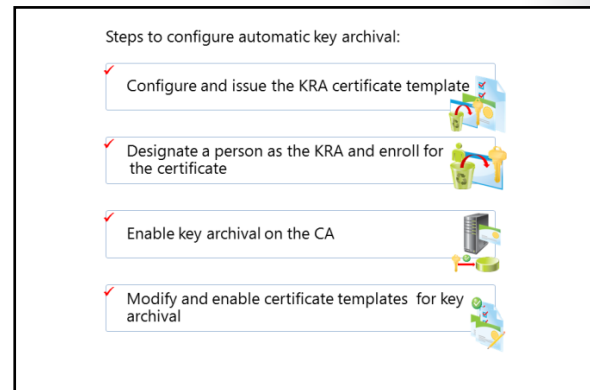
However, manual certificate backup is often unreliable. Some users will forget to back up their certificates or they will lose their backup. Furthermore, some certificate templates do not allow users to perform exports of the private key. This means that you should perform key backup in another way.

AD CS in Windows Server 2012 allows you to configure CAs and certificate templates for private key archival. By using this option, you can back up private keys in the CA database for certificates that are issued based on specific certificate templates. If the private key for a user is lost, you can restore it from the CA database by following a procedure for retrieving an archived key from the CA database. Private key archival must be enabled on the CA, as it is disabled by default. Note that private key archival will be performed only for certificates issued after private key archival has been enabled.

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.

To perform the automatic key archival process, follow these steps:

1. Configure the Key Recovery Agent (KRA) certificate template. Only Enterprise Administrators or Domain Admins are allowed to request a KRA certificate. If you want to enroll another user with a KRA certificate, you must specify it on the template DACL.
2. Configure certificate managers. On the CA, you can define a person to be a certificate manager. Usually, the certificate manager 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.
3. Enable KRA:
 - Log on as an administrator of the server or as a CA administrator, if role separation is enabled.
 - In the Certificate Authority Management console, right-click the CA name, and then click **Properties**. To enable key archival, on the **Recovery Agents** tab, click **Archive the key**.
 - By default, the CA uses one KRA. However, you must first choose the KRA certificate for the CA to begin archival. Do this by clicking **Add**.
 - The system finds the valid KRA certificates, and then displays the available KRA certificates. Typically, an enterprise CA publishes these to AD DS during enrollment. KRA certificates are stored under the KRA container in the Public Key Services branch of the configuration partition in AD DS. Because a CA issues multiple KRA certificates, each KRA certificate will be added to the multivalued user attribute of the CA object.
 - Choose the intended certificate, and then click **OK**.
 - After you have added one or more KRA certificates, click **OK**. The KRA certificates are processed at service start.

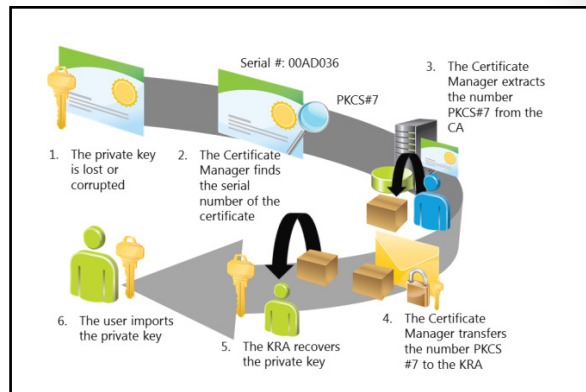


4. Configure certificate templates:
 - In the Certificate Templates Microsoft Management Console (MMC), right-click the certificate template for which you want to enable private key archival, and then click **Properties**.
 - To always enforce key archival for the CA, in the **Properties** dialog box, on the **Request Handling** tab, select the **Archive subject's encryption private key** check box. Also, select the **Use advanced symmetric algorithm to send the key to the CA** option.

Options for Configuring Key Recovery

Archiving a private key by using a manual method is a simple procedure. The user locates the private key backup, which is part of the certificate, and then imports it to his or her user profile. However, if you perform key archival by using options on the CA and certificate templates, follow this procedure to recover the private key:

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 Public Key Cryptography Standards (PKCS) #7 binary large object (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, and then 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 it 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 that must be recovered. Additionally, the holder generates a password-protected .pfx file that contains the certificate and private key.
4. Import recovered keys. The password-protected .pfx file is delivered to the end user. This user imports the .pfx file into the local user certificate store. Alternatively, the KRA or an administrator can perform this part of the procedure on behalf of the user.



Managing Key Archival and Recovery

Certificate or key recovery is a very important management task. You use a key archival and recovery agent for data recovery when you lose your public and private keys. Also, you can use automatic or manual key archival and key recovery methods to ensure that you can access data if keys are lost.

You use key archival and KRAs for data recovery. The only way you can ensure that CA administrators can recover private keys is to archive them. KRAs can retrieve the original certificate, private key, and public key that were used to encrypt the data from the CA database.

Key recovery implies that you can archive and recover the private key portion of a public-private key pair. Private key recovery does not recover any data or messages. It merely enables a user to retrieve lost or damaged keys and an administrator to assume the role of a user for data access or data recovery purposes. In many applications, you cannot recover data without first performing key recovery.

After you configure a CA to issue a KRA certificate, any user with Read and Enroll permissions on the template for the KRA certificate can enroll and become a KRA. As a result, members of the Domain Admins and Enterprise Admins groups have permission by default. Therefore, ensure that you:

- Allow only trusted users to enroll for this certificate.
- Store the KRA's recovery key in a secure location.
- Secure the server where you archive the keys.

When planning and managing key archival and recovery, consider the following:

- Private key recovery does not recover any data or messages
- You should disable the template for the KRA certificate immediately after issuing a KRA certificate
- You should store the KRA certificate in a secure location
- You should secure the the server where keys are archived

Lab: Planning and Implementing an Active Directory Certificate Services Infrastructure

Scenario

The A. Datum Corporation has expanded and so have its security requirements. The security department at A. Datum wants to enable secure access to critical web sites and provide additional security for features such as smart cards and the Windows 7 and Windows 8 DirectAccess features. To address these and other security requirements, A. Datum has decided to implement a PKI by using the AD CS role in Windows Server 2012 R2.

Objectives

After completing this lab, you will be able to:

- Plan the AD CS deployment.
- Deploy the CA infrastructure.
- Implement the certificate templates.
- Implement certificate revocation and distribution.

Lab Setup

Estimated Time: 90 minutes

Virtual machines	20414C-LON-HOST1 20414C-LON-DC1 20414C-LON-SVR1 20414C-LON-CA1 20414C-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 20414C-LON-HOST1 virtual machine, start **Hyper-V Manager**.
2. In Microsoft Hyper-V® Manager, click **20414C-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 two and three for 20414C-LON-SVR1, 20414C-LON-CA1, and 20414C-LON-CL1. Do not log on until instructed to do so.

Exercise 1: Planning an Active Directory Certificate Services Deployment

Scenario

The security department at A. Datum has provided the following requirements for its AD CS deployment:

- The root CA at A. Datum must remain completely isolated from the network at all times.
- Because of the importance that certificates will have in the planned deployment, it is critical that the CA server role be highly available. Users in London and Toronto must be able to obtain certificates and check certificate revocation in the event of a single server or single wide area network (WAN) failure.
- A. Datum has several applications that require the validation of a user's identity through certificates. Only internal employees use the internally developed applications. In addition, because managers use high-privilege accounts, they want to deploy smart cards to managers and force them to use only smart cards when signing on.
- A. Datum has several web servers that run applications for external clients. Communication between server and clients must be secured with SSL. You should configure the web server certificates to expire three years after deployment.
- A. Datum is planning to implement DirectAccess for all users with laptop computers so that they can connect to the internal network without using a virtual private network (VPN). They need to ensure that every computer running the Windows 7 operating system has a computer certificate, and laptop computers should be issued DirectAccess certificates. A. Datum wants each computer's certificate to be valid for six months only.
- The revocation status for all certificates must be available to all computers connected to the internal network in London, Toronto, and Sydney. A. Datum wants to minimize the network traffic necessary to retrieve the certificate revocation information.
- As part of the DirectAccess deployment, the certificate revocation information for the certificates issued to the DirectAccess servers must be available to clients on the Internet. All DirectAccess clients will connect to the London data center. The publicly accessible location is the web server in London, with the public URL www.adatum.com.

The main tasks for this exercise are as follows:

1. Read the supporting documentation
2. Propose a solution and plan a course of action
3. Examine the suggested proposals in the Lab Answer Key

► Task 1: Read the supporting documentation

Read the documentation provided.

► Task 2: Propose a solution and plan a course of action

Based on the lab scenario, propose a solution for a PKI design. Use the following questions as guidance for your PKI design development:

1. How will you address the requirement that the root CA at A. Datum must remain completely isolated from the network at all times?
2. How will you make the CA role highly available?
3. How will you achieve certificate-based authentication for managers and applications?
4. How will you configure a web server certificate to address requirements?

5. How will you address the requirement that only laptops receive certificates for deployment of DirectAccess?
6. How will you address the requirement for revocation checking between sites?
7. How will you address the requirement for Internet-based revocation checking?

► **Task 3: Examine the suggested proposals in the Lab Answer Key**

Compare your solution with the solution that the Lab Answer Key provides. Discuss alternative solutions with the class and instructor.

Results: After completing this exercise, students will have a plan for Active Directory Certificate Services (AD CS) deployment.

Exercise 2: Deploying a CA Infrastructure

Scenario

The first step in deploying an AD CS infrastructure at A. Datum is to deploy the CA infrastructure. To comply with the security requirement that the root CA not be accessible from the network, first you must deploy a stand-alone root CA. Then you will deploy a subordinate enterprise root CA, which will be used to issue the certificates.

The main tasks for this exercise are as follows:

1. Install the Active Directory Certificate Services role on a stand-alone server
2. Configure certificate revocation for the subordinate CA
3. Publish the Root CA certificate in the domain
4. Configure an AD CS role on the subordinate enterprise CA
5. Configure the certificates to enable the subordinate trust

► **Task 1: Install the Active Directory Certificate Services role on a stand-alone server**

1. Sign in 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 **Configure Active Directory Certificate Services on the destination server**.
4. Configure the AD CS role as a stand-alone root CA. Name it **AdatumRootCA**.
5. Set the key length to **4096**, and then accept all other values as default.

► **Task 2: Configure certificate revocation for the subordinate CA**

1. On LON-CA1, open the Certification Authority console.
2. Open the **Properties** window for **AdatumRootCA**.
3. Configure the new location for CDP to be **http://lon-svr1.adatum.com/CertData/<CaName><CRLNameSuffix><DeltaCRLAllowed>.crl**.
4. Select the **Include in the CDP extension of issued certificates** and **Include in CRLs. Clients use this to find Delta CRL locations** options.
5. Configure the new location for AIA to be **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 then copy the RootCA.cer file to `\\lon-svr1\C$`.
9. Copy the contents of the `C:\Windows\System32\CertSrv\CertEnroll` to the `\\lon-svr1\C$` folder.

► **Task 3: Publish the Root CA certificate in the domain**

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 the Trusted Root Certification Authorities store in Computer Configuration\Policies\Windows Settings\Security Settings\Public Key Policies.

► **Task 4: Configure an AD CS role on the subordinate enterprise CA**

1. Sign in to LON-SVR1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. Install the **Active Directory Certificate Services** role on LON-SVR1. Include the **Certification Authority** and **Certification Authority Web Enrollment** role services.
3. After installation is successful, click **Configure Active Directory Certificate Services on the destination server**.
4. Select the **Certification Authority** and **Certification Authority Web Enrollment** role services.
5. Configure LON-SVR1 to be an **Enterprise CA**.
6. Configure the CA Type to be a **Subordinate CA**.
7. For the CA Name, type **Adatum-IssuingCA**.
8. Save the request file to the local drive.

► **Task 5: Configure the certificates to enable the subordinate trust**

1. On LON-SVR1, install the RootCA.cer certificate in the Trusted Root Certification Authorities store.
2. Navigate to **Local Disk (C:)**, and then copy the **AdatumRootCA.crl** and **LON-CA1_AdatumRootCA.crt** files to `C:\inetpub\wwwroot\CertData`.
3. Copy the **LON-SVR1.Adatum.com_Adatum-LON-SVR1-CA.req** request file to `\\lon-ca1\C$`.
4. Switch to LON-CA1.
5. From the Certification Authority console on LON-CA1, submit a new certificate request, by using the .req file that you copied in step three.
6. Issue the certificate, and then export it to p7b format with the complete chain. Save the file to `\\lon-svr1\C$\SubCA.p7b`.
7. Switch to LON-SVR1.
8. Install the **SubCA** certificate on LON-SVR1 by using the Certification Authority console.
9. Start the service.

Results: After completing this exercise, students will have deployed a CA infrastructure.

Exercise 3: Implementing Certificate Templates

Scenario

After deploying the CAs, the next step in deploying an AD CS infrastructure is to configure the certificate templates to meet the security requirements.

The main tasks for this exercise are as follows:

1. Configuring OUs and groups for laptop computers
2. Configure the web server certificate template
3. Configure the DirectAccess certificate template
4. Configure the smart card certificate
5. Issue the certificate templates

► Task 1: Configuring OUs and groups for laptop computers

1. On LON-DC1, open **Active Directory Users and Computers**.
2. In the Active Directory Users and Computer console, create a new organizational unit (OU) called **Laptops**.
3. In the **Laptops** OU, create a security group called **Laptops**.
4. Move the **LON-CL1** machine to the **Laptops** OU.
5. Add **LON-CL1** to the **Laptops** security group.

► Task 2: Configure the web server certificate 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 then name it **Adatum Web Server Certificate**.
4. Configure the validity as **3 years**.
5. Configure the private key as exportable.

► Task 3: Configure the DirectAccess certificate template

1. On LON-SVR1, in the Certificate Templates console, duplicate the **Computer** template.
2. Create a new template, and then name it **DirectAccess Clients**.
3. Configure the validity as **6 months**.
4. Configure the security options on the template so that only members of the Laptops group can read, enroll, and autoenroll for this certificate.

► Task 4: Configure the smart card certificate

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 the **Include e-mail name in subject name** and the **E-mail name** check boxes.
4. Add **Smart Card Logon** to the Application Policies of the new certificate template.
5. Configure this new template to supersede the **User** template.

6. Allow Authenticated Users to **Read, Enroll**, and **Autoenroll** for this certificate.
7. Close the Certificate Templates console.

► **Task 5: Issue the certificate templates**

- Configure LON-SVR1 to issue certificates based on the **Adatum Smart Card User, DirectAccess Clients** and **Adatum Web Server** templates.

Results: After completing this exercise, students will have configured certificate templates to meet security requirements.

Exercise 4: Implementing Certificate Revocation and Distribution

Scenario

After you deploy the certificate templates, you must implement the plan for issuing and revoking certificates. Because your organization's security requirements include several different options for deploying the certificates, you will need to configure web-based certificate enrollment and autoenrollment.

The main tasks for this exercise are as follows:

1. Configure certificate revocation checking
2. Configure certificate revocation checking for DirectAccess clients
3. Configure autoenrollment for user and computer certificates
4. Validate certificate enrollment methods for smart card and DirectAccess certificates
5. Prepare for the next module

► **Task 1: Configure certificate revocation checking**

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 indicates that the installation was successful, 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 **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 on LON-SVR1.
7. Add revocation configuration for **Adatum-IssuingCA**.
8. Enroll for an OCSP Response signing certificate.
9. Ensure that the revocation configuration is listed as working.

► **Task 2: Configure certificate revocation checking for DirectAccess clients**

1. On LON-SVR1, open the **Certification Authority** console.
2. Open the **Properties** window for **Adatum-IssuingCA**.
3. Configure the new location for CDP to be **http://www.adatum.com/CertData/<CaName><CRLNameSuffix><DeltaCRLAllowed>.crl**.

4. Select the **Include in the CDP extensions of issued certificates** and **Include in CRLs. Clients use this to find Delta CRL locations** options.
5. Configure the new location for AIA to be **http://www.adatum.com/CertData/<ServerDNSName>_<CaName><CertificateName>.crt**.
6. Select the **Include in the AIA extension of issued certificates** check box.
7. Restart AD CS services when prompted.

► **Task 3: Configure autoenrollment for user and computer certificates**

1. On LON-DC1, open the **Group Policy Management** console.
2. Create a new Group Policy Object (GPO) called **DirectAccessCert**, and then link it to the Laptops OU.
3. Open the **DirectAccessCert** GPO in Group Policy Management Editor.
4. Navigate to the **Computer Configuration** node, expand **Policies**, expand **Windows Settings**, expand **Security Settings**, click **Public Key Policies**, and then double-click **Certificate Services Client – Auto-Enrollment**.
5. Enable the following configuration options:
 - Renew expired certificates, update pending certificates, and remove revoked certificates
 - Update certificates that use certificate templates
6. Close the Group Policy Management Editor.
7. Create a new GPO called **SmartCardCert**, and then link it to the Managers OU.
8. Open the **SmartCardCert** GPO in **Group Policy Management Editor**.
9. Navigate to the **User Configuration** node, expand **Policies**, expand **Windows Settings**, expand **Security Settings**, click **Public Key Policies**, and then double-click **Certificate Services Client – Auto-Enrollment**.
10. Enable the following configuration options:
 - Renew expired certificates, update pending certificates, and remove revoked certificates
 - Update certificates that use certificate templates
11. Close the **Group Policy Management Editor** and **Group Policy Management** console.

► **Task 4: Validate certificate enrollment methods for smart card and DirectAccess certificates**

1. Sign in to LON-CL1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. On LON-CL1, refresh Group Policy by opening a command prompt, and then typing **gpupdate /force**.
3. Create a new Microsoft Management Console (MMC), and then add the **Certificates** snap-in. Manage certificates for the Computer Account.
4. Verify that the DirectAccess certificate is issued to LON-CL1, and verify that the template used for the certificate is DirectAccess Clients.
5. Sign out from LON-CL1.
6. Sign in as **Adatum\Aidan** with the password **Pa\$\$w0rd**.
7. On LON-CL1, refresh Group Policy by opening a command prompt, and then typing **gpupdate /force**.

8. Create a new MMC, and then add the **Certificates** snap-in.
9. Verify that the smart card certificate is issued to Aidan Delaney, and verify that the template used for the certificate is Adatum Smart Card User.

► **Task 5: Prepare for the next module**

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

1. On the host computer, start Microsoft Hyper-V® Manager.
2. On the **Virtual Machines** list, right-click **20414C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps two and three for 20414C-LON-SVR1, 20414C-LON-CA1, and 20414C-LON-CL1.

Results: After completing this exercise, students will have configured certificate revocation and distribution.

Question: Why is installing an enterprise root CA not recommended?

Question: What is the main benefit of OCSP versus CRL?

Question: What must you do to recover private keys?

Module Review and Takeaways

Review Questions

Question: What is required for using autoenrollment for certificates?

Tools

Tools	Use	Where to find
Certificate Authority console	Managing Certificate Authority	Administrative Tools
Certificate Template console	Managing certificate templates	CA console
Certificates console	Managing local certificate store	MMC snap-in
Certutil.exe	Managing CA and certificates	Command-line tool



Best Practice:

- When deploying a CA infrastructure, deploy a stand-alone root CA—a root CA that is not joined to a domain—and an enterprise subordinate CA—the issuing CA). After the enterprise subordinate CA receives a certificate from the root CA, take the root CA offline.
- Issue a certificate for the root CA for an extended period of time, such as 15 or 20 years.
- Use autoenrollment for certificates that are used widely.
- Use a Restricted Enrollment Agent whenever possible.

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
The CA is not configured to include CRL distribution point locations in the extensions of issued certificates. Clients may not be able to locate a CRL to check the revocation status of a certificate and certificate validation may fail.	Use the Certification Authority snap-in to configure the CRL distribution point extension and to specify the network location of the CRL. The default locations of the CRL are added to the CRL distribution point extension settings during a CA installation. The CA is configured to include the default locations in the extensions of all issued certificates.
The CA was installed as an enterprise CA, but Group Policy settings for user autoenrollment have not been enabled. An enterprise CA can use autoenrollment to simplify certificate issuance and renewal. If autoenrollment is not enabled, certificate issuance and renewal may not occur as expected.	Use the Group Policy Management console to configure user autoenrollment policy settings. Then use the Certificate Templates snap-in to configure autoenrollment settings on the certificate template.

Module 11

Planning and Implementing an Identity Federation Infrastructure

Contents:

Module Overview	11-1
Lesson 1: Planning and Implementing an AD FS Server Infrastructure	11-2
Lesson 2: Planning and Implementing AD FS Claims Providers and Relying Parties	11-15
Lesson 3: Planning and Implementing AD FS Claims and Claim Rules	11-20
Lesson 4: Planning and Implementing Web Application Proxy	11-26
Lab: Planning and Implementing AD FS Infrastructure	11-30
Module Review and Takeaways	11-41

Module Overview

Microsoft has provided a solution for identity federation called Active Directory® Federation Services (AD FS), which provides web-based authentication and authorization services for both internal and external users. Windows Server® 2003 R2 introduced AD FS, and Windows Server 2008 and Windows Server 2008 R2 included it as a server role. Windows Server 2012 also includes AD FS as a server role, with further improvements and changes.

Objectives

After completing this module, you will be able to:

- Plan and implement an identity federation infrastructure, including claims-aware application access.
- Plan and implement an AD FS server infrastructure.
- Plan and configure AD FS claim providers and relying parties.
- Design and deploy AD FS claims and claim rules.

Lesson 1

Planning and Implementing an AD FS Server Infrastructure

Before planning and implementing your AD FS server infrastructure, you should understand the concepts and components of AD FS, including the integration of AD FS with online services.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe AD FS.
- Define the various AD FS components.
- Define the basic AD FS deployment scenarios.
- Describe the options for implementing the AD FS server role.
- Outline the public key infrastructure (PKI) certificate considerations for AD FS.
- Outline the considerations for choosing an AD FS deployment scenario.
- Define the high availability options for AD FS.
- Describe the integration of AD FS with online services.

What Is AD FS?

AD FS is an identity access solution that provides a web-based single sign-on (SSO) experience to one or multiple web applications. AD FS works for both internal users and users from other organizations. An enterprise can use AD FS to provide internal users with SSO to web applications. However, typically you deploy it to allow users from another organization to access your network's web applications by using credentials and attributes defined in the user's organization as the user's *home realm* in identity federation terminology. The user's home realm can be an Active Directory Domain Services (AD DS) forest or a service maintained by other sign-on providers that use standards such as Liberty Alliance, Security Assertion Markup Language (SAML) 2.0 providers, and WS-*, or that use the Microsoft account service.

- AD FS:
 - Is an identity access solution
 - Provides browser-based SSO
 - Can also interact with other SAML 2.0 and WS-* providers
- AD FS enhancements in Windows Server 2012:
 - Dynamic Access Control integration
 - Improved installation experience
 - Enhanced Windows PowerShell cmdlets
- AD FS enhancements in Windows Server 2012 R2:
 - The Workplace Join feature
 - Multifactor authentication
 - Web Application Proxy

This enables each organization to manage its own identities, which include credentials and attributes. AD FS then presents these identities as claims to your web applications so that your application can make an access decision. This eliminates the requirement to maintain accounts in your organization for external users. When you implement federation, each organization can present their identities to other organizations, while continuing to manage those identities and accepting identities that other organizations present.

Windows Server 2003 R2 introduced AD FS 1.0, while Windows Server 2008 and Windows Server 2008 R2 included AD FS 1.1 as a server role. AD FS 2.0 was released to the web in May 2010 as a separate installable package. Windows Server 2012 includes a newer version (AD FS 2.1) as a server role, with functionality that has been added since AD FS 2.0, including:

- Integration with Dynamic Access Control. AD FS can consume AD DS user and device claims presented in Kerberos 5 authentication protocol tickets when a user or device is authenticated in the domain. This enables the use of these claims in Dynamic Access Control scenarios. These AD DS claims are incorporated into Windows® integrated authentication and AD FS can access them. Previously, AD FS could not access device claims, which are computer account attributes.
- Enhanced installation with Server Manager. Windows Server 2012 improves the AD FS server role installation. The Server Manager AD FS Configuration Wizard performs validation checks before installing the AD FS server role, including listing and installing dependent server roles, such as Web Server or Internet Information Services (IIS), and Web Server components, such as ASP.NET.
- Additional Windows PowerShell® cmdlets. AD FS in Windows Server 2012 includes new cmdlets that were not available in AD FS 2.0. These cmdlets provide installation functionality for the AD FS server role and allow configuration of federation servers and federation server proxies.

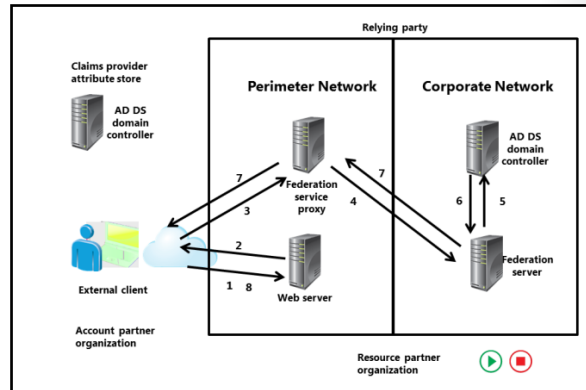
AD FS in Windows Server 2012 R2 includes the latest version of AD FS. The latest version offers a few new features and enhancements, including:


- The Workplace Join feature. Typically, computers have had to be joined to a domain to access company resources. Most information technology (IT) administrators join only company computers to a company AD DS domain. Usually, personal computers are not joined to a domain and do not have access to company resources. While Bring Your Own Device (BYOD) practices have grown tremendously in recent years, companies often restrict personal computing devices to guest networks and Internet access only. Workplace Join is a new feature to bridge the gap between domain-joined computers and personal computers that are not domain-joined. The Workplace Join feature allows employees to register their computing devices with the company domain and then gain access to selected company resources, such as an intranet or other web application.
- Multifactor authentication. This new feature is integrated loosely with the Workplace Join feature and provides support for multiple factors of authentication. In addition to the AD DS username and password, multifactor authentication can examine whether the accessing computing device is workplace-joined, to which network the accessing device is connected, and whether the employee is a member of a specific security group.
- Web Application Proxy. Windows Server 2012 R2 introduced a new role service named Web Application Proxy. The new role is part of the Remote Access role and provides AD FS proxy functionality, reverse proxy functionality, and preauthentication to internal HTTP and HTTPS web applications. The upcoming lesson entitled Planning and Implementing Web Application Proxy will present more information about this role service.


AD FS Overview

Overview of AD FS Components

The following table summarizes the various AD FS components and terminology in Windows Server 2012 and Windows Server 2012 R2.



AD FS component	Definition
Account partner organization	This is the organization whose accounts will be accessing the federated web applications.
Account federation server	This is the federation server in the account partner organization. It is a member of the organization's AD DS forest. The account federation server authenticates the users and then issues a security token that contains the claims representing the user's attributes. AD FS then provides the token to the resource's federation server. Attributes can include the user's group memberships, email address, and surname.
AD FS configuration database	This is either a Microsoft® SQL Server® database or the Windows Internal Database feature in Windows Server 2012. This database represents a single Federation Service. There can be multiple federation servers but only one configuration database in a Federation Service implementation.  Note: If you use SQL Server, you can create the AD FS configuration database with the Fconfig.exe command-line tool. For an implementation that uses the Windows Internal Database, you can use the AD FS Federation Server Configuration Wizard.
AD FS Web Agents	In Windows Server 2012, the AD FS 1.1 Web Agents are an AD FS role service that you can install. Typically, you do not install Web Agents in a new AD FS implementation. Rather, they are for use on application servers that will interact with an existing AD FS 1.x implementation.
Attribute store	In AD FS, an attribute store is a directory or database that you can use to store user accounts and their associated attributes. The Federation Service creates claims based on these attributes by retrieving them from the attribute store. Furthermore, an application or service hosted in a relying party makes authorization decisions on these claims.
Claim	A claim is a statement that a claims provider makes about a client (typically, a user). An example of a claim is the client's privilege, group memberships, name, email address, user principal name (UPN), or Security Accounts Manager account name. When a client is authenticated, the claims provider issues a token to the client containing the claims defined for that claims provider. Then the client provides this token to the relying party. If accepted, the client issues a token for the federated application and presents it to the application.
Claims provider	This is a federation partner organization that contains the users who are accessing the applications configured for the federation.

AD FS component	Definition
Claims provider trust	A claims provider trust is a relationship, created in a resource partner organization, which represents the organization whose accounts will be accessing the web applications that the resource partner organization hosts.
Claim rules	A claim rule is a statement of business logic that produces an outgoing claim by applying conditions (if <i>x</i> , then <i>y</i>) to an incoming claim. An example of a claim rule is to take an incoming UPN claim and transform this to an outgoing email address claim.
Federation metadata	This is the standardized data format for communicating between a relying party and a claims provider. It allows for proper configuration of federation trusts (relying party trusts and claims provider trusts). SAML 2.0 defines the format, with further extension in WS-Federation.
Federation Service	<ul style="list-style-type: none"> • If you install the Federation Service role service, the Windows Server 2012 computer will act in the federation server role by issuing tokens and serving as part of a Federation Service. Federation servers provide several functions, including: • Hosting a security token service (STS) and then issuing tokens based on the credentials that the user or application provides. • Accepting and relaying authentication requests from browser clients or web applications. • Implementing claim mappings, which provide links between the directory and attributes that the presented claims define.
Federation Service or Web Application Proxy	If you install the Federation Service Proxy or Web Application Proxy role service, you will cause the Windows Server 2012 or Windows Server 2012 R2 computer to act in the proxy role. Typically, you deploy the proxy in a perimeter network (also known as a screened subnet). It acts as an intermediary between a browser or other Internet client and the Federation Service.
Primary federation server	<p>If you deploy the AD FS configuration database by using the Windows Internal Database, the first federation server in the Federation Service hosts the only read/write copy of the AD FS configuration database. Therefore, it is referred to as the primary federation server, and all subsequent federation servers that you add to the Federation Service hold read-only copies of the configuration database. If the primary federation server is unavailable, you cannot make any changes to the Federation Service.</p> <p> Note: If you use SQL Server for the AD FS configuration database, the primary federation server concept does not apply. All federation servers can both read and write to an AD FS configuration database that SQL Server hosts.</p>
Relying party	This is the resource partner organization that receives and processes claims and issues tokens for accessing web applications based on those claims.
Relying party trust	The relying party trust establishes a relationship with another Federation Service or application that consumes claims originating from the Federation Service in which you create the relying party trust.
Resource federation server	This generates tokens containing web-based applications that users from the account partner organization deem accessible.
Resource partner organization	This is the resource partner organization that hosts the resource federation server. Typically, this organization hosts the web-based applications that you are federating. In the Federation Service, this organization is a partner in a relying party trust.

AD FS Operational Flow

The following steps describe the communication flow in this scenario:

1. The client computer, which is located outside of the network, must access a web-based application on the Web server. The client computer sends an HTTPS request to the Web server.
2. The Web server receives the request and identifies that the client computer does not have a claim. The Web server redirects the client computer to the Federation Service Proxy.
3. The client computer sends an HTTPS request to the Federation Service Proxy. Depending on the scenario, the Federation Service Proxy might prompt the user for authentication or use Integrated Windows authentication to collect the user's credentials.
4. The Federation Service Proxy passes on the request and the credentials to the federation server.
5. The federation server uses AD DS to authenticate the user.
6. If authentication is successful, the federation server collects AD DS information about the user. Then it uses that information 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.



Note: The federation servers in the account and resource partner organizations do not communicate directly. The only communication is web-based, occurring between the client computer and the federation servers at the account partner and the resource partner organizations, and between the client computer and the web application. Additionally, communication may occur between the client computer and the Web Application Proxy server, if the proxy is deployed.

Overview of AD FS Deployment Scenarios

AD FS Designs

The majority of AD FS implementations will have one of two designs:

- **Web SSO.** The primary focus of this AD FS design is to provide users with web SSO to multiple claims-aware applications or services. All users are external in this design, with your organization hosting all user accounts in AD DS or Lightweight Directory Access Protocol (LDAP), for example. There are no federation trusts in place. Therefore, your organization acts as both the account partner and the resource partner. This is a typical scenario for providing AD FS-secured, claims-aware applications or services to individual customers or consumers over the Internet. Typically, an organization hosts the applications or services in a perimeter network, along with a separate account store for the user's credentials.

The two primary AD FS designs are:

- **Web SSO**
 - Host all users within the organization, including AD DS or LDAP
 - Has no federation trusts in place
- **Federated web SSO**
 - Supports B2B
 - Configures federation trusts
 - Models trusts on business agreements or partnerships

- Federated web SSO. The federated web SSO scenario is the most typical scenario for business-to-business (B2B) collaboration. It is useful if two organizations create a federation trust relationship to allow users in the account partner organization to access web-based applications, which AD FS helps secure in the resource partner organization. After you configure the appropriate federation trusts, either organization can be a resource partner, an account partner, or both simultaneously. A federation trust relationship is modeled on a business-level agreement or a partnership between two organizations that requires end-to-end federation.

The scenario depicted in the previous topic's operational flow is an example of federated web SSO. In the example in that slide, an account federation server in Contoso, Ltd. and a resource federation server in Trey Research route user authentication requests from Contoso, Ltd. to web-based applications in Trey Research.

For the sake of simplicity, the example omits the proxy server, but this scenario can include a proxy in a perimeter network at the account partner, the resource partner, or both. You can use these proxies to avoid connecting directly from the Internet to the federation servers, which are domain members.

Options for Implementing the AD FS Server Role

Implementing a Federation Server

To implement a federation server, you must first install the Federation Service on a server that is joined to AD DS. You must use the same security level on this server as on your AD DS domain controllers, such as in a data center or other secured physical location. This means that you should install the Federation Service role service on the corporate network and behind the firewall, to help prevent exposure from the Internet.

You can apply the Federation Service role service on an account federation server, a resource federation server, or both. Whether the server functions as an account or resource federation server depends on the configuration of AD FS on that server. You can change the configuration at any time.

Windows Server 2012 AD FS provides the following role services:

- Federation Service
- Federation Service Proxy
 - Federation Service and Federation Service Proxy role services are mutually exclusive
- AD FS 1.1 Web Agents
 - Installed on application servers interacting with AD FS 1.x
- Windows Server 2012 R2 does not include the AD FS 1.1 Web Agents and replaces the Federation Service Proxy with the Web Application Proxy

Implementing a Federation Server Proxy

In Windows Server 2012, you install the federation service proxy role service in a perimeter network, on a server in a workgroup, or on a server that is joined to a forest on the perimeter network. You cannot install the Federation Service Proxy on the same server as the Federation Service.

The Federation Service Proxy accepts connections from Internet clients on behalf of the Federation Service, to handle those connections and provide an additional layer of protection.

In Windows Server 2012 R2, the Web Application Proxy role service performs federation server proxy services. You will learn more about this role service in an upcoming lesson entitled "Planning and Implementing Web Application Proxy".



Note: The Federation Service Proxy and Web Application Proxy are optional, but we highly recommend that you use proxy services if client computers are connecting from the Internet. Alternatively, you can use an application publishing solution, such as Microsoft Forefront® Unified Access Gateway.

AD FS 1.1 Web Agents

Windows Server 2012 provides the AD FS 1.1 Web Agents for use on application servers that will interact with an existing AD FS 1.x implementation. Therefore, you do not install them in a new AD FS implementation. Additionally, there is no support for Windows token-based applications in AD FS in Windows Server 2012. Applications and services can convert SAML tokens to Windows tokens for consumption in the application by using the Claims to Windows Token Service within Windows Identity Foundation (WIF).

Windows Server 2012 R2 does not provide the AD FS 1.1 Web Agents as a role service with the AD FS role.

Configuring Certificates

Federation Server Certificates

Federation servers require a certificate for five purposes, although you can use the same certificate for all of the functions. The five certificate uses are:

- Secure Sockets Layer (SSL) for IIS. AD FS uses a standard SSL certificate bound to the default web site in IIS. This certificate help secure the communication between the client computer and the various AD FS components, including the account federation server, the resource federation server, and the web server that hosts the federated application. In a Federation Service implementation that consists of multiple federation servers in a federation farm, you use the same SSL certificate for all federation servers.

AD FS uses the following certificates:

- SSL certificate
- Service communication certificate
- Token-signing certificate
- Token decryption certificate
- Server authentication certificate



Note: All AD FS clients must trust the SSL certificate, so we recommend that you use a certificate issued by a trusted third-party (public) certification authority (CA).

- Service communication between federation servers. This certificate secures Windows Communication Foundation (WCF) messages between federation servers. In a default configuration, the SSL certificate is used as the service communications certificate, but you can change this by using the AD FS Management console.
- Token signing. AD FS uses the token-signing certificate to sign and verify the authentication and application tokens that AD FS uses. A self-signed certificate is created and used by default, but you can change this to a CA-issued certificate in the AD FS Management console.
- Token decryption. This standard SSL certificate decrypts incoming tokens that it receives from a partner federation server. To provide the partner federation server with the ability to encrypt these tokens, the federation metadata publishes the public key of the token decryption certificate.
- Server authentication. You can use this standard SSL certificate on the federation server proxy or the Web Application Proxy server to secure communications between Internet clients and the proxy. All clients must trust it, so a public (third-party) CA issues it. To reduce the number of certificates to manage, you can use the same certificate that you use for the federation server.

For a federation server proxy, you must bind the server authentication certificate to the default web site in IIS before the Federation Server Proxy Configuration Wizard can run successfully.



Note: If any certificate that you use has certificate revocation lists (CRLs), the server with the configured certificate must be able to contact the server that distributes the CRLs.

Identifying Your AD FS Deployment Goals

The success of your AD FS project depends on identifying your deployment goals correctly. Once you determine which goals are relevant to your project, you can map them to a specific AD FS design. You can categorize most AD FS deployments as targeting one or more of the following four deployment goals:

- Providing your AD DS users access to your claims-aware applications and services. In this situation, you are providing SSO access for your own employees to your hosted resources. This deployment goal maps to a web SSO design. In this scenario, your organization is acting as the account partner organization, and is using AD FS to:
 - Secure applications or services that a perimeter network hosts, allowing employees signed in to AD DS on the corporate network to use SSO to access those applications or services.
 - Enable remote employees that AD DS authenticates to gain federated access to internal applications or services.
- Providing your AD DS users access to the applications and services of other organizations. This goal applies to a scenario in which your own employees have federated access to resources (applications or services) that another organization hosts. This deployment goal maps to a federated web SSO design in the resource partner and, optionally, in the account partner. For example:
 - Employees authenticated to AD DS on the corporate network can use SSO to access AD FS–secured services or applications that reside in a different organization.
 - Remote employees with AD DS credentials can access AD FS–secured services or applications that reside in a different organization by obtaining an AD FS token from your Federation Service.
- Providing users in another organization access to your AD FS–secured applications and services. This goal applies to a scenario in which you are hosting claims-aware applications or services in your organization and you need to provide federated access to users from another organization. This deployment goal can map to a federated web SSO design or a web SSO design and has two slightly different variations:
 - This goal enables you to provide federated access to your own users and users in organizations with which you have established a federation trust. This deployment goal maps to the federated web SSO design.
 - This goal enables you to provide federated access to users who have no association with your organization or a trusted organization, such as individual customers. An attribute store on your perimeter network hosts these customer accounts, and the customer uses these credentials to access multiple claims-aware applications or services, which the perimeter network also hosts.
- Providing users in your organization access to cloud-based services such as Microsoft Office 365™. In this situation, you are providing SSO access to Office 365 for your own employees. This deployment goal maps to a federated web SSO design. For example:
 - This goal enables you to provide SSO access to Office 365. Without AD FS, employees would need a second username and password just for Office 365. In such a situation, users would have to authenticate each time they use Office 365 services.

When you identify AD FS deployment goals, you can provide:

- Users access to your applications
- Users access to another organization's applications
- Another organization's users access to your applications

Configuring High Availability for AD FS Services

Federation server farms and proxy farms provide high availability for AD FS to federation servers and federation server proxies, respectively.

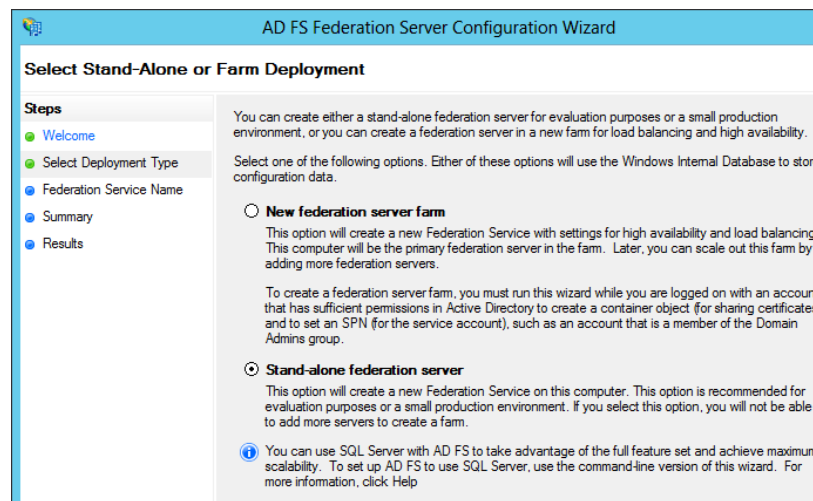
In Windows Server 2012, when you first complete the AD FS Federation Server Configuration Wizard, and you select the Create a New Federation Service option, you can choose to create a new federation server farm or a stand-alone federation server.

When configuring high availability for federation servers:

- Always use a Federation Server farm, even for single server deployments
- You can use WID or SQL Server for the configuration database

When configuring high availability for Web Application Proxy servers:

- Use multiple proxies
- You can cluster proxies with NLB or a hardware NLB solution




We recommend the stand-alone federation server option for evaluation purposes only. Otherwise, the new federation server farm option will create a server farm of one server, to which you can add more federation servers in the future for high availability, load balancing, and scaling.

High Availability for Federation Servers

The federation server farm is the primary unit of high availability for federation servers. There are multiple servers in a federation server farm, so you can distribute connections from AD FS clients across all federation servers in the farm by using Network Load Balancing (NLB) or a hardware NLB solution.

Creating a Federation Server Farm by Using the Windows Internal Database

When you use the AD FS Federation Server Configuration Wizard to create a federation server farm, it uses the Windows Internal Database (WID) for the AD FS configuration database. The first server in the farm is the primary federation server, which holds the only writeable copy of the configuration database. However, when you add other servers to the federation server farm, the configuration database replicates from the primary federation server. All replicas are updated whenever you make changes to the federation server farm.


 **Note:** Using WID provides high availability for the Federation Service, but it does have a single point of failure for making configuration changes to the Federation Service. Additionally:

- A WID federation server farm is limited to five federation servers.
- A WID federation server farm does not support artifact resolution or token replay detection. The SAML protocol defines both of these functionalities.

Creating a Federation Server Farm by Using SQL Server

If you use SQL Server for the AD FS configuration database when creating the federation server farm, configuration data does not replicate to each federation server in the farm. Instead, all federation servers read and write data into the same SQL database, allowing all federation servers in the farm to make configuration changes. SQL Server features provide high availability for the configuration database, including database mirroring, failover clustering, and replication.

When you use SQL Server to store the AD FS configuration database, all federation servers in the farm can read from and write to the configuration database. For this reason, the term primary federation server does not apply when using SQL Server.

 **Note:** When you use SQL Server for the federation server farm, the limitation of five federation servers per farm does not apply. Additionally, it supports artifact resolution and token replay detection.

Prior to Windows Server 2012 R2, if you used SQL Server and its configuration database, you could not use the AD FS Federation Server Configuration Wizard. Instead, you had to use the **fsconfig** command-line tool. If the AD FS administrator did not have the privileges to create SQL Server databases, you could have used the **fsconfig** tool to generate the appropriate SQL Server scripts to create the database.

High Availability for Federation Server Proxies and Web Application Proxy Servers

As with federation servers, the unit of high availability for proxies is the proxy farm. By configuring multiple proxies to protect the same Federation Service, you create a proxy farm.

Before the proxies can provide services as a farm, you must cluster them by using a single IP address and Domain Name System (DNS) fully qualified domain name (FQDN) that utilizes NLB or a hardware NLB solution. Additionally, all proxies in the farm must use the same server authentication certificate.

Integrating AD FS with Online Services

In addition to providing SSO for your internal users, and providing access to claims-aware applications or services for users from other organizations, you can use AD FS to provide SSO functionality for your internal users. This functionality allows them to access online services such as Office 365 and Windows Azure™.

AD FS and Office 365

Deploying and configuring AD FS for use with Office 365 is similar to the federated web SSO scenario described earlier. One difference is that the resource partner is Office 365, and you configure the relying party trust with Office 365 rather than another organization. Another difference is that when you federate with Office 365, we highly recommend that you deploy Federation Service proxies

You can use AD FS effectively with online services, such as:

- Office 365
 - AD FS provides SSO to Office 365 for AD DS users
- Windows Azure
 - AD FS provides SSO to Windows Azure for AD DS users
 - Windows Azure ACS is the integration point to Windows Azure for AD FS

or Web Application Proxy servers. In addition, we recommend that you deploy both the federation server farm and the proxy farm with two or more members for each farm to provide high availability for authentication to Office 365.

Prior to configuring AD FS for use with Office 365, you must install the Microsoft Online Services Sign-In Assistant and the Microsoft Online Services Module for PowerShell on each federation server in the federation server farm. Once you install the Microsoft Online Services Module, you can create the relying party trust with Office 365 by using the **Convert-MsolDomainToFederated** Windows PowerShell cmdlet. If you have a primary domain, such as treyresearch.net, but your AD DS deployment uses subdomains, such as corp.treyresearch.net, the primary domain is registered in Office 365 first. Then you add the subdomains, as necessary, by using the **New-MsolFederatedDomain** cmdlet.



Note: When you convert your primary domain to a SSO domain, all users become federated. Currently, it is not possible to perform a staged rollout of SSO to Office 365.

Configuring AD FS and Office 365 for SSO

Assuming that your AD FS infrastructure is installed and functioning correctly, you can configure SSO for Office 365 by following this procedure:

1. On the first federation server, install the 64-bit version of the **Microsoft Online Services Sign-In Assistant**.
2. By using the Windows Server 2012 Server Manager Add Roles and Features Wizard, install the **.NET Framework 3.5 Features**.
3. Install the 64-bit version of the **Microsoft Online Services Module for PowerShell**.
4. Repeat steps 1 through 3 on all other federation servers in the federation server farm.
5. On the first federation server, run the **Connect-MsolService** Windows PowerShell cmdlet to connect to your Office 365 tenant. Enter your administrative credentials in the **Enter Credentials** dialog box.
6. Convert your primary domain to use federation with the **Convert-MSOLDomainToFederated – DomainName <domainname>** cmdlet.
7. Add any required subdomains used for UPN with the **New-MSOLFederatedDomain –DomainName <sub.domainname>** cmdlet.

In addition to configuring your relying party trust with Office 365, you must notify Office 365 of certain changes by updating the Office 365 federation metadata. These changes include renewing your token-signing certificate and generating a new certificate. You can use the Microsoft Office 365 Federation Metadata Update Automation Installation Tool for this purpose. Follow this procedure:

1. Sign in to the primary federation server with an account with AD FS administrative privileges (typically, a service account).
2. Download the **Microsoft Office 365 Federation Metadata Update Automation Installation Tool** to the primary federation server.
3. Start Windows PowerShell as an administrator, and then switch to the directory containing the **Microsoft Office 365 Federation Metadata Update Automation Installation Tool**.
4. At the Windows PowerShell prompt, type **.\O365-Fed-MetaData-Update-Task-Installation.ps1**, and then press Enter.

5. Enter the federated domain at the Windows PowerShell prompt.
6. Enter your Office 365 administrative username and password when prompted.
7. Enter the password of the AD FS administrative account when prompted.

At this point, the O365-Fed-MetaData-Update-Task-Installation.ps1 script will create a scheduled task that will run once a day to update the Office 365 federation metadata. This task will run in the security context of the account that you used to create the task, as previous steps outlined.



Note: You must run the O365-Fed-MetaData-Update-Task-Installation.ps1 script for each federated domain in your organization.

AD FS and Windows Azure

In Windows Azure, you can configure AD FS as an identity provider by using the Windows Azure Access Control Service. This enables your users to authenticate to ASP.NET applications hosted in Windows Azure by using their AD DS credentials.

Windows Azure Active Control Service

The Access Control Service (ACS) provides a means of integrating Windows Azure with standards-based identity providers such as AD FS, Microsoft accounts (formerly Windows Live ID), and Facebook. This facilitates authenticated and authorized users to access Windows Azure-hosted applications and services. You can configure AD FS as an identity provider in ACS by:

1. Adding AD FS as an identity provider in the ACS Management Portal.
2. Adding a token decryption certificate to ACS, optionally. This decrypts tokens received from AD FS in the ACS Management Portal.
3. Adding a relying party for your ACS namespace in AD FS.
4. Adding claim rules for the ACS namespace, in AD FS.



Additional Reading: For more information on integrating AD FS with the Windows Azure ACS, go to <http://go.microsoft.com/fwlink/?LinkID=285332>.

Demonstration: Preparing DNS and Installing the AD FS Server Role

Consider the following scenario for implementing the proxy. At A. Datum Corporation, the Chief Information Officer has recently announced the following policies:

- The internal web application must be available to external users.
- The internal AD FS server must not be directly accessible from outside of the company network.

You are preparing the infrastructure to support the new policies.

In this demonstration, you will see how to:

- Request and install a certificate for the AD FS server.
- Install the AD FS server role.

Demonstration Steps

Prepare DNS

- Add a new DNS host record to point **adfs.adatum.com** to 172.16.0.10

Install the AD FS server role

1. On LON-DC1, in Server Manager, add the **Active Directory Federation Services** server role.
2. On LON-DC1, run the AD FS Federation Server Configuration Wizard by using the following parameters:
 - Create a new Federation Service.
 - Create a stand-alone deployment.
 - Use the **adfs.adatum.com** certificate.
 - Choose a service name of **adfs.adatum.com**.
3. On LON-DC1, open Windows PowerShell, and then use the **set-ADFSProperties – AutoCertificateRollover \$False** command to enable modification of the assigned certificates.
4. In the AD FS Management console, add the **adfs.adatum.com** certificate as a new token-signing certificate. Verify that the certificate has a subject of **CN=adfs.adatum.com**, and its purposes include **Proves your identity to a remote computer** and **Ensures the identity of a remote computer**.
5. Make the new certificate the primary certificate and remove the old certificate.

Lesson 2

Planning and Implementing AD FS Claims Providers and Relying Parties

Once you deploy your AD FS infrastructure, you can configure it to work with various identity stores by implementing claims providers. You can also configure it to work with applications or services by implementing relying parties.

Lesson Objectives

After completing this lesson, you will be able to:

- Provide an overview of the claims provider and relying party roles.
- Provide an overview of AD FS compatible applications.
- Describe the options for implementing attribute stores.
- Provide an overview of claims provider trust components.
- Provide an overview of relying party trust components.
- Describe the considerations for configuring claims provider trusts.
- Describe the considerations for configuring relying party trusts.

Overview of Claims Provider and Relying Party Roles

Claims Providers

A claims provider is the account partner organization that provides claims to its users. Claims are statements about users that both parties in a federation understand. These statements correspond to values that the claim stores. For example, a claim can be a common name, an email address, or an employee number. You can establish a relationship with a claims provider by configuring a claims provider trust in the resource partner organization.

You can define the way in which federated partners exchange claims in the Federation Service, and then store these claims in the AD FS configuration database. Claims are sourced from attribute stores or they are sourced by applying claim rules to transform an incoming claim into another value. For example, you can transform a claim of Outside Sales into a value of Sales, and then send this new value as an outgoing claim. AD FS supports any claim type, but its default configurations include the following claim types.

A claims provider offers the following functionality:

- Makes statements about users, such as their email address and employee number
- Provides claim definitions that both parties understand
- Enables claim rules to provide translations between parties

Relying parties interact with claims by:

- Consuming claims
- Providing access tokens to services or applications based on claims
- Enabling claims rules to transform incoming claims for applications or services

Claim type	Description
Email address	The email address of the user.
Given name	The given name of the user.
Name	The unique name of the user.
UPN	The UPN of the user.
Common name	The common name of the user.
AD FS 1.x email address	The email address of the user when interoperating with AD FS 1.1 or AD FS 1.0.
Group	A group of which the user is a member.
AD FS 1.x UPN	The UPN of the user when interoperating with AD FS 1.1 or AD FS 1.0.
Role	A role that the user has.
Surname	The surname of the user.
Personal private identifier	The private personal identifier of the user.
Name identifier	The SAML name identifier of the user.
Authentication method	The method used to authenticate the user.
Deny only group SID	The deny-only group security identifier (SID) of the user.
Deny only primary SID	The deny-only primary SID of the user.
Deny only primary group SID	The deny-only primary group SID of the user.
Group SID	The group SID of the user.
Primary group SID	The primary group SID of the user.
Primary SID	The primary SID of the user.
Windows account name	The domain account name of the user in the form of <domain>\<user>.

Relying Party

A *relying party* is a Federation Service or an application that is consuming claims in a particular transaction. The relying party consumes claims that a claims provider presents, and then acts as a web service to request claims from a trusted claims provider. You can establish a relationship with a relying party by configuring a relying party trust in the account partner organization.

Typically, a relying party also issues tokens to the claims-aware services or applications in its own organization, based on the claims contained in the security tokens that the AD FS client presents to it, assuming that a trusted claims provider produces the tokens.

You can configure claim rules for transforming incoming claims from the claims provider into other values, so that you can present compatible claims to the application or service.

Overview of AD FS Compatible Applications

Windows Identity Foundation

WIF provides an application programming interface (API) for utilizing federation and claims in browser-based applications. Applications and services use WIF to process tokens that AD FS issues, and make identity-based decisions based on the tokens presented. You can use WIF to translate federation claims to Windows NT tokens for interaction with applications that are not claims-aware.



Additional Reading: For more information on WIF, go to Microsoft Developer Network (MSDN) at <http://go.microsoft.com/fwlink/?LinkID=285330>.

WIF provides the following:

- API for claims-aware browser applications
- Translation of Windows NT tokens for applications that are not claims-aware

Third-party applications and services can be based on the following protocols:

- SAML
- WS-Federation
- WS-Trust protocols

Third-Party Applications and Services

AD FS is standards-based, so it can interoperate with applications and services based on the SAML, WS-Federation, and WS-Trust protocols. Some examples applications and services are RSA SecurID, IBM Tivoli Federated Identity Manager, Shibboleth 2.0, and Ping Identity PingFederate.

Options for Implementing Attribute Stores

An attribute store is a directory or database in AD FS that stores user accounts and their associated attributes. The Federation Service retrieves attributes from the attribute store and then creates claims based on these attributes. An application or service that a relying party hosts then makes authorization decisions based on these claims. AD FS in Windows Server 2012 supports attribute stores based on the following directories or databases:

- Active Directory Domain Services (Windows Server 2003 and newer)
- LDAP
- Microsoft SQL Server 2005 and newer
- Custom attribute stores

Attribute stores for AD FS in Windows Server 2012 can be based on the following:

- AD DS
 - Windows Server 2003 or newer
- LDAP
 - Including AD LDS
- SQL Server 2005 or newer
- Custom attribute stores
 - Connect to other external data stores, such as CSV files

An example of an LDAP-based attribute store is a web SSO deployment for providing applications or services to individual customers, in which you deploy the federation and application servers in a perimeter network, and store customer accounts in an AD LDS database. In this case, an LDAP attribute store is created in AD FS to use the AD LDS instance as its source.

You can use a custom attribute definition to connect to other external data stores, such as a comma-separated value (CSV) file, to query for claim information. Defining and using custom attribute stores is custom development work, which is beyond the scope of this course.



Additional Reading: For more information on custom attribute stores, go to MSDN at <http://go.microsoft.com/fwlink/?LinkID=285331>

Overview of Claims Provider Trust Components

A claims provider trust is conceptually similar to a Windows trust, but it is entirely web-based. Once a claims provider trust is established, a resource partner accepts as valid the claims and attributes (authentications) that the account partner provides. This is similar to a Windows trust, in which the trusting domain accepts as valid the authentications from the trusted domain. However, there is no connectivity between the account and resource federation servers, and all network traffic in an AD FS scenario is web-based. An administrator can configure the information for establishing the trust on the account and resource federation servers in three ways:

- By entering the information manually.
- By importing a policy file that the trust partner provides.
- By importing data about the claims provider from federation metadata published online via a URL that the claims provider organization provides.

A claims provider trust is comprised of the following:

- Federation metadata. This is data about the claims provider. You obtain it from a URL for the claims provider, import it from a file that the claims provider provides, or configure it manually.
- Claims. The claims provider defines the claims that are available.
- Claim rules. Claim rules are defined on a claims provider trust. They specify how AD FS and its relying parties use claims received from the claims provider.

Claims provider trusts are conceptually similar to Windows trusts, but are entirely web-based

Claims provider trusts are comprised of the following components:

- Federation metadata
- Claims
- Claim rules

Overview of Relying Party Trust Components

A relying party trust defines a relationship with another Federation Service or application to allow that entity to consume claims originating from the Federation Service in which the relying party trust is created. It is similar to a claims provider trust, and includes the following components:

- Federation metadata. This is data about the relying party, which you obtain from a URL for the relying party, import from a file that the relying party provides, or configure manually.
- Claims. The claims provider defines the available claims.
- Claim rules. The relying party trust defines claim rules to determine how to present claims that pass through the trust to the relying party.

A relying party trust is a relationship with another Federation Service or application

The components of a relying party trust include:

- Federation metadata
- Claims
- Claim rules

Configuring Claims Provider and Relying Party Trusts

Claims Provider Trusts

A claims provider trust is used to issues claims about users. When you add a claims provider trust, you need to specify the data source. You should also take into account the following considerations about claims providers trusts:

- Use automatic retrieval for federation metadata. If possible, you should configure a claims provider trust by using federation metadata that is published to a URL or in a file that the claims provider provides. This automatic retrieval method simplifies the trust configuration process greatly, and it avoids the inconvenience of troubleshooting a manual configuration.
- Ensure that certificates are valid. The token-signing certificate configured for the trust must be a valid certificate (current and trusted) and must be unique to that claims provider trust. You use it to verify that the configured claims provider issued the tokens received.
- Use FQDNs. Although this may be out of your control, certificates configured in claims provider trusts should use FQDNs, such as myserver.contoso.com, and not an unqualified host name, such as myserver. This helps reduce potential impersonation issues.


When configuring claims provider and relying party trusts:

- Use federation metadata automatic retrieval where possible
- Ensure that all certificates that you use are valid
- Use FQDNs and not unqualified host names in certificates
- Be aware that federation metadata is automatically updated when using the automatic retrieval option

Remember that most AD FS trust problems are related to URLs or certificates

Relying Party Trusts

When you configure relying party trusts, you must consider many of the same factors as when you create claims provider trusts. However, if you configure a relying party trust by using automatic retrieval of federation metadata from a URL, you must remember that AD FS updates the federation metadata every 24 hours. Therefore, you do not have to edit the federation trust manually when changes are made at the relying party, such as renewing, adding, or removing certificates.

 **Note:** Usually, when troubleshooting most configuration issues with AD FS trusts, you will discover that the problem is related to either certificates or incorrect URLs.

Lesson 3

Planning and Implementing AD FS Claims and Claim Rules

Once you have your claims provider and relying party trusts in place, you can configure claims and claim rules to define what claims you provide to AD FS clients and how to transform incoming claims for use with your applications or services.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the options for configuring claim rules.
- Configure claim rules by using the claims rule language.
- Design AD FS claims.
- Describe the options for integrating with previous AD FS versions.

Options for Configuring Claim Rules on a Federation Trust

Claim Rule Template

When configuring a claim rule by using the Add Transform Claim Rule Wizard, first you must determine which claim rule template to use. Claim rules are specific to a federation trust; you cannot use them on other trusts that you define in the Federation Service. Therefore, AD FS includes claim rule templates that simplify the creation of claims rules. These templates provide preconfigured claim rule language frameworks through a user interface–based form. These templates are based on the most common types of created claim rules. The following table provides examples of claim rule templates.

When configuring claim rules, remember:

- Claim rule templates provide a preconfigured claim rule language framework
- Descriptive claim rule names should be descriptive helps to manage the Federation Service in the future

You must specify incoming claims with one of the following templates:

- Pass Through or Filter an Incoming Claim
- Transform an Incoming Claim

You must specify outgoing claims with all other templates

Claim rule template	Description
Pass Through or Filter an Incoming Claim	Use to accept or issue unchanged claims.
Transform an Incoming Claim	Map an incoming claim to a different claim type, or map a claim value to a new claim value. For example, you can use this template on an incoming claim to replace an email suffix with a new value, or to transform an incoming UPN claim to an email address claim.
Send LDAP Attributes as Claims	Map LDAP attributes to claims so that you can source claims from any LDAP store and from AD DS. A single rule can issue multiple claims, but performance can be slow due to LDAP lookups. Additionally, you cannot use custom LDAP filters for your queries.

Claim rule template	Description
Send Group Membership as a Claim	Use with AD DS attribute stores. Rules based on this template can create a single claim based on the specified group. In addition, because the rule uses group SIDs to issue the claim, performance is very fast. This is because this rule does not require account lookup and can issue only one claim based on one group.
Send Claims Using a Custom Rule	Use to create custom rules by using the AD FS claim rule language. You can use this for additional scenarios that the predefined templates do not address, such as rules used with a custom attribute store, sending claims based on multiple incoming claims, or allowing use of a custom LDAP filter.

Claim Rule Name

Regardless of the template you use, you must define a claim rule name. Choosing a descriptive name will assist in managing the Federation Service in the future, as no one will know how to use a rule named Claim Rule 1.

Incoming Claims

You must specify an Incoming claim type when using the following two claim templates:

- Pass Through or Filter an Incoming Claim
- Transform an Incoming Claim

You can specify whether the rule will act on all claim values, only on specific claim values, only on claim values matching a specific email suffix, or on claim values that begin with a specific value.

Outgoing Claims

You define an Outgoing claim type when creating your rule, except for rules that use the Pass Through or Filter an Incoming Claim template. Depending on the outgoing claim type that you select, you may need to specify an outgoing claim value also. For example, you may use an outgoing claim type of Group with a value of Admin if the user is a member of the AD DS Domain Admins group.

Creating Claim Rules by Using the Claims Rule Language

Claims Rule Language

You can create the majority of claim rules by using the predefined claim rule templates. However, other scenarios may require that you use a custom rule. The Send Claims Using a Custom Rule template allows you to define a rule by using the AD FS claims rule language.

Syntax for Claims Rule Language

In a structure similar to many programming languages, the claims rule language is usually structured as an if-then statement, as follows:

```
If (condition is true), Then (issue a claim with this value)
```

When you create claim rules with the claims rule language use the following template:

- Send Claims by Using a Custom Rule
- Structure rules as if-then statements:
 - If (condition is true), Then (issue a claim with this value)
- Use claim rule language statements from existing rules as starting blocks

In claims rule language, a special operator (`=>`) separates the condition from the issuance statement, and a semicolon ends the statement. Here is an example of a claims rule:

```
c:[Type == http://contoso.com/department] => issue(Type =
"http://adatum.com/department", Issuer = c.Issuer, OriginalIssuer = c.OriginalIssuer,
Value = c.Value, ValueType = c.ValueType);
```

In this example, `c:[Type == http://contoso.com/department]` is the condition, and `issue(Type = "http://adatum.com/department", Issuer = c.Issuer, OriginalIssuer = c.OriginalIssuer, Value = c.Value, ValueType = c.ValueType)` is the issuance statement.

You can combine multiple conditions by using the conjunction logical operator `&&`. For example, the following rule issues a claim only if the user has the Editor role and an email claim:

```
c1:[Type == "http://contoso.com/role", Value=="Editors"] &&
c2:[Type == "http://contoso.com/email"]
=> issue(claim = c1);
```

Additionally, you can combine multiple incoming claims into a single outgoing claim. For example, the following rule combines incoming claims of location and role into a single outgoing targeted claim:

```
c1:[Type == "http://contoso.com/location"] &&
c2:[Type == "http://contoso.com/role"]
=> issue(Type="http://contoso.com/targeted", Value=c1.value+" "+c2.value);
```

In this example, incoming claims of location equaling Seattle and role equaling Editor would result in an outgoing targeted claim of Seattle Editor.



Additional Reading: For more information on the claims rule language, go to AD FS 2.0 Claims Rule Language Primer on TechNet at <http://go.microsoft.com/fwlink/?LinkID=285329>.

Creating Rules with Claims Rule Language

If you are unfamiliar with the AD FS claims rule language, an easy way to create custom rules and learn the syntax is to create more basic rules by using the available claim rule templates. Then you can use the claim rule language statements from these rules as your custom rule's building blocks.

You can obtain the claim rule language for an existing rule by opening the Edit Rule dialog box. Click View Rule Language, and then copy the rule language to a text editor.

Designing AD FS Claims

As you determine which claims you want your AD FS deployment to use, you should consider several guidelines. The following sections provide details about these guidelines.

Claims Source

The Federation Service defines claims exchanged between federation partner's Federation Service. These claims are sourced in one of two ways:

When designing AD FS claims, remember:

- Claims can be sourced from the following:
 - An attribute store
 - A transformed incoming claim
- Claims flow in this way:
 - Received from a claims provider trust
 - Processed by claim rules (acceptance transform rules and issuance authorization rules)
 - Sent to the relying party via the relying party trust
- A claim is incoming or outgoing from the perspective of the trust it references
- Typically, claim types are expressed as URIs
- Claim descriptions specify claim types that are accepted
- Claims will be designed based on application needs; some applications only need authentication while others need additional information

- By using values that you obtain from an attribute store, such as a claim of Sales Executive retrieved from the title attribute of an AD DS user account.
- By a claim rule transforming an incoming claim value into another outgoing claim value, such as transforming a UPN inbound claim into an email outgoing claim.

How Claims Flow

The Federation Service processes and delivers claims from a claims provider to a relying party through the *claims pipeline*. The claims pipeline is a logical construct that provides the following flow:

1. Claims are received from the claims provider and processed by claim rules (acceptance transform rules) that are defined in the claims provider trust. These rules specify which of the claims that the claims providers issues to accept.
2. Claims that the acceptance transform rules process are submitted to the issuance authorization rules, which determine if the user has permission to access the relying party. These rules are defined as part of the relying party trust.
3. Acceptance transform rules output is used as input to the issuance transform rules, which are also defined on the relying party trust. These rules determine which claims to send to the relying party that the trust defines.

Incoming vs. Outgoing Claims

When you define claim rules, you need to consider the source of the incoming claims and the destination for outgoing claims, which the following table describes.

Type of federation trust	Incoming claims	Outgoing claims
Claims Provider Trust	Claims received from claims provider	Claims sent to Federation Service
Relying Party Trust	Claims received from corresponding claims provider trust via Federation Service	Claims sent to relying party (application or another Federation Service)

Claim Types

The claim type provides the context for a claim value. Typically, the claim type is expressed as a Uniform Resource Identifier (URI). The following table lists the default claim types in AD FS.

Name	Description	URI
Email Address	The user's email address	http://schemas.xmlsoap.org/ws/2005/05/identity/claims/emailaddress
Given Name	The user's given name	http://schemas.xmlsoap.org/ws/2005/05/identity/claims/givenname
Name	The user's unique name	http://schemas.xmlsoap.org/ws/2005/05/identity/claims/name
UPN	The user's UPN	http://schemas.xmlsoap.org/ws/2005/05/identity/claims/upn
Common Name	The user's common name	http://schemas.xmlsoap.org/claims/CommonName

Name	Description	URI
AD FS 1.x Email Address	The user's email address, when interoperating with AD FS 1.1 or AD FS 1.0	http://schemas.xmlsoap.org/claims/EmailAddress
Group	A group of which the user is a member	http://schemas.xmlsoap.org/claims/Group
AD FS 1.x UPN	The user's UPN when interoperating with AD FS 1.1 or AD FS 1.0	http://schemas.xmlsoap.org/claims/UPN
Role	A role that the user has	http://schemas.microsoft.com/ws/2008/06/identity/claims/role
Surname	The user's surname	http://schemas.xmlsoap.org/ws/2005/05/identity/claims/surname
Personal Private Identifier	The user's private identifier	http://schemas.xmlsoap.org/ws/2005/05/identity/claims/privatepersonalidentifier
Name Identifier	The user's SAML name identifier	http://schemas.xmlsoap.org/ws/2005/05/identity/claims/nameidentifier
Authentication Method	The method used to authenticate the user	http://schemas.microsoft.com/ws/2008/06/identity/claims/authenticationmethod
Deny Only Group SID	The user's deny-only group SID	http://schemas.xmlsoap.org/ws/2005/05/identity/claims/denyonlysid
Deny Only Primary SID	The user's deny-only primary SID	http://schemas.microsoft.com/ws/2008/06/identity/claims/denyonlyprimarysid
Deny Only Primary group SID	The user's deny-only primary group SID	http://schemas.microsoft.com/ws/2008/06/identity/claims/denyonlyprimarygroupsid
Group SID	The user's group SID	http://schemas.microsoft.com/ws/2008/06/identity/claims/groupsid
Primary Group SID	The user's primary group SID	http://schemas.microsoft.com/ws/2008/06/identity/claims/primarygroupsid
Primary SID	The user's primary SID	http://schemas.microsoft.com/ws/2008/06/identity/claims/primarysid

Claim Descriptions

Claim descriptions define the claim types that the Federation Service supports. In addition, they specify which claim types the Federation Services accept and which types they can send. This list of claim types is stored in the AD FS configuration database and published to the federation metadata. Each claim description contains a name, claim type URI, description, and publishing state (Accepted or Sent).

Options for Interoperation with Previous Versions of AD FS

AD FS 1.x Interoperation

AD FS in Windows Server 2012 can interoperate with both AD FS 1.0 (Windows Server 2003 R2) and AD FS 1.1, which was sometimes used with Windows Server 2008 or Windows Server 2008 R2. Windows Server 2012 supports interoperation in three ways:


1. By sending claims from an AD FS 1.x Federation Service to a Windows Server 2012 Federation Service.
2. By sending claims that are compatible with AD FS 1.x from a Windows Server 2012 Federation Service to an AD FS 1.x Federation Service.
3. By sending claims that are compatible with AD FS 1.x from a Windows Server 2012 Federation Service to web servers running the AD FS 1.x claims-aware web agent.

AD FS in Windows Server 2012 can interoperate with the following AD FS 1.0 or 1.1 components:

- AD FS 1.x Federation Service
- AD FS 1.x claims-aware web agent

Windows Server 2012 provides the following Name ID claim types to AD FS 1.x:

- AD FS 1.x email address
- AD FS 1.x UPN
- Common name
- Group


 **Note:** Windows Server 2012 does not support the AD FS 1.x Windows NT token-based web agent in AD FS.

Name ID Claim Type

AD FS 1.x uses the concept of an *identity claim type*. The Windows Server 2012 Name ID claim type is the equivalent of this identity claim type, and you must use an identity claim whenever interoperating with AD FS 1.x. An identity claim indicates that the claim belongs to a particular user or security principal. AD FS in Windows Server 2012 provides the following Name ID claim type formats that map to the equivalent AD FS 1.x identity claims:

- AD FS 1.x Email address
- AD FS 1.x UPN
- Common name
- Group

An identity claim sent to AD FS 1.x can contain only one Name ID claim. You must configure it in the appropriate format, as described above.

 **Note:** Additionally, the Federation Service in AD FS 1.x can only accept incoming claims with a URI that begins with <http://schemas.xmlsoap.org/claims/>.

Lesson 4

Planning and Implementing Web Application Proxy

Once you have your AD FS infrastructure in place, you may expand access to the applications for external parties. You can use the new Web Application Proxy role service to extend AD FS to the perimeter network to give access to the external parties.

Lesson Objectives

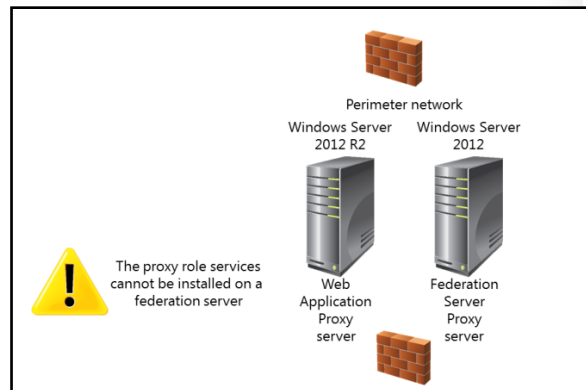
After completing this lesson, you will be able to:

- Describe the Web Application Proxy service.
- Describe the scenarios when a Web Application Proxy server provides value.
- Plan the integration of AD FS and the Web Application Proxy.
- Plan Web Application Proxy authentication.
- Understand how to install and configure the Web Application Proxy.

Web Application Proxy Overview

Prior to Windows Server 2012 R2, the Federation Service Proxy role service provided proxy services for AD FS. In Windows Server 2012 R2, the Web Application Proxy role service replaces the Federation Service Proxy role service and proxies connections for AD FS federation servers. These two role services provide proxy services in much the same way. The characteristics of a Web Application Proxy are:

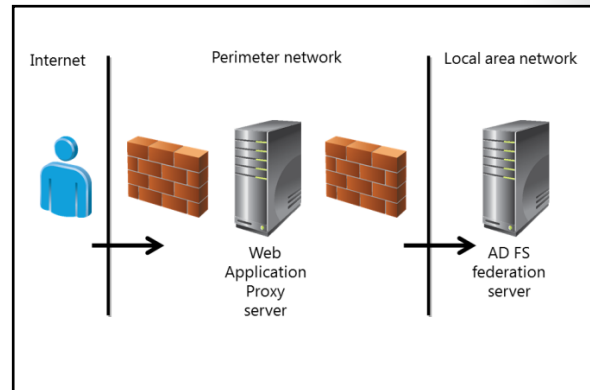
- Most commonly, you deploy Web Application Proxy servers to a perimeter network to enhance security of a federated environment by requiring all outside communication to go through the proxy before reaching the internal servers. This allows the federation servers to remain on the internal local area network without direct communication to the Internet.
- You cannot install the Web Application Proxy role service on an AD FS federation server. Instead, you must install it on another server. Often, such a server is dedicated to the proxy service only.
- You can deploy more than one Web Application Proxy server to provide high availability. This ensures that proxy services are available in the event that a proxy server fails.
- In addition to providing proxy services for AD FS, a Web Application Proxy server can be used to proxy basic web applications such as Microsoft Outlook® Anywhere.



Web Application Proxy Scenarios

The Web Application Proxy is designed for use in a perimeter network. The primary purpose of the proxy is to enhance security of the environment by not exposing the federation servers on the local area network to the Internet. The following scenarios describe some common uses for the proxy:

- You have a web application that both internal employees and external partners use. Internal employees need to use their AD DS credentials to authenticate to the web application. External partners need to use their company AD DS credentials to authenticate to the web application. You can deploy federation servers to the local area network to authenticate internal employees directly. You can install proxy servers in the perimeter network to proxy requests from the internet to the internal federation servers.
- Your company subscribes to Office 365. You need to deploy a SSO solution so that internal employees can take advantage of SSO when using Office 365. So far, your scenario has required federation servers only on the local area network. If you want external users to take advantage of SSO for Office 365 also, you will require at least one proxy server.



A common use for proxy servers is to connect external users, whether employees who are working remotely, partners, or employees from other organizations. You should consider using the Web Application Proxy server when a scenario involves external users.

Planning the Integration of AD FS and Web Application Proxy

The Web Application Proxy server is an important component of your infrastructure. Because it resides in the perimeter network, you must plan the deployment and integration with AD FS carefully. It is especially important to consider security concerns in your planning. You should consider the following design and deployment options for the proxy server before the deployment:

- Domain join. Unlike federation servers, the proxy server does not need to be joined to an AD DS domain. This is important because in most environments, the internal AD DS domain is not available in the perimeter network. You should consider whether the proxy server is joined to a perimeter-only AD DS domain.
- Network communication. The proxy server communicates with federation servers over port 443. Thus, if a firewall separates the local area network and the perimeter network, you must open port 443 to allow the communication. The proxy server communicates to the Internet on port 443 too. If a firewall separates the perimeter network from the Internet, you must open port 443 to the Internet to allow the communication.

Considerations for integrating AD FS and a Web Application Proxy include:

- Domain join
- Network communication such as firewalls and ports
- SSL certificates
- DNS and name resolution

- SSL certificates. Internal federation servers often use an internal PKI to obtain certificates. However, internal certificates are often unsuitable for use on Internet-facing proxy servers because the certificates are not trusted. In most cases, a third-party SSL certificate is preferred for proxy servers since third-party certificates are trusted. The proxy servers must use the same certificate subject name that the federation servers use.
- DNS and name resolution. It is important to remember that applications that communicate with AD FS are often not aware of the existence of proxy servers. DNS plays a key role in ensuring that the right communication goes to the right server. In addition, in most situations, organizations want to send internal traffic straight to the internal federation servers, instead of through the proxy servers. Split DNS, whereby internal DNS servers handle internal name resolution requests and external DNS servers handle external name resolution requests, is one solution for ensuring that internal traffic stays internal and external traffic goes to a proxy server.

Planning Web Application Proxy Authentication

AD FS federation servers handle all the authentication traffic coming to the proxy. The proxy serves as an intermediary in the communication between clients and AD FS federation servers. When you plan Web Application Proxy authentication, consider the following:

- If all communications must be authenticated, AD FS preauthentication can deliver that functionality. AD FS preauthentication requires users to enter their credentials before they access a published application. This ensures that traffic is authenticated before it is sent to the application, which can reduce attacks such as a Denial of Service (DoS) attack.
- If all communication needs to be authenticated for external users only, you can use pass-through preauthentication. You might use pass-through preauthentication to allow users to access an application without entering their credentials. This works only for internal users who authenticate with AD DS as part of their computer sign-in. You might also use it to allow the backend server, which hosts the published application, to handle the authentication. In such a case, the backend server decides if authentication is required and the Web Application Proxy server forwards requests directly to the backend server.

- The Web Application Proxy server is used to authenticate requests going to web-based applications
- When planning Web Application Proxy authentication, consider the following questions:
 - Should all communication be authenticated? AD FS preauthentication can require users to enter their credentials before reaching a published application
 - Should preauthentication be passed through? You can use this option for internal users who are authenticated to their work computer already
- The Web Application Proxy can take advantage of new features in Windows Server 2012 R2 such as Workplace Join, SSO, and multifactor authentication

Because the Web Application Proxy uses AD FS for authentication, it can take advantage of some of the authentication features in Windows Server 2012 R2:

- Workplace Join is a new feature that allows you to associate personal computing devices with an AD DS domain. Once associated with a domain, they can gain access to resources and applications.
- Single sign-on (SSO) is a feature that allows users to enter credentials once and then access any of the published applications without authenticating manually thereafter.
- Multifactor authentication is a new feature that works together with Workplace Join. It allows AD FS to authenticate user credentials and require that the consuming device be associated with the domain.

Demonstration: Installing and Configuring the Web Application Proxy

Consider the following scenario for implementing the Web Application Proxy. At A. Datum Corporation, the Chief Information Officer has recently announced the following policies:

- The internal web application must be available to external users.
- The internal AD FS server must not be accessible directly from outside the company network.

In this demonstration, you will see how to:

- Install and configure the Web Application Proxy role.
- Configure certificates for the Web Application Proxy server.
- Configure the proxy to allow access to an internal web application.

Demonstration Steps

Install the Web Application Proxy role

- On LON-SVR2, in Server Manager, add the **Remote Access Server** role and the **Web Application Proxy** role service.

Configure certificates for the Web Application Proxy server

1. On LON-DC1, open a Microsoft Management Console, and then add the **Certificates** snap-in for the **Local Computer**.
2. From the Personal folder, export the **adfs.adatum.com** certificate and use the following export options:
 - **Yes, export the private key**
 - File format: **Personal Information Exchange – PKCS #12 (.PFX)**
 - Password: **Pa\$\$w0rd**
 - File name: **C:\adfs.pfx**
3. On LON-SVR2, open a Microsoft Management Console, and then add the **Certificates** snap-in for the **Local Computer**.
4. From the Personal folder, import the **lon-dc1.adatum.com** certificate and use the following import options:
 - File name: **\\LON-DC1\c\$\adfs.pfx**
 - Password: **Pa\$\$w0rd**
 - Certificate store: **Personal**

Configure the proxy to allow access to an internal web application

5. On LON-SVR2, run the **Remote Access Management** console, and then click **Publish**.
6. Publish **https://lon-svr3.adatum.com/AdatumTestApp/** and use the **adfs.adatum.com** certificate.

Lab: Planning and Implementing AD FS Infrastructure

Scenario

A. Datum Corporation has established 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 partners and customers can access only the resources to which they require access. Another requirement is that implementing the solution should not increase the workload significantly for the operations team.

To meet these business requirements, A. Datum has decided 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 has entered into a partnership with another company, Trey Research, whose users must also be able to access the same application.

Objectives

After completing this lab, you will be able to:

- Design the AD FS deployment.
- Configure prerequisite components for AD FS.
- Deploy AD FS for internal users.
- Deploy AD FS for a partner organization.
- Deploy the Web Application Proxy.

Lab Setup

- Estimated Time: 75 minutes
- Virtual machines: 20414C-LON-HOST1, 20414C-LON-DC1.
- 20414C-LON-SVR2, 20414C-LON-SVR3.
- 20414C-LON-CL1, 20414C-TREY-DC1
- User names: **Adatum\Administrator**, **TreyResearch\Administrator**
- Password: **Pa\$\$w0rd**

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

1. On LON-HOST1, start **Hyper-V Manager**.
2. In Hyper-V® Manager, click **20414C-LON-DC1**, and then in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Sign in by using the following credentials:
 - User name: **Adatum\Administrator**
 - Password: **Pa\$\$w0rd**

5. Repeat steps 2 to 4 for **20414C-LON-SVR2**, **20414C-LON-SVR3**, **20414C-LON-CL1**, and **20414C-TREY-DC1**:
 - Sign in to **20414C-LON-SVR2** and **20414C-LON-SVR3** as **Adatum\Administrator**.
 - Do not sign in to **20414C-LON-CL1** until directed to do so.
 - On **20414C-TREY-DC1**, sign in as **TreyResearch\Administrator** with the password **Pa\$\$w0rd**.

Exercise 1: Designing the AD FS Deployment

Scenario

The security group at A. Datum has identified the following requirements for the AD FS deployment:

- Only users from the A. Datum's research and managers departments should be able to access the AD FS-protected applications.
- Members of the research department should be identified with the HighSecurity role in the AD FS application.
- Members of the managers department should be identified with the Managers role in the AD FS application.
- Users from the executives group at Trey Research should have access to the AD FS-protected application. Only these users should be identified with the ExternalSecure role in the AD FS application.
- Users that work from a home computer must be able to access AD FS-protected applications by using their AD DS credentials.

The main tasks for this exercise are as follows:

1. Choose the appropriate AD FS deployment scenario.
2. Identify the number of AD FS servers that the deployment requires, and the role and location of each of the federation servers that you will deploy.

► Task 1: Choose the appropriate AD FS deployment scenario

Based on the requirements identified, determine whether A. Datum requires a web SSO or a federated web SSO deployment to enter into a partnership with Trey Research.

► Task 2: Identify the number of AD FS servers that the deployment requires, and the role and location of each of the federation servers that you will deploy

1. Identify how many AD FS servers you will require for the deployment of AD FS in A. Datum and Trey Research.
2. Determine the role and location of each of the federation servers that you will deploy.

Results: In this exercise, you should have identified the appropriate Active Directory® Federation Services (AD FS) deployment scenario to use to meet the defined requirements. You also should have identified the number of AD FS servers required for the deployment, their locations, and the role of each AD FS server in the deployment.

Exercise 2: Configuring Prerequisite Components for AD FS

Scenario

A. Datum has deployed a WIF application as a proof-of-concept application for the AD FS deployment. You must prepare the prerequisite components for the AD FS deployment at A. Datum by configuring DNS name resolution, certificates, and certificate trusts.

The main tasks for this exercise are:

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 then verify application access.

► Task 1: Configure DNS forwarders and add DNS records

1. On LON-DC1, create a new conditional forwarder for the **TreyResearch.net** domain, by using the DNS server IP address of **172.16.10.10**. Then, create a new host record for **adfs.adatum.com** pointing to **172.16.0.10**.
2. On TREY-DC1, create a new conditional forwarder for the **Adatum.com** domain, by using the DNS server IP address of **172.16.0.10**. Then, create a new host record for **adfs.treyresearch.net** pointing to **172.16.10.10**.

► Task 2: Exchange root certificates to enable certificate trusts

1. On LON-DC1, copy **TREY-DC1.TreyResearch.net_TreyResearchCA.crt** from **\\TREY-DC1.treyresearch.net\certenroll** to the **Documents** folder.
2. Create a new Microsoft Management Console (MMC), and then add the **Group Policy Management Editor**.
3. Edit the **Default Domain Policy** Group Policy Object (GPO), and then import the copied root certificate to the **Trusted Root Certification Authorities** folder.
4. On TREY-DC1, copy the **LON-DC1.Adatum.com_AdatumCA.crt** from **\\LON-DC1.Adatum.com\certenroll** to the **Documents** folder.
5. Create a new MMC, and then add the **Certificates** snap-in focused on the local computer.
6. Import the copied root certificate to the **Trusted Root Certification Authorities** folder.

► Task 3: Request and install a certificate for the web server

1. On LON-SVR3, open the Internet Information Services (IIS) Manager console.
2. Request a new domain certificate for the server by using the following parameters:
 - Common name: **LON-SVR3.adatum.com**
 - Organization: **A. Datum**
 - Organizational unit: **IT**.
 - City/locality: **London**
 - State/province: **England**
 - Country/region: **GB**
 - Friendly name: **LON-SVR3.adatum.com**
3. Request the certificate from **AdatumCA**.

► **Task 4: Bind the certificate to the claims-aware application on the web server, and then verify application access**

1. On LON-SVR3, in IIS, create a new HTTPS site binding, and then select the newly created certificate.
2. On LON-DC1, open Internet Explorer®, and then connect to **https://lon-svr3.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.

► **Task 5: Prepare a certificate template and install a certificate for the Trey Research AD FS server**

Prepare the certificate template

1. On Trey-DC1, duplicate the Web Server template and use **Trey Research Web Server** as the new template name.
2. Set the permissions on the new template so that authenticated users can enroll for a certificate by using the new template.
3. Issue the new template to make it available for enrollment.

Request the certificate

- On Trey-DC1, enroll for a new certificate by using the **Trey Research Web Server** template. Specify a common name of **adfs.treyresearch.net** and mark the private key as exportable.

Results: In this exercise, you should have configured Domain Name System (DNS) forwarding to enable name resolution between A. Datum and Trey Research. In addition, you should have exchanged root certificates between the two organizations, and installed and configured a web certificate on the application server.

Exercise 3: Deploying AD FS for Internal Users

Scenario

In this scenario, you will deploy AD FS for internal users only. To implement AD FS, you must install and configure the AD FS server role and configure the claim provider trust. You will configure the WIF application to trust the AD FS deployment, and then configure a relying party trust. Finally, you will configure the claims rules, and then verify that only approved internal users can access the application.

The main tasks for this exercise are:

1. Installing and configuring the AD FS server role.
2. Configuring the claim provider trust.
3. Configuring the WIF application to trust the AD FS deployment.
4. Configuring a relying party trust and claims rules.
5. Verifying that approved internal users can access the application.

► **Task 1: Install and configure the AD FS server role**

1. On LON-DC1, click the Windows PowerShell shortcut on the taskbar. At the Windows PowerShell prompt, run the **Add-KdsRootKey -EffectiveTime (Get-Date).AddHours(-10)** command.
2. Close the Windows PowerShell window.

3. On LON-DC1, in Server Manager, add the **Active Directory Federation Services** server role.
4. On LON-DC1, run the AD FS Federation Server Configuration Wizard by using the following parameters:
 - Create a new Federation Service.
 - Create a stand-alone deployment.
 - Use the **adfs.Adatum.com** certificate.
 - Choose a display name of **adfs.Adatum.com**.
 - Create a managed service account named **ADFS**.
 - Use the Windows Internal Database.
5. On LON-DC1, open Windows PowerShell, and then use the **set-ADFSProperties – AutoCertificateRollover \$False** command to enable modification of the assigned certificates.
6. In the AD FS Management console, add the **adfs.adatum.com** certificate as a new token-signing certificate.
7. Make the new certificate the primary certificate and remove the old certificate.

► **Task 2: Configure the claim provider trust**

1. On LON-DC1, in the AD FS Management console, navigate to **Claims Provider Trusts**, highlight the **Active Directory** store, and then navigate to **Edit Claim Rules**.
2. In the **Edit Claim Rules for Active Directory** dialog box, on the **Acceptance Transform Rules** tab, launch the Add Transform Claim Rule Wizard, and then complete the wizard with the following settings:
 - 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 section, select the following values:
 - E-mail addresses to **E-mail Address**
 - User-Principal-Name to **UPN**
 - Display-Name to **Name**

► **Task 3: Configure the WIF application to trust the AD FS deployment**

1. On LON-SVR3, 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 WIF sample application by pointing to **C:\inetpub\wwwroot\AdatumTestApp\web.config**.
 - Specify an **Application URI** by typing **https://lon-svr3.adatum.com/AdatumTestApp/**.
 - Select **Use an Existing STS**, and then enter the path **https://adfs.adatum.com/federationmetadata/2007-06/federationmetadata.xml**.
 - Select **No Encryption**.

► **Task 4: Configure a relying party trust and claims rules**

1. On LON-DC1, 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-svr3.adatum.com/adatumtestapp**.
 - Specify a **DisplayName** 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**.
3. In the **Edit Claim Rules for Adatum Test Application** properties dialog box, choose to add a rule on the **Issuance Transform Rules** tab.
4. 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**.
5. Create three more rules to pass through **Email Address, UPN, and Name**.

► **Task 5: Verify that approved internal users can access the application**

1. Sign in to LON-CL1 as **Adatum\Brad** with the password **Pa\$\$w0rd**.
2. On LON-CL1, click to the desktop, open Internet Explorer, and then connect to **https://lon-svr3.adatum.com/AdatumTestApp/**.
3. Verify that you can access the application.
4. If you are prompted for credentials, type **Adatum\Brad** with the 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.
5. Close Internet Explorer.

Results: In this exercise, you should have installed and configured the AD FS server role on LON-DC1. You should also have configured the claim provider and relying party trusts, and then configured the necessary claims rules. Finally, you should have configured the Windows Identity Foundation (WIF) application to trust the AD FS deployment, and then verified that approved internal users could access the application.

Exercise 4: Deploying AD FS for a Partner Organization

Scenario

In the second part of the deployment scenario, you must enable Trey Research users to access the web application. To meet the security requirement of preventing direct access to your AD FS server from the Internet, you need to deploy and configure a federation proxy server. You must configure the integration of AD FS at Trey Research with AD FS at A. Datum, and then verify that Trey Research users can access the application. In addition, you should confirm that you can configure access based on user groups.

The main tasks for this exercise are as follows:

1. Install and configure the AD FS server role on Trey-DC1.
2. Add a claims provider trust for the TreyResearch.net AD FS server.
3. Configure a relying party trust on TREY-DC1 for the A. Datum claims-aware application.
4. Verify access to the A. Datum test application for Trey Research users.
5. Configure claim rules for the claim provider trust and the relying party trust to allow access only for a specific group.
6. Verify restrictions and accessibility to the claims-aware application.

► Task 1: Install and configure the AD FS server role on Trey-DC1

1. On Trey-DC1, click the Windows PowerShell shortcut on the taskbar. At the Windows PowerShell prompt, run the **Add-KdsRootKey –EffectiveTime (Get-Date).AddHours(-10)** command.
2. Close the Windows PowerShell window.
3. On Trey-DC1, in Server Manager, add the **Active Directory Federation Services** server role.
4. On LON-DC1, run the AD FS Federation Server Configuration Wizard by using the following parameters:
 - Create a new Federation Service.
 - Create a stand-alone deployment.
 - Use the **adfs.treyresearch.net** certificate.
 - Choose a service name of **adfs.treyresearch.net**.
 - Create a managed service account named **ADFS**.
 - Use the Windows Internal Database.
5. On LON-DC1, open Windows PowerShell, and then use the **set-ADFSProperties –AutoCertificateRollover \$False** command to enable modification of the assigned certificates.
6. In the AD FS Management console, add the **adfs.treyresearch.net** certificate as a new token-signing certificate. Verify that the certificate has a subject of **CN= adfs.treyresearch.net**.
7. Make the new certificate the primary certificate and remove the old certificate.

► Task 2: Add a claims provider trust for the TreyResearch.net AD FS server

1. On LON-DC1, in the AD FS Management console, navigate to **Trust Relationships**, navigate to **Claims Provider Trusts**, and then click **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://adfs.treyresearch.net** as the data source.
 - In **Display Name**, type **adfs.treyresearch.net**.

3. In the **Edit Claim Rules for the adfs.treyresearch.net** properties dialog box, enter the following values:
 - On the **Acceptance Transform Rules** tab, click **Add Rule**.
 - 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 **Pass through all claim values**.
 - Complete the rule.
4. On LON-DC1, run the following command in Windows PowerShell:

```
Set-ADFSClaimsProviderTrust -TargetName "adfs.treyresearch.net" -
SigningCertificateRevocationCheck None
```

5. Close the Windows PowerShell window.



Note: You should disable certificate revocation checking in test environments only. In a production environment, you should enable certificate revocation checking.

► Task 3: Configure a relying party trust on TREY-DC1 for the A. Datum claims-aware application

1. On TREY-DC1, in the AD FS Management console, open the Add Relying Party Trust Wizard, and then complete it with the following settings:
 - Choose **Import data about the relying party published online or on a local network** and type **https:// adfs.adatum.com**.
 - Specify a **Display name** of **Adatum TestApp**.
 - Do not configure multi-factor authentication.
 - Choose **Permit all users to access this relying party**.
 - Accept the option to open the Edit Claim Rules for lon-dc1.adatum.com when the wizard completes.
2. In the **Edit Claim Rules for lon-dc1.adatum.com properties** dialog box, on the **Issuance Transform Rules** tab, click to add a rule with the following settings:
 - 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 the **Incoming claim type** drop-down list.
 - Choose to **Pass through all claim values**.
 - Complete the wizard.

► Task 4: Verify access to the A. Datum test application for Trey Research users

1. On TREY-DC1, open Internet Explorer, and then connect to **https://lon-svr3.adatum.com/adatumtestapp/**.
2. Select **adfs.treyresearch.net** as the home realm, and then sign in as **TreyResearch\April**, with the password **Pa\$\$w0rd**.

3. Verify that you can access the application.
4. Close Internet Explorer, and then connect to the same web site. Verify that this time you are not prompted for a home realm.



Note: You will not be prompted for a home realm again. Once users have selected a home realm and a realm authority has authenticated them, the relying party federation server issues an `_LSRealm` cookie to the user. The default lifetime for the cookie is 30 days. Therefore, to sign in multiple times, you should delete that cookie after each sign-in attempt to return to a clean state.

► Task 5: Configure claim rules for the claim provider trust and the relying party trust to allow access only for a specific group

1. On TREY-DC1, open the AD FS Management console, and then 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 **adfs.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 then 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 then select **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**, and an Incoming Claim Value of **@adatum.com**. Select **Permit Access to Users With This Incoming Claim**, and then click **Finish**.
8. Edit the Temp rule, and then copy the claim rule language into the clipboard.
9. Delete the Temp rule.
10. Create a new rule that sends claims by using a custom rule named **ADatum User Access Rule**.
11. Click 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 6: Verify restrictions and accessibility to the claims-aware application

1. On TREY-DC1, open Internet Explorer, and then connect to **https://lon-svr3.adatum.com/adatumtestapp/**.
2. Verify that **TreyResearch\April** no longer has access to the A. Datum test application.
3. Clear the browsing history in Internet Explorer.
4. Connect to **https://lon-svr3.adatum.com/adatumtestapp/**.
5. On the **Sign In** page, click **adfs.treyresearch.net**.
6. Verify that **TreyResearch\Morgan** does have access to the A. Datum test application. Morgan is a member of the Production group.

Results: In this exercise, you should have configured a claims provider trust for Trey Research on Adatum.com and a relying party trust for Adatum on Trey Research. You should also have verified access to the A. Datum claims-aware application and configured the application to restrict access from Trey Research to specific groups.

Exercise 5: Deploy the Web Application Proxy

Scenario

The next step in this scenario is to enable Trey Research users to access the web application while meeting the security requirement of preventing direct access to your AD FS server from the Internet. You must deploy and configure a federation proxy server. You need to configure the integration of AD FS at Trey Research with AD FS at A. Datum, and then verify that Trey Research users can access the application. You also want to confirm that you can configure access based on user groups.

The main tasks for this exercise are as follows:

1. Configure certificates for the Web Application Proxy server.
2. Install the Web Application Proxy role.
3. Configure access to an internal website.
4. Verify access to the internal web site from the client computer.
5. To prepare for the next module.

► Task 1: Configure certificates for the Web Application Proxy server

1. On LON-DC1, open a Microsoft Management Console, and then add the **Certificates** snap-in for the **Local Computer**.
2. From the Personal folder, export the **adfs.adatum.com** certificate and use the following export options:
 - **Yes, export the private key**
 - File format: **Personal Information Exchange – PKCS #12 (.PFX)**
 - Password: **Pa\$\$w0rd**
 - File name: **C:\adfs.pfx**
3. On LON-SVR1, open a Microsoft Management Console, and then add the **Certificates** snap-in for the **Local Computer**.

4. From the Personal folder, import the **adfs.adatum.com** certificate and use the following import options:
 - File name: **\\LON-DC1\c\$\adfs.pfx**
 - Password: **Pa\$\$w0rd**
 - Certificate store: **Personal**
- ▶ **Task 2: Install the Web Application Proxy role**
 - On LON-SVR2, in Server Manager, add the **Remote Access Server** role and the **Web Application Proxy** role service.
- ▶ **Task 3: Configure access to an internal website**
 1. On LON-SVR2, run the **Remote Access Management** console, and then click **Web Application Proxy**.
 2. Click **Run the Web Application Proxy Configuration Wizard**.
 3. Configure the wizard with the following settings:
 - Federation Server: **adfs.adatum.com**
 - User name: **Administrator**
 - Password: **Pa\$\$w0rd**
 - AD FS Proxy certificate: **adfs.adatum.com**
 4. Publish **https://lon-svr3.adatum.com/AdatumTestApp/** and use the **adfs.Adatum.com** certificate.
- ▶ **Task 4: Verify access to the internal web site from the client computer**
 1. On LON-CL1, sign in as **Adatum\Administrator** and update the HOSTS file to point lon-svr3.adatum.com to 172.16.0.13.
 2. Go to **https://lon-svr3.adatum.com/adatumtestapp/** in Internet Explorer. Accept the security certificate warning and continue to the site.
 3. Authenticate by using **TreyResearch\Morgan** with the password **Pa\$\$w0rd**.
 4. After the application loads, close Internet Explorer.
- ▶ **Task 5: To prepare for the next module**

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

 1. On the host computer, start Hyper-V® Manager.
 2. In the **Virtual Machines** list, right-click **20414C-LON-DC1**, and then click **Revert**.
 3. In the **Revert Virtual Machine** dialog box, click **Revert**.
 4. Repeat steps 2 and 3 for **20414C-LON-SVR3**, **20414C-TREY-DC1**, and **20414C-LON-CL1**.

Results: In this exercise, you should have installed the Web Application Proxy role, configured certificates for the proxy server, configured access to an internal web site, and verified access to the internal web site.

Question: When does A. Datum's AD FS server act as an account federation server, and when does it act as a resource federation server?

Question: How could a proxy improve the security of A. Datum's AD FS deployment?

Module Review and Takeaways

Review Question

Question: What is the difference between a claims provider trust and a relying party trust?

Tools

- AD FS Management console. This console is available as part of the Remote Server Administration Tools or on AD FS servers.
- Public key infrastructure (PKI) certification authority. The PKI management tools are available as part of the Remote Server Administration Tools or on Certification Authority servers.
- Windows PowerShell. Windows PowerShell is available on Windows-based computers. You can add modules to manage specific roles such as AD FS.
- Server Manager. Server Manager is available on Windows-based servers. It is installed by default and can manage the local server and remote servers.

Module 12

Planning and Implementing Data Access for Users and Devices

Contents:

Module Overview	12-1
Lesson 1: Planning and Implementing DAC	12-2
Lab A: Implementing DAC and Access-Denied Assistance	12-13
Lesson 2: Planning Workplace Join	12-22
Lesson 3: Planning Work Folders	12-27
Lab B: Implementing Work Folders	12-33
Module Review and Takeaways	12-39

Module Overview

The Windows Server® 2012 operating system introduces new features that enhance access control for file-based and folder-based resources, including features for accessing work data from various locations. These features—Dynamic Access Control (DAC), Workplace Join, and Work Folders—extend traditional access control. These three features also enable administrators to use claims, resource properties, policies, and conditional expressions to manage access. By using these features, administrators can support the Bring Your Own Device (BYOD) capabilities. These features also enable users to have more flexible data access. In this module, you will learn how to plan and implement DAC, Workplace Join, and Work Folders.

Objectives

After completing this module, you will be able to:

- Plan and implement DAC.
- Plan and implement Workplace Join.
- Plan and implement Work Folders.

Lesson 1

Planning and Implementing DAC

The Windows Server 2012 operating system introduces DAC for enhancing access control for file-based and folder-based resources. DAC extends access control based on New Technology File System (NTFS) attributes by enabling you to use claims, resource properties, rules, and conditional expressions to manage access.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe DAC.
- Identify prerequisites for DAC.
- Access resources with DAC applied.
- Describe identity, claims, resource properties, and a Central Access Policy.
- Plan for a Central Access Policy.
- Create Central Access Rules and Policies.
- Describe file classifications and how to plan for them.
- Describe file access auditing and how to plan for it.
- Describe access-denied assistance and how to plan for it.
- Implement DAC.

DAC Overview

In the past, we used NTFS file permissions to help secure access to file server-based data files.

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

DAC combines multiple criteria into access decisions. This augments the NTFS file system access control list (ACL) so that users need to satisfy Share permissions, the NTFS file system ACL, and the Central Access Policy to gain access to a file. However, DAC also can work independently from NTFS file system permissions.

DAC is a new access control mechanism in Windows Server 2012 for file system resources that combines multiple criteria into access decisions

- You can use DAC to help to achieve four goals:
 - Establish a central access policy for file access
 - Auditing for compliance and analysis
 - Protecting sensitive information
 - Access-denied remediation
- Most common reasons for implementing DAC:
 - NTFS does not fit your corporate environment
 - To reduce security group complexity
 - Adhere to compliance regulations
 - To protect sensitive information
 - To upgrade to access control based on attributes

DAC provides:

- Data classification. You can use automatic and/or manual classification of files to tag data across your organization.
- Access control to files. Central access policies enable you to define who can access particular files. For example, you can define who can access human resources (HR) information.
- Auditing of access to files. DAC provides central audit policies that you can use for compliance reporting and for security analysis. For example, you can identify who accessed sensitive HR data files.
- Optional Active Directory Rights Management Services (AD RMS) protection integration. You can implement automatic RMS encryption for Microsoft® Office files that contain sensitive information. For example, you can configure RMS to encrypt all documents containing HR information.

Reasons to Implement DAC

The most common reasons for implementing DAC are to reduce security group complexity, comply with regulations, and protect sensitive information. In addition, with DAC, you can extend the functionality of an existing model for access control management. Most companies use NTFS and Share permissions to implement access control for file and folder resources. In most cases, NTFS is sufficient, but it does not work in all situations. For example, you cannot use an NTFS ACL to protect a resource on a file server by stipulating that a user must be member of two groups simultaneously to access the resource.

In general, when you want to use more specific information for implementing access control and authorization, you cannot use the usual methods. NTFS and Share permissions use only user or group objects. If you want to implement more complex access control scenarios, like using attributes for access control, you should use DAC.

Usage Scenarios

You can use DAC to help address four types of usage scenarios:

- Central access policy for access to files. Enables you to define safety net policies that reflect business and regulatory compliance.
- Auditing for compliance and analysis. Enables you to target auditing across your file servers for compliance reporting and security analysis.
- Protecting sensitive information. Enables you to identify and protect sensitive information while it is stored in a Windows Server 2012 environment and after it leaves that environment.
- Access-denied remediation. Enables you to customize and thereby improve the user experience if the user is denied access to files. This can reduce calls to the helpdesk.

Prerequisites for DAC

Before you implement DAC, you must ensure that your organization's network infrastructure meets certain prerequisites. To implement claims-based authorization for resource access, you must implement the following:

- Windows Server 2012, with the File Server Resource Manager (FSRM) feature enabled and installed on the file server that hosts the resources that DAC is protecting. The file server hosting the share must be a Windows Server 2012 or Windows Server 2012 R2 file server. This enables the file server to read claims and device authorization data from a Kerberos authentication ticket, to translate those security identifiers (SIDs) and claims 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. This is necessary to store the central definitions for resource properties and policies, which must be accessible by a client computer running the Windows® operating system in the user's domain.
- Optionally, one Windows Server 2012 domain controller. If you do not need to use user claims for security groups, you should include at least one Windows Server 2012 domain controller in the user domain that is accessible by the file server. Then the file server can retrieve claims on behalf of the user.
- Optionally, Windows® 8. If you are using device claims, then all client computers must have Windows 8 or newer installed. This is because device claims work for Windows 8-based devices and Windows 8.1-based devices only.
- Windows Server 2012 domain controllers in each domain when using claims across a forest trust. Although a Windows Server 2012 domain controller is required if you are using user claims, there is no requirement for a Windows Server 2012 domain and forest functional level, unless you want to use claims across forest trust. This means that you can also have domain controllers on Windows Server 2008 and Windows Server 2008 R2 with forest functional level on Windows Server 2008.

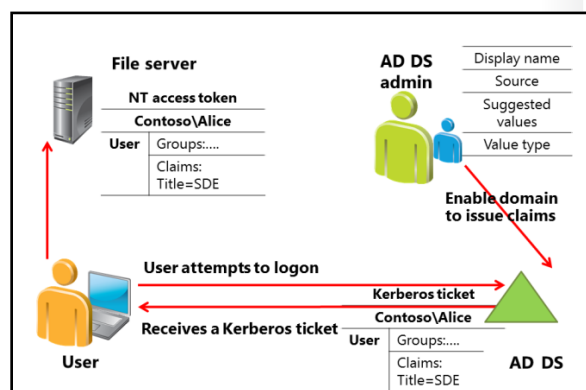
To deploy DAC, you need these technologies:

- A domain controller running on Windows Server 2012 or newer
- A file server running Windows Server 2012 or newer with FSRM
- Windows 8 or Windows 8.1 desktop (for device claims)

Accessing Resources with DAC

DAC is the new authorization and auditing mechanism that provides required extensions to Active Directory® Domain Services (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 V5 authentication protocol Key Distribution Center (KDC).

When you use the NTFS file system to manage access control, the user's access token contains the user's SID and the SIDs of all the groups of which the user is a member. When the user tries to access the resource, the ACL of that resource is



evaluated. If at least one SID from the user's token matches the SID on the ACL, the appropriate rights are assigned to the user.

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

The Windows Server 2012 KDC contains Kerberos protocol enhancements that are required to transport the claims within a Kerberos ticket, including the use of compound identity. Also, Windows Server 2012 KDC includes an enhancement to support *Kerberos armoring*. Kerberos armoring is an implementation of Flexible Authentication Secure Tunneling, which provides a protected channel between the Kerberos client and the KDC.

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

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

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

What Are Identity, Claims, Resource Properties and Central Access Policy?

To plan and implement DAC, you must understand some fundamental concepts. These concepts include identity, claims, Resource Property, and Central Access Policy

What Is Identity?

Identity is information that a trusted source provides about an entity. This information is considered authoritative because the source is trusted. Older versions of Windows Server, such as Windows Server 2008, use the user and group account SIDs to represent the identity of a user or

- Identity is information that a trusted source provides about an entity
- Claims are statements that AD DS makes about specific users or computer objects
- Resource properties are the information you provide about the resource
- Central Access Policy contains one or more Central Access Rules that determine applicability and permissions

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 provides a token with the SID of the security principal and the SIDs of all the groups of which the principal is a member. The domain controller claims the user's SID is valid and should be used as the identity of the user. Because all domain members trust the domain controller, they treat the response as authoritative.

However, 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 the source of the information.

What Is a Claim?

Claims provide information from a trusted source about an entity. Windows Server 2008 and Windows Server 2003 use claims in Active Directory Federation Services (AD FS). These claims are statements made about users, which both partners in an AD FS federation understand.

Some examples of claims are the user's department and security clearance. These claims state something about a specific entity. Specifically, claims state the value of a particular attribute of a user or computer object. An entity can contain more than one claim. When configuring resource access, you can use any combination of those claims to control access to resources.

Windows Server 2012 introduces two new types of claims:

- **User claim.** A *user claim* is information that a Windows Server 2012 domain controller provides about a user. Windows Server 2012 domain controllers can use most AD DS user attributes as claim information. This provides you with many ways to configure and use claims for access control.
- **Device claim.** A *device claim* is information that a Windows Server 2012 domain controller provides about a device that is represented by a computer account in AD DS. As with a user claim, a device claim, often called a *computer claim*, can use most of the AD DS attributes that are applicable to computer objects.

What Is a Resource Property?

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


As with user or device claims, you have to define the attributes of the resource that you want to use. You do this by configuring the resource properties. You manage resource properties in the resource properties container, which is displayed in the DAC node in the Active Directory Administrative Center.

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

When evaluating file authorization and auditing, the Windows operating system uses the values in these properties and the values from user and device claims.

What Is a Central Access Policy?

The Central Access Policy is a feature in Windows Server 2012 that enables you to create a policy that applies to one or more file servers. You create this policy in the Active Directory Administrative Center, store it in AD DS, and then apply it by using Group Policy Objects (GPOs). Central Access Policy contains one or more Central Access Policy rules. Each rule contains settings that determine applicability and permissions.

 **Note:** Before you create a Central Access Policy, you must create at least one Central Access Rule. A *Central Access Rule* defines all parameters and conditions that control access to specific resources. Central access policies have three configurable elements:

- **Name.** For each Central Access Rule, you should provide a meaningful name.
- **Target resources.** Defines the data to which the policy applies by specifying an attribute and its value. For example, a particular central policy might apply to any data classified as Sensitive.
- **Permissions.** A list of one or more access control entries (ACEs) that define who can access the data. For example, you can specify Full Control Access to a user with the attribute `EmployeeType` populated with the value `FTE` for a full time employee. This is the key component of each Central Access Rule. You can combine any group conditions that you place in a Central Access Rule. Also, you can also set permissions as Current or Proposed, for staging purposes.

After you have configured one or more Central Access Rules, you place these rules in a Central Access Policy. Then the rules are applied to the resources.

Central Access Policies enhance but do not replace discretionary access control lists (DACs) 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 on the file restricts that user's access, the user cannot obtain access to the file. Likewise, if the Central Access Policy allows access but the DACL does not, the user cannot obtain access to the file.

Before you implement a Central Access Policy, you must follow these steps:

1. Create claims and connect them with attributes on user or computer objects.
2. Create file property definitions.
3. Create one or more Central Access Rules.
4. Create a Central Access Policy object and place rules in it.
5. Use Group Policy to deploy the policy to file servers. By doing this, you make file servers aware that a Central Access Policy exists in AD DS.
6. On the file server, apply that policy to a specific shared folder. You can also use the Microsoft Data Classification Toolkit to apply central policies automatically across multiple file servers, and to generate a report on which policies are applied on which shares.

Planning for Central Access Policy

Implementing Central Access Policy is not mandatory in all DAC scenarios, but we recommend it for consistent configuration of access control across all file servers. If you choose to implement Central Access Policy, you should plan carefully before deployment.

When planning Central Access Policy, you must identify and understand the business requirements for implementing both Central Access Policy and DAC. Follow these steps when planning central access policy:

When planning Central Access Policy, you should:

- Identify relevant business cases
- Identify resources to be protected
- Understand business requirements
- Translate business requirements to conditional expressions
- Define claim types, resource properties, and rules

1. Identify resources that you want to protect:
 - If all of the resources are on one file server or in one folder, then you may not need to implement Central Access Policy. Instead, you can configure conditional access on the folder's ACL.
 - If resources are distributed across multiple servers or folders, you should deploy Central Access Policy.
2. Define criteria for protection. Usually, these are defined by business requirements. For example:
 - All documents that have their confidentiality property set to high must be available to managers only.
 - HR documents from each country should be accessible to HR staff from the same country only.
 - Only full-time staff should be able to access technical documentation from previous projects.
3. Translate the policies that you require into expressions. In the case of DAC, expressions are attributes that are associated with both the resources, such as files and folders, and the user or device that wants access to these resources.
4. Break down the expressions that you created and determine which security groups, claim types, resource properties, and device claims you must create to deploy your policies. In other words, you must identify the attributes for access to filtering.

Demonstration: Creating Central Access Rules and Policies for DAC

This demonstration shows how to create Central Access Rules and Policies.

Demonstration Steps


1. On LON-DC1, open **Server Manager**.
2. In the Active Directory Administrative Center, create claims for the **Department** and **employeetype** attributes.
3. Enable **Resource Type** for department.
4. Create a Central Access Rule to enable members of the IT group to access resources if the user department attribute matches the Resource Department.
5. Create a Central Access Policy.

Planning for File Classifications

You should include file classifications in your planning for DAC implementation. Although file classifications are not mandatory for DAC, they can enhance the automation of the entire process. For example, you may require security-critical documents to be accessible to top management only, and classified with the attribute Confidentiality set to high, regardless of the server on which the documents exist. In this case, you should consider how to identify these documents and how to classify them appropriately.

The File Classification Infrastructure (FCI) uses location-based classification, such as For this Folder, Structure Confidentiality is High. Also, you can use

- Resource Property definitions are defined in AD DS
 - Resource Property definitions can be used during file classifications
 - File classifications can be run automatically
- When planning classification implementation:
- Identify classifications
 - Determine method for classification
 - Determine schedule
 - Perform review



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

To implement DAC effectively, you must have well-defined claims and resource properties. Although attributes define claims for a user or a device, you typically create and define resource properties manually. File classifications enable administrators to define automatic procedures for defining a desired property on the file, based on conditions specified in a classification rule. For example, you can set the Confidentiality property to high on all documents with contents that contain the word "secret." Then you can use this property in DAC to specify that only employees with 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 two tasks, you can automate file and folder maintenance tasks, such as cleaning up old data or protecting sensitive information.

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

You can use file classification to perform the following actions:

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

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

Planning File Access Auditing

In Windows Server 2012, you can implement file access auditing together with DAC to use new Windows security auditing capabilities. By using conditional expressions, you can configure auditing to be implemented in specific cases only. For example, you may want to audit attempts to open shared folders by users located in countries other than the country where the shared folder is located.

With Global Object Access Auditing, you can define computer system access control lists (SACLs) per object type for either the file system or registry. The specified SACL is then applied automatically to every object of that type.

File access auditing features:

- Track changes to user and machine attributes
- Obtain more information from user logon events
- Provide more information from Object Access Audit policy
- Track changes to Central Access Policies, Central Access Rules, and claims
- Track changes to file attributes

You can use a Global Object Access Auditing Policy to enforce the object access audit policy for a computer, file share, or registry without configuring and propagating conventional SACLs. Configuring and propagating SACLs is a more complex administrative task. It is difficult to verify, particularly if you must verify to an auditor that security policy is being enforced. Instead, 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 diagnostics. For example, setting a Global Object Access Auditing policy to log all activity for a specific user and enabling the Access Failures audit policies in a resource can help administrators quickly identify which object in a system is denying a user access.

You should make an audit plan before you implement any auditing. In an auditing plan, you 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 live in AD DS, and, therefore, you can track changes to their attributes using Directory Service Access Auditing.
- Retrieving more information from user logon events. In Windows Server 2012, user logon event (4624) contains information about the attributes of the user who logged on. Also, you can use audit log management tools to correlate user logon events with object access events, and enable event filtering based on both file attributes and user attributes.
- Providing more information from object access auditing. In Windows Server 2012, file access events (4656, 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. These objects define the central policy that you can use to control access to critical resources. Tracking changes to these objects could be important for the organization. Because these objects are stored in AD DS, you can audit them by using Directory Service Access Auditing. This is the same method you use for auditing any other securable object in AD DS.
- Tracking changes to file attributes. File attributes determine which Central Access Policy applies to the file. Potentially, a change to the file attributes can 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 Audit policy for file systems. Event 4911 was introduced to differentiate this event from other Authorization Policy Change events.

Planning Access-Denied Assistance

Access-denied assistance helps end users determine the reason why they cannot access a resource. Also, it helps you to diagnose a problem and initiate a resolution. Windows Server 2012 enables you to customize access-denied messages and allows users to request access without contacting the helpdesk or Information Technology (IT) team. In combination with DAC, access-denied assistance informs the file administrator of any user and resource claims, enabling the file administrator to make educated decisions regarding adjustments to policies or user attributes. An example of an adjustment to a user attribute would be if a department name is written as "HR" instead of "Human Resources".

When planning access-denied assistance, plan carefully for the following:

- Message that users see when they are denied access
- Email text that users generate to request access
- Recipients of emails requesting access
- Target operating systems – access denied assistance works only with Windows 8 or newer

When planning for access-denied assistance, you should include the following:

- Carefully plan the message wording that users see when they try to access resources to which they do not have access rights. The message should be easy to understand and have specific instructions for remediation, if applicable.
- Determine whether users can send a request for access via email. If so, you have the option of configuring text to be added to the end of their email messages.
- Decide who should receive access request email messages. You can choose folder owners, file server administrators, or any other specified recipient. You should ensure that messages are always directed to the right person. If you have a helpdesk tool or monitoring solution that allows emails, you can also direct those emails to generate user requests in your helpdesk solution automatically.
- Plan the target operating systems. Access-denied assistance works with Windows 8 or Windows Server 2012 and newer operating systems.

Implementing DAC

When deciding whether to implement DAC, consider your organization's needs. In some situations, NTFS file permissions and, optionally, AD RMS can meet your file access management goals. However, in certain circumstances, DAC can help you address more complex business needs. Consider the following scenario:

The Research Department at Tailspin Toys wants to ensure that archived research files are only available to users who are members of the Research Department. Specifically, these data files should be accessible with read-only permissions to Research Department members in the same country only. Additionally, a central administrative group must have read-only permissions on files from all countries.

You could achieve these objectives with DAC by following these steps:

1. Plan the DAC implementation:
 - a. Define the access policy:
 - Research files must be read-only for members of the Research Department.
 - Members of the Research Department must access documents in their own country only.
 - Only Research Administrators should have Write access.
 - An exception will be allowed for members of the Research_Exceptions group.
 - The Research_Exceptions group will have Read access.
 - b. Express the access policy in Windows Server 2012 constructs:
 - Targeting: Resource.Department Contains Research
 - Access rules:
 - Allow read User.Country=Resource.Country AND User.department=Resource.Department
 - Allow Full control User.MemberOf(ResearchAdmin)
 - Exception: Allow read memberOf(Research_Exceptions)

You can implement DAC for complex business needs by:

- Defining the access policy
- Expressing the access policy in Windows Server 2012 constructs
- Determining the file properties required to support the policy
- Determining the claims types and groups required to support the policy
- Determining the servers on which to apply the policy

You can configure DAC by:

- Creating claim types
- Creating resource properties
- Configuring a Central Access Rule
- Configuring a Central Access Policy
- Targeting Central Access Policy to the file servers

- c. Determine the file properties required to support the policy:
 - Tag files with: Department and Country
 - d. Determine the claims types and groups required to support the policy:
 - Claim types:
 - Country
 - Department
 - User groups:
 - ResearchAdmin
 - Research_Exceptions
 - e. Determine the servers on which to apply the policy:
 - Apply to all research servers.
2. Configure DAC:
- a. Create claim types:
 - Department
 - Country
 - b. Create resource properties:
 - Department
 - Country
 - c. Configure a Central Access Rule:
 - Create a Research Documents rule that includes the policy determined in step one.
 - d. Configure a Central Access Policy:
 - Create a Central Access Policy called Research Policy and add the Research Documents rule to that Central Access Policy.
 - e. Target Central Access Policy to the file servers:
 - Publish the Research Policy Central Access Policy to the research file servers.

You could not achieve the preceding objectives by using only NTFS file permissions without the requirement for very complex groups, NTFS permissions, and folder structures.

Lab A: Implementing DAC and Access-Denied Assistance

Scenario

A. Datum Corporation has created a new regional sales team in Europe. As a result, branch offices have been equipped to support various regional sales functions. You are responsible for planning the network infrastructure of the branch offices. The National Sales Manager, Dan Drayton has contacted you for industry standards on access to file-based resources from the head office in London.

The research team at A. Datum performs highly confidential work, which they store on the company's file servers. The Security Department wants to ensure that these confidential files are accessible to authorized personnel only and that all access to these files is audited.

As a senior network administrator, you are responsible for addressing these security requirements by implementing DAC on the file servers. You plan to work closely with the business groups and the Security Department to identify which files must be secured and who should have access to those files. Then you plan to implement DAC based on the company requirements.

Objectives

After completing this lab, you will be able to:

- Plan and implement DAC.
- Prepare for a DAC deployment.
- Implement DAC.
- Validate and remediate DAC.

Lab Setup

Estimated Time: 70 minutes

Virtual machines	20414C-LON-DC1 20414C-LON-SVR1 20414C-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 Microsoft Hyper-V® Manager, click **20414C-LON-DC1**, and, in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Sign in using the following credentials:
 - User name: **Administrator**
 - Password: **Pa\$\$w0rd**
 - Domain: **Adatum**
5. Repeat steps two through four for 20414C-LON-SVR1 and 20414C-LON-CL1.

Exercise 1: Planning and Implementing DAC

Scenario

You must ensure the security of documents that the Research Department and managers use. Review the following proposal, and then plan your DAC implementation.

File Security Strategy	
Document Reference Number: BS1002/1	
Document Author	Ethan Rincon
Date	2 nd October
<p>Requirements Overview</p> <p>Design a file security strategy to support the following objectives:</p> <ul style="list-style-type: none"> • Only members of the Research Department should access and modify folders that belong to the Research Department. • Only managers should access documents that are classified as highly confidential. • Also, the A. Datum Security Department is concerned that users in the Managers Department are accessing files using unmanaged computers, which may not be highly secure. You must create a plan for securing the documents regardless of where they are located, and ensure that the documents can be accessed from authorized computers only. • Authorized computers for managers must be members of the security group ManagersWks. • The Support Department reports a high number of calls from users who cannot access resources. You must implement a technology that helps users understand the error messages that they receive and enables them to request access automatically. 	
<p>Additional Information</p> <ul style="list-style-type: none"> • Most of the files that these departments use are stored in shared folders dedicated to these departments, but confidential documents are occasionally saved to other shared folders. 	
<p>Proposal</p> <ul style="list-style-type: none"> • How will you design DAC to fulfill requirements for access control as described in the scenario? 	

The main tasks for this exercise are as follows:

1. Read the supporting documentation
2. Update the proposal document with your planned course of action
3. Examine the suggested proposals in the Lab Answer Key
4. Discuss your proposed solution with the class, as guided by your instructor

► **Task 1: Read the supporting documentation**

Read the documentation provided.

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

Answer the questions in the proposals section of the File Security Strategy document.

► **Task 3: Examine the suggested proposals in the Lab Answer Key**

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

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

Be prepared to discuss your proposals with the class.

Results: After completing this exercise, students should have successfully planned and implemented DAC.

Exercise 2: Preparing DAC Deployment

Scenario

The first step in implementing DAC is to configure the claims for the users and devices that access the files. In this exercise, you will review the default claims and create new claims based on the department and computer group attributes. Also, you will configure the Resource Property lists and the Resource Property definitions. Then you will use the resource properties to classify files.

The main tasks for this exercise are as follows:

1. Prepare AD DS for DAC, and review the default claim types
2. Configure claims for users and devices
3. Configure resource properties for files
4. Classify files and folders

► Task 1: Prepare AD DS for DAC, and review the default claim types

1. On LON-DC1, from Windows 8 Server Manager, open **Active Directory Users and Computers**.
2. Make a new organizational unit (OU) named **Test**.
3. Move the LON-CL1 and LON-SVR1 computer objects into the Test OU.
4. Open the Group Policy Management Console.
5. Edit the **Default Domain Controllers Policy** GPO.
6. In the Group Policy Management Editor, navigate to **Computer Configuration\Policies\Administrative Templates\System\KDC**.
7. Enable the **KDC support for claims, compound authentication and Kerberos armoring** policy setting.
8. In the Options section, click **Supported**.
9. Refresh Group Policy.
10. Open **Active Directory Users and Computers**, and then, in the Users container, create a security group named **ManagersWKS**.
11. Add LON-CL1 to the ManagersWKS group.
12. Verify that the user Aidan Delaney is a member of the Managers Department, and then verify that Allie Bellew is a member of the Research Department.
13. On LON-DC1, in Server Manager, open **Active Directory Administrative Center**.
14. In the Active Directory Administrative Center, click the **Dynamic Access Control** node.
15. Open the **Claim Types** container, and then verify that no default claims are defined.
16. Open the **Resource Properties** container, and then note that all properties are disabled by default.
17. Open the **Resource Property Lists** container, and then open the properties of the **Global Resource Property List**.
18. In the Resource Properties section, review the available resource properties.
19. Click **Cancel**.

► Task 2: Configure claims for users and devices

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 by using the following settings:
 - Source Attribute: **Department**
 - Display name: **Company Department**
3. Select the **User** and **Computer** check boxes.
4. In the Active Directory Administrative Center, in the Tasks pane, click **New**, and then click **Claim Type**.
5. Create a new claim type for computers by using the following settings:
 - Source Attribute: **description**
 - Display name: **description**
 - Clear the **User** check box
 - Select the **Computer** check box

► Task 3: Configure resource properties for files

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 the **Properties for Department** property.
4. Add **Research** as the suggested value in the **Value** and **Display name** text boxes.
5. Open the **Global Resource Property List**, and then ensure that both **Department** and **Confidentiality** are included in the list.
6. Click **Cancel**.
7. Close the Active Directory Administrative Center.

► Task 4: Classify files and folders

1. Switch to LON-SVR1.
2. On LON-SVR1, in Server Manager, use the Add Roles and Features Wizard to add the File Server Resource Manager service role.
3. Create the following folders and files:
 - Folder: C:**Docs**
 - File: C:\Docs**Doc1.txt**
 - File: C:\Docs**Doc2.txt**
 - File: C:\Docs**Doc3.txt**
4. Populate Doc1 and Doc2 with the following text:
 - **This is a secret document.**
5. Populate Doc3 with the following text:
 - **This is a document.**

6. Share the new Docs folder by using the following properties:
 - Share with: **Specific people**
 - Authenticated Users: **Read/Write**
7. Open the **File Server Resource Manager**.
8. Refresh **Classification Properties**, and then verify that the **Confidentiality** and **Department** properties are included in the list.
9. Create a Classification rule with following values :
 - Name: **Set Confidentiality**
 - Scope: **C:\Docs**
 - Classification method: **Content Classifier**
 - Property: **Confidentiality**
 - Value: **High**
 - Classification Parameters: **String "secret"**
 - Select **Re-evaluate existing property values**, and then click **Overwrite the existing value**.
 - Run the classification rule.
10. Open **File Explorer**, and then open the **Properties** files for files **Doc1.txt**, **Doc2.txt**, and **Doc3.txt**.
11. Verify values for Confidentiality. Doc1.txt and Doc2.txt should have confidentiality set to **High**.
12. Open **File Explorer**.
13. Create a folder named **C:\Research**.
14. Create a text document named **C:\Research\Research1**, and add the following text to the file:
 - **This is a research document.**
15. Share the new folder using the following properties:
 - Share with: **Specific people**
 - Authenticated Users: **Read/Write**
16. Browse to **C:\Research**, and open its properties.
17. On the **Classification** tab, set the **Department** value to **Research**.

Results: After completing this exercise, students will have prepared for DAC deployment.

Exercise 3: Implementing DAC

Scenario

The next step in implementing DAC is to configure the Central Access Rules and Policies that link the claims and property definitions. You will configure rules for DAC to address requirements from the lab scenario. After you configure DAC rules and policies, you will apply the policy to a file server.

The main tasks for this exercise are as follows:

1. Configure Central Access Policy rules
2. Create and publish the Central Access Policy
3. Apply a Central Access Policy

► Task 1: Configure Central Access Policy rules

1. Switch to LON-DC1.
2. On LON-DC1, in Server Manager, click **Tools**, and then click **Active Directory Administrative Center**.
3. In the Active Directory Administrative Center console, click **Dynamic Access Control**, and then open the **Central Access Rules** container.
4. Create a new Central Access Rule with the following values:
 - Name: **Department Match**
 - Target Resource: create a new condition as follows: **Resource-Department-Equals-Value-Research**
 - Permissions:
 - Click **Use following permissions as current permissions**
 - Remove **Administrators** group
 - Add **Authenticated Users** group
 - Assign this group the right to **Modify, Read and Execute, Read**, and **Write**, with condition **User-Company Department-Equals-Resource-Department**
5. Create another Central Access Rule with the following values:
 - Name: **Access Confidential Docs**
 - Target Resource: create new condition as follows: **Resource-Confidentiality-Equals-Value-High**
 - Permissions:
 - Click **Use following permissions as current permissions**
 - Remove **Administrators** group
 - Add **Authenticated Users** group
 - Assign this group the right to **Modify, Read and Execute, Read**, and **Write**, with the condition **User-Group-Member of each-Value-Managers**
 - Set second condition to **Device-Group-Member of each-Value-ManagersWKS**

► Task 2: Create and publish the Central Access Policy

1. On LON-DC1, in Active Directory Administrative Center, create a new Central Access Policy with the following values:
 - Name: **Protect confidential docs**
 - Rules included: **Access Confidential Docs**
2. Create another Central Access Policy with the following values:
 - Name: **Department Match**
 - Rules included: **Department Match**
3. Open the Group Policy Management Console.
4. Create a new GPO named **DAC Policy**, and link it to the **Test** OU.
5. Edit the DAC Policy and navigate to **Computer Configuration/Policies/Windows Settings/Security Settings/File System**, and then right-click **Central Access Policy**.
6. Click **Manage Central Access Policies**.

7. Click **Department Match** and **Protect confidential docs**, click **Add**, and then click **OK**.
8. Close the Group Policy Management Editor and the Group Policy Management Console.

► **Task 3: Apply a Central Access Policy**

1. Switch to LON-SVR1.
2. On LON-SVR1, start Windows PowerShell®.
3. Refresh Group Policy by opening Windows PowerShell and running **gpupdate /force**.
4. Open File Explorer, and browse to the **C:\Docs** folder.
5. Open the properties for this folder, and then navigate to **Security\Advanced\Central Policy**.
6. Apply the **Protect confidential docs** central access policy to the **C:\Docs** folder.
7. Browse to the **C:\Research** folder, and then open its properties.
8. Navigate to **Security\Advanced\Central Policy**.
9. Apply the **Department Match** central policy to the **C:\Research** folder.

Results: After completing this exercise, students should have implemented DAC.

Exercise 4: Validating and Remediating DAC

Scenario

In order to ensure that the DAC settings are configured correctly, you will test various scenarios in which users access the files. You will try both approved users and devices, as well as unapproved users and devices. Also, you will validate the access remediation configuration.

The main tasks for this exercise are as follows:

1. Configure access-denied remediation settings
2. Verify DAC functionality
3. View effective permissions
4. Prepare for the next lab


► **Task 1: Configure access-denied remediation settings**

1. Switch to LON-DC1.
2. On LON-DC1, open the Group Policy Management Console.
3. In the Group Policy Management Console, browse to **Forest:Adatum.com\Domains\Adatum.com\Group Policy Objects**, and then 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.
9. Review the other options without making any changes, and then click **OK**.


10. In the right pane of Group Policy Management Editor, double-click **Enable access-denied assistance on client for all file types**, click **Enabled**, and then click **OK**.
11. Close the Group Policy Management Editor, and then close the Group Policy Management Console.
12. Switch to LON-SVR1.
13. On LON-SVR1, use Windows PowerShell to refresh Group Policy.

► Task 2: Verify DAC functionality

1. Restart LON-CL1.
2. Sign in to LON-CL1 as **Adatum\April** with the password **Pa\$\$w0rd**.
3. Click the **Desktop** tile, and then open **File Explorer**.
4. In File Explorer, browse to **\\LON-SVR1\Docs**.
5. Verify that you can open only **Doc3**.
6. Try to access **\\LON-SVR1\Research**. You should be unable to access it.
7. Click **Request Assistance**, review the available options, and then click **Close**.
8. Sign out of LON-CL1.
9. Sign back in to LON-CL1 as **Adatum\Allie** by using the password **Pa\$\$w0rd**.
10. Open File Explorer, and try to access **\\LON-SVR1\Research**.

 **Note:** You should be able to access this folder and open documents inside, because Allie is a member of the Research Department.

11. Sign out of LON-CL1.
12. Sign back in to LON-CL1 as **Adatum\Aidan** with the password **Pa\$\$w0rd**.
13. Open File Explorer and try to access **\\LON-SVR1\Docs**.

 **Note:** You should be able to access this folder and open documents inside, because Aidan is a member of the Managers Department, and he is accessing the documents from a computer that is a member of the ManagersWKS group.

► Task 3: View effective permissions

1. On LON-SVR1, open the **Properties** dialog box for **C:\Research**.
2. Open **Advanced options** for **Security**.
3. Click the **Effective access** tab.
4. Select **April** as a user, and then select **View effective access**. April should not have access to this folder.
5. Add a user claim: **Include a user claim:**
 - **Company Department = Research**
6. Verify that April now has access.

► **Task 4: Prepare for the next lab**

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

1. On the host computer, start Microsoft Hyper-V® Manager.
2. On the **Virtual Machines** list, right-click **20414C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machines** dialog box, click **Revert**.
4. Repeat steps two and three for 20414C-LON-SVR1 and 20414C-LON-CL1.

Results: After completing this exercise, students should have validated functionality of DAC.

Question: What was your approach to the data access design?

Question: What was your approach to the DAC design?

Question: How does your organization implement data access for branch offices?

Lesson 2

Planning Workplace Join

Many third-party applications are available to support mobile device users and roaming users. However, this means that security is an increasingly important concern. For devices that are non-domain joined, you can now implement Workplace Join. Workplace Join creates an object in AD DS to represent the device. That object authenticates the device in a similar way to how a computer object authenticates a domain-joined computer.

In this lesson, you will learn about the components of Workplace Join. Also, you will learn about the process that you use to perform a Workplace Join. Finally, you will learn about the authentication considerations for Workplace Join.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe a business case for using Workplace Join.
- Describe Workplace Join technology.
- Describe Workplace Join components.
- Describe how to register and enroll devices.
- Describe single sign-on (SSO) considerations.

Business Case for Using Workplace Join

More and more people expect to be able to use their own devices to both pursue consumer purposes and access business data. As we use information technology constantly throughout the day, we blur the traditional boundaries between work and home.

In the past, users accessed email mostly from private devices such as laptops, smart phones, and tablets. Today, more and more companies are promoting the BYOD concept, which allows users to choose the devices on which they work. In the BYOD concept of work, users do not have a

dedicated business or private computer. On the contrary, they can use any device for both private and business purposes. Users select and customize devices to fit their personalities, activities, and schedules.

In such scenarios, it is important for administrators to enable users with non-company devices to access company resources when they are at their workplace, at home, or travelling. Most users will not want or be able to join their private devices to the company AD DS domain. Also, some might use devices other than Microsoft devices, such as a tablet or a smartphone.

In Windows Server 2012 R2, Microsoft has provided a new technology to address BYOD scenarios. With this technology, users can access business data with their own devices. While doing so, they can have the same or similar user experience as if they were using domain-joined computers. This technology, named Workplace Join, follows the BYOD concept from both the administrator's perspective and the user's perspective.

- The BYOD concept allows users to use their private devices to do their work
- Connecting non-domain, non-managed devices to company networks and resources can pose a security risk
- Administrators need technology to provide users with flexibility while maintaining security
- Windows Server 2012 R2 provides Workplace Join technology

Workplace Join Overview

With Workplace Join technology, users are able to join their devices to company networks in a new way. Instead of joining the device to the AD DS domain, workplace-joined devices become *known devices*. A known device is one that is allowed access to company resources. The user of the device is given an SSO experience when accessing company resources and applications. Known devices store a subset of their attributes in AD DS. This means that these attributes can be used to manage conditional access for purposes of authorization. Also, it is possible to implement multifactor authentication for additional security.

- Workplace Joined devices become known devices to AD DS
- Known devices store a subset of their attributes in AD DS
- Device registration service provisions a device object in AD DS and issues a certificate to known devices
- Users on known devices have an SSO experience
- Windows Server 2012 R2 with AD FS role service is required
- Windows 8.1 client operating system or iOS-based devices are supported
- Device registration service can be published externally by using Web Application Proxy

To implement Workplace Join, you must have Windows Server 2012 R2 with the AD FS role service installed. On the client side, you must use the Windows 8.1 client operating system or iOS-based devices such as the iPad.

A new service in AD FS called device registration service is responsible for making Workplace Join possible. When users initiate a Workplace Join process from their machines, device registration service provisions a device object in AD DS and also issues a certificate for the workplace-joined device. This certificate is used later to represent device identity when accessing company resources. This service works internally, by default, which means that you have to connect a user's device to the internal network to make it a known device. However, you use device registration service together with Web Application Proxy, also a new functionality in Windows Server 2012 R2, so you can publish this service to the Internet.

From a user's perspective, one of the most important benefits of Workplace Join is the SSO experience. When they use known devices, users will be prompted for their domain credentials only once during the lifetime of the SSO session, as if they were using domain-joined devices. However, the administrator can enforce a password prompt or reauthentication of some resources based on certain criteria.



Note: When using workplace-joined known devices, users are still required to have valid domain user credentials, and each user of the same device must perform Workplace Join separately.

Workplace Join Components

Workplace Join enables you to control access to claims-enabled applications based on device information. To enable Workplace Join, you must have domain controllers from Windows Server 2003 R2 or newer to support the necessary schema extensions.

A Workplace Join scenario involves the following components:

- Clients. The supported clients for Workplace Join are Windows 8.1 and iOS devices, such as iPhone and iPad. Android support is planned.

The Workplace Join process involves the following components:

- Clients
- Claims-aware application support
- A Web Application Proxy
- An AD FS server
- A device registration service

- Claims-aware application support. A claims-aware application is an application that uses claims from AD FS for authentication. AD FS provides the information about the registered device. As a result, only applications that are claims-aware can control access based on information about devices that have performed a Workplace Join. Also, only claims-aware applications can require that devices have performed a Workplace join.
- A Web Application Proxy. A Web Application Proxy server is installed in the perimeter network. This facilitates communication between devices on the external network and AD FS. Devices on the internal network perform a Workplace Join by communicating directly with the AD FS server.
- An AD FS server. This server hosts the device registration service that performs the Workplace Join process for clients.
- A device registration service. This service performs the Workplace Join process for clients. It is responsible for creating the object in AD DS that represents the workplace-joined object. This service also distributes a certificate to the client.



Note: You can use Workplace Join to provide additional security for the new Work Folders feature in Windows Server 2012. Work Folders synchronizes files between multiple devices.

Registering and Enrolling Devices

When you are using Workplace Join functionality, you must establish the relation between the device you're your AD DS. Regardless of the client type, the client must trust the service communication certificate configured for AD FS. Because the organization does not already manage devices that perform a Workplace Join, you should use a certificate from a trusted third-party certification authority. This avoids the need to configure each device to trust your internal certification authority (CA).

The Workplace Join process requires clients to perform a certificate revocation check on the certificate used by the AD FS server or Web Application Proxy with which they are communicating. If the certificate revocation check fails, the Workplace Join will also fail. When using a third-party CA, you do not have to configure a certificate revocation list distribution point (CDP) for your internal CA that is accessible from the Internet.

Workplace Join for Windows-based Devices

During Workplace Join, you are prompted to provide your email address and password. The required information is actually your user principal name (UPN) and not your email address. To simplify this process, we strongly recommend that a user's UPN match his or her email address.

Windows devices locate the server for Workplace Join automatically based on the provided UPN. The server used for Workplace Join is `enterpriseregistration.upndomainname.com`. You must configure Domain Name System (DNS) to resolve this record to the IP address of your AD FS server or Web Application Proxy that is configured to support Workplace Join.

The certificate for the AD FS server and AD FS proxy functionality of Web Application Proxy needs to include the `enterpriseregistration.upndomainname.com` domain name. The configuration process is easier if you include this name in the certificate used during the installation of AD FS and Web Application Proxy, instead of changing the certificate after installation.

- To perform a Workplace Join, devices must trust the service communication certificate for AD FS
- A certificate is placed on the device for authentication

Devices running Windows	Devices running iOS
Require a UPN for authentication	Use Safari to install a configuration profile
Access <code>enterpriseregistration.upndomainname.com</code>	

Workplace Join for iOS Devices

To perform a Workplace Join for an iOS device, you need to set up a configuration profile on the iOS device. You create an iOS configuration profile by providing an XML file. For a Workplace Join, a website delivers the XML file. This is referred to as *over-the air profile delivery*.

The website that iOS devices use to download the configuration profile is located on the AD FS server where the device registration service is enabled. An example of a URL used to configure an iOS device is <https://adfs.contoso.com/enrollmentservice/otaprofile>. On the website, you are prompted to sign in by using your email address as a user name. Like the process for devices that run Windows, you should enter your UPN rather than your email address. After signing in, you install the profile on the iOS device. If the iOS device requires you to enter a PIN to unlock the device, it will prompt you to enter the PIN before the profile is installed.

Certificates on Devices

The Workplace Join process places a certificate on the device. The device uses this certificate to prove its identity. This certificate is used to authenticate to the object created for the device in AD DS.

Single Sign-On Considerations for Workplace Join

When Workplace Join has been completed for a device, the authentication process is modified. Using a workplace-joined device provides SSO and two-factor authentication.

SSO

The benefits of SSO vary depending on the scenario. In some scenarios, SSO means that you can use a single set of credentials to access multiple applications, but you must authenticate to each application separately. In other scenarios, SSO means that authentication is required only once to access multiple applications.

Using the Workplace Join feature provides:

- SSO
- Two-factor authentication

You can implement additional multi-factor authentication by configuring AD FS

- Example: Windows Azure Multi-factor Authentication

When you use Workplace Join from a particular device, you provide authentication credentials to AD FS the first time you access an application. Then your authentication is cached as a cookie on that device. This cookie is used for subsequent access to the same application and other applications that use the same AD FS server farm.

The cookie for authentication is cached for an extended period, until you are required to provide credentials again. To enhance security, you can shorten the length of time that the cookie is retained.

Two-Factor Authentication

If an application is configured to allow access from devices that are workplace-joined only, it requires a second factor for authentication. To access the application, you must have valid credentials and be accessing the application from a device that has been linked to your credentials. The Workplace Join process applies to a specific combination of a user account and a device. For a second user to use a workplace-joined device to access the application, the second user must also perform a Workplace Join.

Multifactor Authentication

Workplace-joined devices can also perform multifactor authentication by using additional authentication providers such as Windows Azure™ Multi-Factor Authentication. However, this multifactor authentication is a function of AD FS rather than Workplace Join.

If you integrate Windows Azure Multi-Factor Authentication with AD FS, you can implement the following methods for additional authentication:

- Phone calls. With this method, you receive a call on your phone to confirm your authentication. You press the pound symbol (#) to confirm after receiving the call.
- Text messages. With this method, you receive a text message with a passcode. You respond to the text message and include the passcode.
- Mobile App. With this method, an authentication prompt that you must acknowledge appears in the mobile app.

Lesson 3

Planning Work Folders

In today's business environment, it has become more and more common for people to perform work duties on their own computers, tablets, and smart phones. All day, users use the same UI to perform both personal and business-related tasks. This is the result of the BYOD approach that many companies have adopted recently. BYOD is the policy of permitting employees to use personally owned mobile devices, such as laptops, tablets, and smart phones to access privileged company information and applications in the workplace.

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

Lesson Objectives

After completing this lesson, you will be able to:

- Describe Work Folders.
- Implement Work Folders.
- Compare Work Folders to similar current technologies.
- Describe how to integrate Work Folders with other BYOD technologies.
- Configure Work Folders.
- Describe considerations for implementing Work Folders.

Work Folders Overview

A local hard disk drive of a computer or a tablet is often an unsecure and inefficient place to store a user's individual data. Because users commonly use more than one device, it is difficult to keep these devices both synchronized with business data and backed up to protect this data efficiently.

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

Alternatively, users with mobile computers or laptops that are members of a company's AD DS domain often need to access company data while they are offline. They can use Offline Files to keep important data available locally on their computers, even when they are not connected to the network. However, Offline Files are synchronized only when users connect to the company's local network. If users are offline for a long time, they may be working on old copies of data.

Work Folders:

- Enable users to access business data securely from any location and from any device, including Windows devices and devices other than Windows
- Are managed by administrators in AD DS
- Are currently supported on Windows 8.1 devices
- Are planned to be supported by iOS-based and Android-based devices

To overcome these problems, Microsoft has implemented a new technology named Work Folders in Windows Server 2012 R2 and Windows 8.1. This technology enables users to access their individual business data independently of their location. It enables administrators to manage the data and settings of this technology.

The main purpose of Work Folders is to provide a user access to the latest data, no matter where the user is located, whether on the corporate network or logged in remotely. Also, by using Work Folders, administrators can manage data, including a user's connections to Work Folders. The administrator can enforce the encryption of Work Folders and can control which users can use this functionality. The administrator can enforce some security settings on a device that uses Work Folders, even if it is not a domain member.

Users can use Work Folders while they are in a local network, but also when they are out of the network—for example, while they are at home or traveling.

You can publish Work Folders to the Internet by using the Web Application Proxy functionality, which is also new to Windows Server 2012 R2. Web Application Proxy functionality enables users to synchronize their data whenever they have an Internet connection. This overcomes the limitation of Offline Files and enables users to synchronize files in their Work Folders from any location.



Note: Currently, Work Folders are available only for Windows 8.1 client operating systems. However, Microsoft plans to add support for Work Folders for devices running Windows 7, Windows 8, iOS and Android operating systems in the future.

Implementing Work Folders

To use Work Folders, you should have at least one Windows Server 2012 R2 file server and at least one Windows Server 2012 R2 domain controller in your network. Work Folders is a service of the File and Storage Services server role. You can install it easily by using Server Manager. Although it is not mandatory, it is a best practice that you also install FSRM and Data Deduplication functionality if you want to manage user data more efficiently.



Note: When you install Work Folders functionality, Microsoft Internet Information Services (IIS) Hostable Web Core and IIS

Management tools are also installed. You do not have to configure any IIS settings, but you must assign a trusted Secure Sockets Layer (SSL) certificate to your file server in the IIS Console and bind it to port 443 on the default web site. The certificate should have a file server name and the name under which you plan to publish your Work Folders, if different.

After you install Work Folders functionality, you should provision a share to store users' data. You can locate a share on any location that is accessible and controlled by the file server where you installed Work Folders. When you create a root share, we recommended that you leave the default values on Share and NTFS file system permissions and that you enable access-based enumeration.

After you create a root share where users' Work Folders are located, you should start the New Sync Share Wizard to create the Work Folders structure. You should select the root folder that you provisioned as a

To use Work Folders, you should:

- Have at least one Windows Server 2012 R2 file server
- Have at least one Windows Server 2012 R2 domain controller
- Install Work Folders functionality on the file server
- Provision a share where users' data is stored
- Run the New Sync Share Wizard to create the Work Folders structure
- Configure clients to use Work Folders manually or by using Group Policy

share. Also, you should choose the format for the subfolders' names. It can be a user alias, or *alias@domain*. If you have more than one domain in your AD DS forest, we recommend that you use the *alias@domain* naming format.

You can control Sync Access by listing users who use your Work Folders structure or by specifying a group. We recommend that you specify a group for easier administration. In addition, we recommend that you disable permission inheritance for Work Folders so that each user has exclusive access to his or her files.

During the setup wizard for Work Folders provisioning, you can enforce some security settings on devices that access Work Folders. You can enforce Work Folders encryption on a client device. Also, you can enforce use of a lock screen with a password for additional security. This is especially useful for devices that are non-domain joined, as you cannot control them in any other way. When a user on a device that is non-domain joined is configuring Work Folders, he or she must accept usage of security policies before the Work Folders synchronization begins.



Note: You cannot use Group Policy to enforce security settings related to Work Folders, which can be accessed from unmanaged devices. These settings are enforced when a user establishes the Work Folders connection, and they are applied on domain-joined and non-domain joined computers.

Configuring Clients to Use Work Folders

You can configure Windows 8.1 clients manually to use Work Folders or to use Group Policy. For domain-joined computers, it is easier to configure settings by using Group Policy, but you must configure non-domain clients manually.

If you use Group Policy to configure Work Folders automatically, you should check two locations. Because Work Folders are user-based, you perform configuration in the user part of the GPO. When you open the Group Policy Management Editor, you should navigate to User Configuration\Policies\Administrative Templates\Windows Components\Work Folders. There you will open the Specify Work Folders settings and enable the policy. You must also configure the Work Folders URL. This URL is the location of your file server where you enabled Work Folders, followed by your fully qualified domain name (FQDN), as in <https://fileserverFQDN>. In this same GPO setting, you have the option to force automatic setup for each user. You should use this option with caution. If you enable it, all users to whom this GPO applies will have their Work Folders configured on each device to which they log on, without being prompted. This is undesirable in some scenarios.

You can manage some Work Folders settings in the computer part of the GPO. If you navigate to Computer Configuration\Policies\Administrative Templates\Windows Components\Work Folders, you can find the option to Force automatic setup of Work Folders for all users. Computers that have this GPO setting applied configure Work Folders for every user that signs on.

After you apply these Group Policy settings to the users' domain, and optionally the computers' domain, users can start using Work Folders.

If you also want to enable Work Folders on a non-domain joined computer—for example, on a personal tablet that an employee is using—you have to configure this manually by using the Work Folders item in the Control Panel of Windows 8.1. You will have to provide a valid user name and password for the domain account that is allowed to use Work Folders, as well as a file server URL. If you want to avoid having users type a file server URL, you can configure Work Folders on non-domain joined devices by typing the user's UPN. However, in this case, you must have a workfolders CNAME record in your public DNS, that resolves to the name of your Work Folders file server.

Work Folders Vs. Alternative Technologies

It is important to know how Work Folders differ from other similar technologies. Also, you should be aware that Work Folders cannot replace all functionalities that other technologies offer. For example, although Microsoft OneDrive™ for Business, formerly known as Microsoft SkyDrive Pro, can provide connectivity to team data on your company Microsoft SharePoint® locations, Work Folders do not provide this benefit. On the other hand, Work Folders can replace the functionality of Folder Redirection and Client-Side Caching by using Offline Files.

- Work Folders do not replace similar existing technologies
- Alternative technologies:
 - One Drive
 - OneDrive for Business
 - Folder Redirection/Client-Side Caching
 - Offline Files
- Identify key differences and understand which technologies meets your business requirements

The following table compares similar technologies for managing and accessing user data.

	Personal data	Individual work data	Team/group work data	Personal devices	Data location
OneDrive	Yes			Yes	Public cloud
OneDrive for Business		Yes	Yes	Yes	Microsoft SharePoint / Microsoft Office 365™
Work Folders		Yes		Yes	File server
Folder Redirection/ Client-Side Caching		Yes			File server

You should consider ways to protect data stored in Work Folders. This is especially important because users can access their Work Folders from unmanaged devices and devices other than Windows devices.

Integrating Workplace Join, Web Application Proxy, and Work Folders

Because Workplace Join, Web Application Proxy, and Work Folders are technologies that support the BYOD concept, you can implement them together to improve users' experiences when they use their own devices to do business work.

As long as users have valid AD DS credentials, they can configure Workplace Join, Work Folders, or both technologies on their private devices.

Implementing Workplace Join together with Work Folders will provide users a SSO experience, including access to their data. However, when users are outside your internal network, you must provide them with secure access to a current copy of their data. In other words, you have to provide a secure way for the external users to access and sync their Work Folders content even when they are not connected to the internal network.

- The Workplace Join feature, Work Folders, and Web Application Proxy can work together to provide a full BYOD experience for end users
- These three technologies can also work independently
- Security should be the main consideration when planning deployment of BYOD technologies

You can use Web Application Proxy functionality to publish your Work Folders structure securely for external users. By doing so, you allow them to sync their data. Work Folders publishing with Web Application Proxy makes these folders accessible from any place that has an Internet connection.

When used together, these three technologies provide a full BYOD experience for end users. However, the BYOD concept can introduce potential security threats that administrators should consider when planning these technologies.

Demonstration: Configuring Work Folders

In this demonstration, you will see how to configure Work Folders.

Demonstration Steps

1. On LON-SVR1, in Server Manager, open **File and Storage Services**, and then select **Work Folders**.
2. Start the New Sync Share Wizard.
3. Select the share that you created in preparation tasks, WF-Share.
4. Specify user alias for the structure for user folders.
5. Grant access to the WFSync user group.
6. Switch to LON-DC1.
7. Open the Group Policy Management Console.
8. Create a new GPO and name it **Work Folders GPO**.
9. Open the Group Policy Management Editor for Work Folders GPO.
10. Expand **User Configuration/ Policies/Administrative Templates/Windows Components**, and then click **Work Folders**.
11. Enable Work Folders support, and then type **https://lon-svr1.adatum.com** as the Work Folders URL.
12. Link the Work Folders GPO to the domain.

Considerations for Implementing Work Folders

Before you implement Work Folders, you should be aware of several factors that can affect both the functionality and security of this BYOD technology. When planning your Work Folders implementation, consider the following guidelines:

- Do not enable Work Folders for all users at once. Instead, implement this technology in phases, implementing it first for the users who have more than one device.
- Use Group Policy to enable and configure Work Folders automatically for the domain-joined client computers.

When you plan your Work Folders implementation, consider the following guidelines:

- Do not enable Work Folders for all users at once
- Use Group Policy to configure Work Folders automatically
- Ensure that you have a reliable backup strategy
- Always enforce security policies
- Consider Work Folders encryption
- Publish the Work Folders server securely
- Issue an appropriate and trusted certificate

- Ensure that you have an appropriate and reliable backup strategy for locations that contains Work Folder data.
- Always enforce security policies to devices that use Work Folders and that are not part of your AD DS environment.
- Consider encryption of Work Folders for additional security.
- If you allow users to access Work Folders from the Internet, ensure that you publish the Work Folders server securely. Use Web Application Proxy or a similar technology for publishing.
- Ensure that you issue an appropriate certificate to the Work Folders server. In addition, ensure that the certificate is trusted on the client side, and that clients that are connecting from an external network, such as the Internet, can access at least one CDP specified in the certificate.

If you have already implemented Folder Redirection or Offline Files, you can maintain these technologies to work together with Work Folders. Although they have a similar purpose, these technologies are not analogous alternatives to Work Folders. Also, if users in your organization currently use consumer-oriented public cloud storage services for business data, ensure that they move to Work Folders as soon as possible.

Lab B: Implementing Work Folders

Scenario

You are an administrator for the A. Datum Corporation. The company has a wide range of sales staff and general users who work remotely and use different types of personal devices.

The sales staff needs to be able to access their work data on their personal devices, from any location. This will allow them to manage clients while on the road.

You must implement a feature that allows users to access their data from anywhere, and that supports a variety of corporate-owned and personal devices.

Objectives

After completing this lab, you will be able to:

- Plan and implement an infrastructure for Work Folders.
- Configure AD FS and Web Application Proxy for Work Folder publishing.
- Validate Work Folders functionality.

Lab Setup

Estimated Time: 60 minutes

Virtual machines	20414C-LON-DC1 20414C-LON-SVR1 20414C-LON-SVR2 20414C-LON-CL1 20414C-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 **20414C-LON-DC1**, and, in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Sign in using the following credentials:
 - User name: **Administrator**
 - Password: **Pa\$\$w0rd**
 - Domain: **Adatum**
5. Repeat steps two through four for 20414C-LON-SVR1 and 20414C-LON-SVR2. Do not sign in to 20414C-LON-CL1 and 20414C-LON-CL2 until instructed to do so.
6. On LON-SVR2, right-click the **Start** button, and then click **Network Connections** in the menu.
7. In the Network Connections window, right-click **Ethernet 2** network connection object, and then click **Enable**.
8. Close the Network Connections window.

Exercise 1: Preparing and Implementing an Infrastructure for Work Folders

Scenario

To address the requirements from the lab scenario, you decide to implement Work Folders technology. The first step in implementing Work Folders is currently enabling a Sync Server Share on an existing File Server on your network.

The main tasks for this exercise are as follows:

1. Installing Work Folders functionality and configuring a Secure Sockets Layer certificate
2. Provision a share for Work Folders
3. Configure and implement Work Folders

► Task 1: Installing Work Folders functionality and configuring a Secure Sockets Layer certificate

1. Sign in to LON-SVR1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. Start Server Manager.
3. Add the **Work Folders** role service by using the Add Roles and Features Wizard.
4. Open the **Internet Information Services (IIS) Manager** console.
5. Create a domain certificate for lon-svr1.adatum.com as follows:
 - Common name: **lon-svr1.adatum.com**
 - Organization: **Adatum**
 - Organizational Unit: **IT**
 - City/locality : **Seattle**
 - State/province : **WA**
 - Country/region: **US**
6. Bind this certificate to the https protocol on the Default Web Site.

► Task 2: Provision a share for Work Folders

1. On LON-SVR1, in Server Manager, expand **File and Storage Services**, and then click **Shares**.
2. Start the New Share Wizard.
3. Select the **SMB Share – Quick** profile.
4. Name the share **WF-Share**.
5. Enable access-based enumeration.
6. Leave all other options on default values.

► Task 3: Configure and implement Work Folders

1. On LON-SVR1, in Server Manager, expand **File and Storage Services**, and then select **Work Folders**.
2. Start the New Sync Share Wizard.
3. Select the share that you created in previous step: **WF-Share**.
4. Specify user alias for the structure for user folders.
5. Grant access to the WFSync user group.
6. Do not apply device policies.
7. Switch to LON-DC1.

8. Open the Group Policy Management Console.
9. Create a new GPO and name it **Work Folders GPO**.
10. Open the Group Policy Management Editor for Work Folders GPO.
11. Expand **User Configuration/Policies/Administrative Templates/Windows Components**, and then click **Work Folders**.
12. Enable the Work Folders support and type <https://lon-svr1.adatum.com> as the Work Folders URL.
13. Link the Work Folders GPO to the domain adatum.com.

Results: After completing this exercise, students will have configured a Work Folders server infrastructure.

Exercise 2: Configuring AD FS and Web Application Proxy for Work Folders Publishing

Scenario

To publish Work Folders functionality for external users and their private devices, you need to deploy AD FS and Web Application Proxy. With Web Application Proxy, you will publish Work Folders hosted on LON-SVR1 to an external network by using the public name wf.adatum.com and pass-through authentication.

The main tasks for this exercise are as follows:

1. Install AD FS
2. Configure AD FS
3. Install Web Application Proxy
4. Configure Web Application Proxy
5. Publish Work Folders through Web Application Proxy

► Task 1: Install AD FS

1. On LON-DC1, use the DNS Manager to add a new host record for AD FS:
 - Forward lookup zone: **Adatum.com**
 - Name: **adfs**
 - IP address: **172.16.0.10**
2. On LON-DC1, in Server Manager, add the Active Directory Federation Services role.
3. On LON-DC1, at a Windows PowerShell command prompt, type the following cmdlet: **Add-KdsRootKey -EffectiveTime (Get-Date).AddHours(-10)**. This enables usage of managed service accounts for AD FS.

► Task 2: Configure AD FS

1. In the Server Manager notifications, click **Configure the federation services on this server**.
2. Use the following options to configure the AD FS server:
 - **Create the first federation server in a federation server farm**
 - Account for configuration: **Adatum\Administrator**
 - SSL Certificate: **ADFS.adatum.com**
 - Federation Service Display Name: **A. Datum Corporation**
 - Create a Group Managed Service Account: **Adatum\ADFS**
3. **Create a database on this server using Windows Internal Database**

► Task 3: Install Web Application Proxy

1. On LON-SVR2, in Server Manager, add the Remote Access server role and the Web Application Proxy role service.

► Task 4: Configure Web Application Proxy

1. On LON-DC1, open the Microsoft Management Console, and then add the **Certificates** snap-in for the **Local Computer**.
2. From the Personal folder, export the **ADFS.adatum.com** certificate by using the following settings:
 - **Yes, export the private key**
 - File format: **Personal Information Exchange – PKCS #12 (.PFX)**
 - Password: **Pa\$\$w0rd**
 - File name: **C:\adfs.pfx**
3. On LON-SVR2, open the Microsoft Management Console, and then add the **Certificates** snap-in for the **Local Computer**.
4. From the Personal folder, import the **adfs.adatum.com** certificate:
 - File name: **\\LON-DC1\c\$\adfs.pfx**
 - Password: **Pa\$\$w0rd**
 - Certificate store: **Personal**
5. On LON-SVR2, in Server Manager, click the **Notifications** icon, and then click **Open the Web Application Proxy Wizard**.
6. In the Web Application Proxy Wizard, provide the following configuration settings:
 - Federation service name: **adfs.adatum.com**
 - User name: **Adatum\Administrator**
 - Password: **Pa\$\$w0rd**
 - Certificate to be used by the AD FS proxy: **adfs.adatum.com**
7. Leave the Remote Access Management Console open for the next task.

► Task 5: Publish Work Folders through Web Application Proxy

1. On LON-SVR2, in the Remote Access Management Console, publish a new application with the following settings:
 - **Preauthentication:** Pass-through
 - **Name:** A. Datum Work Folders
 - **External URL:** <https://wf.adatum.com/>
 - **External certificate:** adfs.adatum.com
2. **Backend server URL:** <https://lon-svr1.adatum.com/>

Results: After completing this exercise, students will have configured Active Directory® Federation Services (AD FS) and Web Application Proxy services for Work Folders publishing.

Exercise 3: Validating Work Folders Functionality

Scenario

After configuring a Work Folders infrastructure and publishing it on both internal and external networks, you want to validate access to Work Folders. You will use one domain-joined client device, LON-CL1, and one device that is non-domain-joined, LON-CL2, to access Work Folders. A domain-joined client will be accessing Work Folders from an internal network by using an internal server name, while a device that is non-domain-joined will connect from an external network, through Web Application Proxy, and by using a public wf.adatum.com name.

The main tasks for this exercise are as follows:

1. Validating Work Folders from a domain-joined client device
2. Validating Work Folders from a non-domain-joined device
3. Validate Work Folders data synchronization
4. Prepare for the next module

► Task 1: Validating Work Folders from a domain-joined client device

1. Sign in to LON-CL1 as **Adatum\Aidan** with the password **Pa\$\$w0rd**.
2. Open **Windows PowerShell**, and then refresh Group Policy.
3. Open **File Explorer** and ensure that a folder called Work Folders has been created.
4. Create a few text files and save them in Work Folders.
5. Open the Control Panel, and then open the Work Folders applet and ensure that Work Folders are working properly.

► Task 2: Validating Work Folders from a non-domain-joined device

1. Use local account **Delaney** and password **Pa\$\$w0rd** to sign in locally on the LON-CL2 virtual machine.
2. Use Windows PowerShell to ping lon-svr1.adatum.com and wf.adatum.com. The first ping should be unsuccessful while the second one should provide you with a reply from 131.107.0.2. (Note: The ping to lon-svr1.adatum.com verifies that you cannot access a server that hosts Work Folders by using its internal name. The ping to wf.adatum.com verifies that you can access the Web Application Proxy server external interface and the name used to publish Work Folders.)
3. Open the Work Folders applet from the Control Panel, and start the Set Up Work Folders wizard. Provide the following data:
 - Work Folders URL : **https://wf.adatum.com**
 - Domain credentials : **aidan@adatum.com** and **Pa\$\$w0rd**
 - Accept policies for Work Folders
4. Ensure that the Work Folders local folder opens and that files created by Aidan on LON-CL1 are present.

► Task 3: Validate Work Folders data synchronization

1. In the Work Folders folder on LON-CL2, create a few new text documents.
2. Switch to LON-CL1, make changes to the existing documents in the Work Folders folder, wait for a few minutes, and ensure that the new documents, created on LON-CL2, appear on LON-CL1.

► Task 4: Prepare for the next module

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

1. On the host computer, start Hyper-V Manager.
2. On the Virtual Machines list, right-click **20414C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machines** dialog box, click **Revert**.
4. Repeat steps two and three for 20414C-LON-SVR1, 20414C-LON-SVR2, 20414C-LON-CL1, and 20414C-LON-CL2.

Results: After completing this exercise, students will have validated Work Folders functionality from both domain-joined devices and devices that are non-domain joined.

Question: What is the purpose of implementing Web Application Proxy?

Question: Why did you implement AD FS in this lab?

Question: Why do you need to accept security policies when configuring Work Folders on a non-domain joined device?

Module Review and Takeaways

Review Questions

Question: What is the BYOD concept?

Question: What is the main purpose of access-denied assistance technology?

Question: What is the main benefit of Workplace Join for end users?

Question: What is a claim?

Question: What is the purpose of a Central Access Policy?

Tools

Tool	Use	Location
Active Directory Administrative Center	Administering and creating claims, resource properties, rules, and policies	Administrative tools
Group Policy Management Console (GPMC)	Managing Group Policy	Administrative tools
Group Policy Management Editor	Editing Group Policy Objects (GPOs)	GPMC
Active Directory Federation Services (AD FS) Management Console	AD FS Management	Administrative tools
Web Application Proxy Console	Secure publishing of internal resources	Administrative tools
Work Folders administration interface	Work Folders provisioning and management	Server Manager



Best Practice:

- Use Central Access Policies instead of configuring conditional expressions on resources.
- Enable access-denied assistance settings.
- Always test changes that you have made to Central Access Rules and to Central Access Policies before implementing them.
- Use file classifications to assign properties to files.
- Use Work Folders to synchronize business data across devices.
- Always use HTTPS for Work Folders client connections.
- Use Workplace Join in BYOD scenarios.

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
Claims are not populated with the appropriate values.	Verify that the correct attribute is selected for the claim. In addition, check that the attribute value for a specific object is populated.
A conditional expression does not allow access.	Verify that the expression is well defined. In addition, try using the Effective Access tab to troubleshoot the problem.
Work Folders do not sync on non-domain joined devices.	Ensure that a non-domain joined device trusts the certificate installed on the server that hosts Work Folders.

Module 13

Planning and Implementing an Information Rights Management Infrastructure

Contents:

Module Overview	13-1
Lesson 1: AD RMS Overview	13-2
Lesson 2: Planning and Implementing an AD RMS Cluster	13-8
Lesson 3: Planning and Implementing AD RMS Templates and Policies	13-19
Lesson 4: Planning and Implementing External Access to AD RMS Services	13-29
Lesson 5: Planning and Implementing AD RMS Integration with Dynamic Access Control	13-43
Lab: Planning and Implementing an AD RMS Infrastructure	13-46
Module Review and Takeaways	13-57

Module Overview

Organizations today face many security-related challenges, but one of the biggest is protecting sensitive data and information from careless mishandling and malicious use. Increasingly, instances of information theft, and the rise of new legislative requirements to protect data, highlight the need for better protection of digital content. The growing use of computers to create and work with sensitive information, the introduction of extensive connectivity through private and public networks (including the Internet), and the appearance of increasingly powerful computing devices have made protecting organizational data an essential security consideration. Sensitive digital content can include dynamic, database-driven reports on an information portal, confidential email messages, strategic planning documents, military defense reports, and other sensitive files. To help organizations protect this valuable intellectual property, Microsoft has developed the Active Directory® Rights Management Service (AD RMS) as a server role designed to help ensure the protection of this kind of data. This module describes how you can plan and implement an AD RMS deployment to protect content.

Objectives

After completing this module you will be able to:

- Describe AD RMS.
- Plan and implement an AD RMS deployment.
- Plan and manage AD RMS templates and access.
- Plan and implement external access to AD RMS services.
- Plan and implement AD RMS integration with Dynamic Access Control.

Lesson 1

AD RMS Overview

By using AD RMS and the AD RMS client, you can improve an organization's security strategy by protecting information through persistent usage policies, which remain with the information no matter where you move it. Additionally, you can use AD RMS to help prevent inappropriate access to sensitive information, such as financial reports, product specifications, customer data, and confidential email messages. This lesson will detail AD RMS concepts and components, usage scenarios, and implementation and administration.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe and understand AD RMS.

What Is IRM?

Information Rights Management (IRM) is a technology that provides persistent file-level-based permissions and authorization to help prevent unauthorized users from printing, forwarding, or copying sensitive information. IRM may also be referred to as Enterprise Rights Management, Enterprise Digital Rights Management (DRM), Document Rights Management, and Intelligent Rights Management.

Once you restrict permission for a document or message by using this technology, the usage restrictions travel with the document or email message as part of the file's contents. IRM technology allows you to tag information, usually in the form of documents and email messages. These tags determine what users can and cannot do to that information. IRM enables you to provide controls so that you can create, view, edit, and distribute such information separately. Generally, an enterprise uses an IRM system to protect information, such as financial data, intellectual property, and executive communications, in a business-to-business scenario.

Major functional uses of IRM:

- Provides business-level encryption of information
- Enables information protection while in use
- Allows for simple mapping of business classifications
- Provides offline use without users needing network access for particular amounts of time
- Provides full auditing of access to documents and changes that business users make to usage rights

IRM helps businesses and users address two essential needs:

- Restricted permission for sensitive information. IRM helps prevent unauthorized users from accessing and reusing sensitive information. Organizations protect their sensitive intellectual property by relying on firewalls, security measures related to signing in to corporate networks, and other network technologies. These technologies share one shortcoming: legitimate users who have access to the information can share it with unauthorized people. This could lead to a possible breach of security.
- Information privacy, control, and integrity. Business computer users often work with private or sensitive information. When you use IRM, employees do not have to depend on others' discretion to ensure that sensitive materials remain within the company. IRM reduces users' ability to forward, copy, or print confidential information by disabling those functions in documents and messages that have restricted permissions.

For information technology (IT) managers, IRM facilitates the enforcement of existing corporate policies about document confidentiality, workflow, and email retention. For chief executive officers (CEOs) and

security officers, IRM reduces the risk of unauthorized users acquiring key organizational information, whether accidentally, through carelessness, or through malicious intent.

The major uses of IRM include:

- Business-level encryption of information.
- Protection of information while in use, such as limiting copy and paste, blocking screen shot copying, and preventing printing.
- A usage rights policy, which allows for simple mapping of business classifications to information.
- Offline use without required network access for particular lengths of time.
- Full auditing of both access to documents and changes to the usage rights that business users make.

How AD RMS Works

AD RMS is an IRM technology in Windows Server® 2012 that works with AD RMS-enabled applications to help safeguard digital information from unauthorized use, both online and offline, and inside and outside the firewall. Organizations can use AD RMS to protect sensitive and proprietary information, such as financial reports, product specifications, customer data, and confidential email messages.

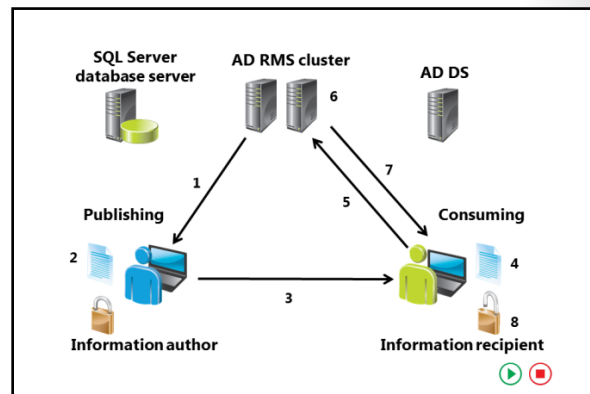
AD RMS applies usage policies, which include usage rights and conditions that remain with the information regardless of its location. The usage policies protect information and enhance an organization's security strategy. AD RMS persistently protects any data in binary format, so the usage rights remain with the information and not on the organizational network. Once an authorized recipient accesses the information, whether online or offline, inside or outside the organization, AD RMS enforces the usage rights.

When a user protects a document and wants to specify a particular level of access to it for other users or groups, he or she starts a multistep process. Although the end user may not be aware of this process, you should understand what happens when a user uses AD RMS to protect a document.

When the author protects content initially, the AD RMS cluster issues a Rights Account Certificate (RAC) and a client licensor certificate, which establish the author's AD RMS credentials. Then the author can publish information that AD RMS helps secure.

The following sequence outlines the AD RMS workflow steps:

1. The author creates a file and specifies the usage rights and conditions by using an AD RMS-enabled application. AD RMS generates a publishing license that contains the usage policies, and then ties the publishing license to the protected file.
2. The application encrypts the file with a symmetric key, which the public key of the AD RMS cluster encrypts. AD RMS inserts the key into the publishing license and then binds it to the file.
3. The author distributes the file. Distribution methods include email, Microsoft® SharePoint® document libraries, and file servers.
4. A recipient obtains the file and opens it by using an AD RMS-enabled application. If the recipient does not have a RAC on the current computer, the AD RMS cluster issues a RAC.



5. The application requests a use license through the AD RMS cluster that issued the publishing license for the secured information.
6. The AD RMS cluster confirms the recipient's authorization level, checks that the recipient is a named user, and creates a use license. The server decrypts the symmetric key by using the private key of the server, and then reencrypts the symmetric key by using the public key of the recipient. Then the server adds the encrypted symmetric key to the use license, which also includes the content expiration date, if applicable.
7. After the AD RMS cluster confirms the recipient's authorization level, the licensing server returns the use license to the recipient's client computer.
8. After receiving the use license, the application verifies the license and the recipient's account certificate, which helps determine whether any certificate, in either chain of trust, requires a revocation list. If it does, the application checks for a local copy of the revocation list that is not expired. If necessary, the application retrieves a current copy of the revocation list. The application then applies any relevant revocation conditions in the current context. If the revocation conditions allow access to the file, the application renders the data. Users can apply their granted rights.

AD RMS Deployment Scenarios

You can deploy AD RMS differently depending on your organization's scenario. You should consider the different trust relationships, such as those between computers in a domain, in a forest of domains, or across untrusted domains, because these relationships require different types of AD RMS structure.

When you design your AD RMS environment, you should consider the following aspects of your Active Directory Domain Services (AD DS) implementation:

- The scope of an AD RMS installation is the AD DS forest. If you have deployed users in multiple forests, then each forest requires its own AD RMS server.
- It is good practice to use a virtual name for AD RMS certification cluster. Typically, this name can be a load-balancing cluster name.

AD RMS deployment scenarios:

- AD RMS in a single forest
- AD RMS licensing-only cluster in a single forest
- AD RMS in multiple forests
- AD RMS in an extranet
- AD RMS with AD FS

Deploying AD RMS in a Single Forest Scenario

AD RMS requires AD DS to manage users and groups and to assign specific privileges to the documents. The healthy management of AD DS is critical for an AD RMS deployment.

Deploying an AD RMS Licensing-Only Cluster Scenario

Optionally, you can deploy servers to address specific licensing requirements, such as the following:

- Support for a department's unique rights management requirements. For instance, a group within your organization may have a different set of rights policy templates that they should not share with the entire organization. Because a forest can have only one root cluster, setting up a separate root cluster is not possible unless you create a new forest. In this case, you could set up a licensing-only cluster that you dedicate to this group's needs, and then set up rights policy templates separately for that licensing-only cluster.
- Support for rights management for external business partners, as part of an extranet that requires strong separation and tracking of resources for specific business partners.

Deploying AD RMS In a Multiple Forest Environment Scenario

If you have AD DS forest trusts with other partners, you should create an AD RMS multiforest deployment. This requires one AD RMS installation per forest. After assigning a Secure Sockets Layer (SSL) certificate to each website that hosts an AD RMS cluster, you will need to extend the AD DS forest schema to include the AD RMS objects. If you are using Microsoft Exchange in the forests, you will have these extensions already. Lastly, you must ensure that all forests involved trust the AD RMS service account.

Deploying AD RMS in an Extranet Scenario

An extranet is an extension of your organization's network to an external source. You can extend the AD RMS cluster to the Internet so that users can consume rights-protected content even when they are not connected to the internal network.

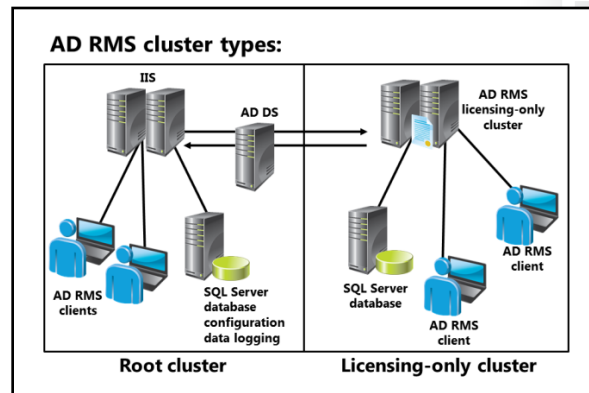
Deploying AD RMS with AD FS Scenario

You can establish a federated trust between two forests by using Active Directory Federation Services (AD FS). This is useful if one forest does not have AD RMS installed, but its users need to consume rights-protected content from another forest. We recommend this connection method between two partners that are running Windows Server 2008 or newer.

Question: Have you implemented AD RMS in your environment? If yes, which deployment scenario are you using? If you are not using it currently, which deployment scenarios might you use in the future?

AD RMS Components

In an AD RMS cluster, all AD RMS servers are either root certification servers or licensing servers. The first AD RMS server in an AD DS forest assumes the Active Directory root certificate role. There can be only one root-certification server in each Active Directory forest. If you add any additional or secondary AD RMS servers, these servers will assume the licensing server role. You might use these additional servers to provide independent policy options to certain groups within an Active Directory forest.



The AD RMS workflow contains several components, which the following table details.

Server component	What does it do?
AD RMS certification server cluster	You use this component to administer and configure AD RMS. It handles all of the major AD RMS functions, including licensing, publishing, account certification, and recovery. Each AD DS forest has a limit of one AD RMS certification server cluster or <i>root cluster</i> .
AD RMS licensing-only cluster	Servers in a licensing-only cluster provide only licensing and publishing services.
Administration web service	The AD RMS server computer hosts this web service, which you can use to manage AD RMS through either the AD RMS administration console or Windows PowerShell® commands for AD RMS.

Server component	What does it do?
Account certification	AD RMS servers generate RACs that associate users with specific computers.
Licensing	AD RMS servers issue end-user licenses, which enable AD RMS client-enabled applications to access protected content, within the user restrictions that the content publisher sets.
Publishing	AD RMS servers create client licensor certificates that enable content publishers to define the policies that you can enumerate in an end-user license.
Precertification	This enables a server to request a RAC on behalf of a user so that Exchange can prelicense content to Microsoft Outlook® users.
Service locator	Provides the URL of the account certification, licensing, and publishing services to AD DS so that AD RMS clients can discover them.

Client component	What does it do?
AD RMS client	Windows® 8, Windows 7, Windows Vista®, Windows Server 2012, and Windows Server 2008 include the client. There is an add-on client for earlier versions of the Windows operating system and the Internet Explorer® browser. The AD RMS client serves as the client component and interacts with the AD RMS certificate server cluster to encrypt and decrypt data.
AD RMS-enabled application	Some applications are enabled for AD RMS and can interact with AD RMS. An author can use these applications to create and help protect content. Additionally, recipients can use these applications to read protected content.
Author	The user or service that generates the rights-protected document.
Client licensor certificate	Grants an author permission to publish rights-protected content without connecting to the corporate network.
Consumer	The user or service that accesses the rights-protected document.
End-user license	This license decrypts the content and enforces the rights enumerated in the license for a specific user.
Issuance license	This lists the users who can decrypt protected content and the rights that AD RMS can make available to them.

Additional components	What does it do?
AD DS	AD DS is an AD RMS prerequisite, and AD RMS uses AD DS to store user accounts and groups. Clients query AD DS for the service connection point to discover registered AD RMS services.
SQL Server database	The AD RMS database stores the configuration and log data. You can use Microsoft Windows Internal Database (WID) in place of the Microsoft SQL Server® software, but we do not support WID in a production environment.

Lesson 2

Planning and Implementing an AD RMS Cluster

When you use AD RMS and the AD RMS client, you can increase an organization's security strategy by protecting information through persistent usage policies, which remain with the information no matter where you move it. Additionally, you can use AD RMS to help prevent sensitive information, such as financial reports, product specifications, customer data, and confidential email messages from intentionally or accidentally getting into the wrong hands. This lesson will detail AD RMS concepts and components, usage scenarios, and its implementation and administration.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe options for configuring AD RMS clusters.
- Install an AD RMS cluster.
- Understand guidelines for designing AD RMS services.
- Configure high availability for AD RMS.
- Plan AD RMS management.
- Implement an AD RMS backup and recovery strategy.
- Understand how to decommission and remove AD RMS.
- Upgrade an AD RMS cluster.

Options for Configuring AD RMS Clusters

An AD RMS cluster is either a single server running AD RMS or a group of servers that share AD RMS publishing and licensing requests from AD RMS clients. The simplest form of a cluster is one AD RMS server. In deployments that are more complex, you can configure multiple servers as a cluster behind a single, shared URL.

When you provision the first AD RMS server in an Active Directory forest, that server always becomes the root cluster. You can provision more servers and add them to the AD RMS cluster at any time.

You can configure two types of clusters in Windows Server 2012 AD RMS:

- Root-certification clusters, which handle all certification and licensing requests for the AD DS domain in which you install the cluster.
- Licensing-only clusters, in which only the servers provide licensing and publishing services only.

There are two types of clusters in Windows Server 2012:

- Root certification cluster
- Licensing-only cluster

Root-certification cluster:

- The first server that you install always becomes the root certification cluster
- It handles all certification and licensing requests for the domain

The simplest cluster is one AD RMS server

To create complex clusters, you can:

- Configure multiple servers as a cluster behind a single, shared URL
- Create licensing-only clusters, in addition to the root-certification cluster

Most organizations deploy licensing-only clusters to address specific licensing requirements, such as:

- Supporting a department's unique rights management requirements. For instance, a department within your organization might have a different set of rights policy templates that they cannot share with the rest of the organization. Because a forest can have only one root-certification cluster, you cannot configure a separate root-certification cluster unless you create a new forest. Alternatively, you can set up a licensing-only cluster dedicated to this department's needs, and then you can set up separate rights policy templates for that licensing-only cluster.
- Supporting rights management for external business partners, as part of an extranet that requires strong separation and resource tracking for specific business partners.
- Removing licensing tasks from the root-certification cluster.

You cannot load balance across root-certification and licensing-only clusters. For this reason, we recommend that you use a root-certification cluster only, not a licensing-only cluster, and that you join additional AD RMS servers to the root-certification cluster. The root-certification cluster (or server) is the first component of an AD RMS deployment, and all other components depend on it. Each Active Directory forest can have only one root-certification cluster. The root-certification cluster runs all the AD RMS services, including the account certification service that provides Rights Account Certificates to the AD RMS clients. You can add servers to the AD RMS installation to create a root-certification cluster, which you can use for redundancy and load balancing.

A licensing-only cluster is optional, and it is not part of the root-certification cluster. It is subenrolled under the root-certification cluster. A licensing-only cluster depends on the root-certification cluster for certification and other services. To establish redundancy and load balancing, you can add servers to the installation to create a licensing-only cluster. Most often, you will deploy licensing-only clusters to address specific licensing requirements.

Demonstration: Installing an AD RMS Cluster

The goal of an AD RMS deployment is to protect information, regardless of where users move it. Once you add AD RMS protection to a digital file, the protection stays with the file. By default, only the content owner is able to remove the protection from the file. The owner grants rights to other users to perform actions on the content, such as the ability to view, copy, or print the file.

Demonstration Steps

Prepare the AD RMS infrastructure

1. Open Active Directory Users and Computers, and then create the **ADRMSSRV** account. Add the account to the **Domain Administrators** group.
2. Create and share a folder named **Public**, and then select **Read** for the **Everyone** group.
3. Add email addresses for the Hani Loza user account and the Research group, both located in the Research organizational unit (OU). The addresses are hani@adatum.com and Research@adatum.com, respectively.

Install and configure AD RMS

1. Start the Add Roles and Features Wizard, and then install the **AD RMS** role.
2. After the wizard completes successfully, run the configuration required for the Active Directory Rights Management Services task in Server Manager, as stated under **Notifications**. Use the following values:
 - AD RMS Cluster: **Create a new AD RMS root cluster**
 - Configuration Database: **Windows Internal Database**
 - Service account:
 - Name: ADATUM\ADRMSSRVC
 - Password: Pa\$\$w0rd
 - Cryptographic mode: **Cryptographic Mode 2**.
 - Cluster Key Storage: **centrally managed**
 - Cluster Key Password: **Pa\$\$w0rd**
 - Cluster Web Site: Default web site
 - Cluster address: **http://LON-DC1.Adatum.com**
 - Licensor certificate: **LON-DC1**
 - SCP Registration: **Register the SCP now**.

After the configuration task completes successfully, sign out of LON-DC1, sign in again, and then open the Active Directory Rights Management Services console.

Guidelines for Designing AD RMS Clusters

When designing AD RMS clusters, follow these guidelines:

- Use a single-server cluster in a small environment with limited resources. A single-server cluster is sufficient for an environment in which you do not expect a large load on the AD RMS servers.
- Add servers in a cluster behind a single load-balanced URL to provide high availability. If AD RMS is a critical resource, you should make it highly available. Be aware that an AD RMS cluster is not the same as a failover cluster in Windows Server 2012. In the case of AD RMS, a single AD RMS server can be a cluster in its own right. You can achieve high availability by adding servers to the AD RMS cluster behind a single load-balanced URL, but not by configuring a failover cluster in Windows Server.
- Use a root-certification cluster only, whenever possible, and then join additional AD RMS servers to it. You cannot load balance across root-certification and licensing-only clusters.
- Create a licensing-only cluster for a complex environment or for an environment that has specific licensing requirements. To provide load balancing and to avoid having a single point of failure, we recommend that you separate the licensing feature by deploying a licensing-only cluster.
- Provide redundancy and load balancing for licensing by adding servers to the installation, and then creating a licensing-only cluster.

When designing AD RMS clusters, follow these guidelines:
 Use a single-server cluster in small environments
 In larger environments:
 Add servers in a cluster behind a single URL
 Use a root certification cluster, with additional AD RMS servers added as required
 Create a licensing-only cluster for a complex environment

Configuring High Availability for AD RMS Services

You should ensure that your AD RMS infrastructure is highly available, in most cases. One way you can ensure high availability of AD RMS is to deploy it onto Hyper-V virtual machines that you can easily back up and start. During day-to-day operations, applications such as Microsoft Office do not need to contact AD RMS to consume an active license. However, the loss of AD RMS would mean that an author could not create client licensor certificates for new content or that the system could not issue new end-user licenses. Therefore, providing high availability is desirable.

To achieve high availability for AD RMS services:

- Use a DNS CNAME for the RMS services URL
- Configure NLB for two or more AD RMS servers
- Ensure that the AD RMS database is configured for high availability, based on:
 - Microsoft SQL Server failover clustering
 - Log shipping
 - Database mirroring
 - AlwaysOn
- Consider sizing carefully:
 - Tests show a CPU can handle 100,000 license requests per hour at maximum load
 - To accommodate potential fluctuations in peak demand, cut by 50% or 50,000 requests per hour by configuring more CPUs or servers in the cluster

Network Load Balancing

As previously stated, the simplest form of a cluster is one AD RMS server. However, in more-complex deployments, you can configure multiple servers as a cluster behind a single, shared URL. To achieve high availability of the service, you should use Network Load Balancing (NLB) for those servers. When using servers in an NLB, you must create either a Domain Name System (DNS) alias or canonical name (CNAME), or at least a Host (A) record for the AD RMS cluster URL so that AD RMS clients can properly find the URL on all the servers in the cluster. Certain AD RMS tasks, such as group expansion, require AD RMS connectivity to a group-catalog domain controller, so ensure redundant access is available. Add the NLB feature after installing AD RMS.

Database Servers

AD RMS cluster servers integrate closely with the database server. The AD RMS database server stores configuration, logging, and AD DS information that AD RMS uses. AD RMS uses the following databases:

- **Configuration database.** The configuration database is a critical component of an AD RMS installation. It stores, shares, and retrieves all configuration data and other data that you need to manage a cluster's account certification, licensing, and publishing services. The way that you manage your configuration database directly affects the security and availability of rights-protected content. Each AD RMS cluster has one configuration database. The configuration database for the root cluster contains a list of Windows user identities and RACs. If AD RMS manages the cluster key centrally, AD RMS encrypts the certificate key pair with the AD RMS cluster key before it stores the certificate key pair in the database. The configuration databases for licensing-only clusters do not contain this information.
- **Logging database.** By default, for each root or licensing-only cluster, AD RMS installs a logging database in the same database server instance that hosts the configuration database. This database can grow quite large. Therefore, you need a plan to help maintain adequate service and performance.

While providing high availability to the AD RMS database is vitally important, immediate, uninterrupted service is not always necessary. Your scenario may not require that you apply a SQL Server failover cluster for AD RMS, but you can use one if the server is clustered for other reasons. Generally, failover clustering requires at least two SQL Server database servers and a storage area network (SAN) disk subsystem. Log shipping is a cheaper alternative to a failover cluster, but it does require some manual intervention should a SQL Server database server failure occur. As long as the database service can recover within a few minutes, most AD RMS functionality will continue. Database mirroring involves running the same insert, update, and delete operation that occurs on the principal database onto the mirror database as quickly as possible. You can accomplish this by sending a stream of active transaction log records to the mirror server, which applies log records to the mirror database, in sequence and as quickly as possible. Database

mirroring works at the level of the physical log record. Both log shipping and database mirroring require two SQL Server database servers. SQL Server 2012 has a new feature, AlwaysOn Availability Groups, which will replace database mirroring in future releases. AlwaysOn allows as many as four SQL Server database servers to host a mirrored database. However, you must configure these four SQL Server database servers in a failover cluster.

Windows Server 2012 features a redesign of the AD RMS setup process to enable better support for remote deployment of AD RMS and SQL Server database servers. In prior Windows Server releases, AD RMS Setup required that the account used to install the AD RMS server role had local administrator privileges on any computers that were hosting a SQL Server installation for support of AD RMS. This was because AD RMS needed to read SQL database settings from the Windows Registry during the setup process.

In Windows Server 2012, AD RMS now has the following requirements for SQL Server access:

- The AD RMS installer account must have System Administrator permissions in the SQL Server instance.
- For assistance in accessing and locating available SQL Server instances, the SQL Server Browser service must be running on the server computer where you plan to install AD RMS.
- AD RMS supports SQL Server named instances on both Windows Server 2008 R2 and Windows Server 2012. To use SQL Server named instances, you must ensure that the SQL Server Browser service is running on the database server.
- The computer running SQL Server that supports AD RMS must have firewall exceptions enabled for well-known ports used by SQL Server processes. For example, the SQL Server Browser service uses User Datagram Protocol (UDP) port 1434, and the default SQL Server Transmission Control Protocol (TCP) port is 1433. If you use these default ports for your SQL Server installation, they must have port exceptions in Windows Firewall.
- You must enable nondefault TCP ports for exceptions that are configured with your SQL Server installation, if you want these ports to access SQL Server instances. Usually, for default SQL instances, TCP port 1433 is assigned. If you have configured any SQL Server instances intended for use with AD RMS to use a nondefault TCP port, you must enable those ports for Windows Firewall exceptions, so that AD RMS setup can connect to your targeted SQL Server installation.

Additionally, if you want to be able to recover the AD RMS database in another system, you must configure your AD RMS database to use a DNS CNAME record that refers to the database server. Do not call the database server by its proper name. Use an alias during setup to point AD RMS to the proper database server.

AD RMS Sizing

Scalability on the AD RMS licensing servers should be your first consideration. Licensing servers are involved in the issuance of every content publishing or access license, and in the creation of the user's RACs. The load that they handle is proportional to the deployment's size, the number of documents to protect, and their consumption rate.

AD RMS is processor-intensive because it uses cryptography when it generates and validates certificates. Typically, the abilities of AD RMS server's central access policies are processor-bound. The AD RMS server has an almost linear performance increase from one processor up to four cores, and good scalability of up to 16 cores. Eight cores provide a very good cost/performance ratio.

Microsoft laboratory tests and production deployments show that each core in a typical server central processing unit (CPU) can handle 100,000 license requests per hour at maximum load, which may occur in a large organization. However, you should target 50 percent of maximum load, so that you can provide some room for fluctuations. In this case, 50,000 requests per hour should be the maximum load for each CPU core.

You can use random access memory (RAM) in AD RMS servers to cache certificates, store requests during processing, and manage the queue of outbound requests to SQL Server. Increasing the memory in each AD RMS server reduces the CPU load and increases the ability to cope with peak demand. Still, a configuration of 2 gigabytes (GB) of RAM is able to provide a significant amount of caching and is sufficient in most cases. 4 GB of RAM or more is necessary only for servers with a high number of processors.

Typically, an enterprise will not use a local disk intensively in the AD RMS servers unless it also uses a local installation of SQL Server or SQL Server Express (a configuration not recommended in high load environments). However, your disk must be fault tolerant to provide the availability that most environments require. Therefore, a two-drive Redundant Array of Independent Disks (RAID) array is the minimum recommendation for a production server. This configuration should serve most scenarios well. You can use virtualization, assuming the SQL Server workload is less than 27 requests per second, which should be sufficient if you monitor disk usage carefully.

Network bandwidth can be a limiting aspect for large servers, since each licensing request uses between 24 kilobytes (KB) and 30 KB. Therefore, a server that is receiving 200,000 license requests per hour can use more than 10 megabytes per second (Mbps) on the network. This scenario may occur when a large organization receives a protected email and many users attempt to read it at the same time. A high-end server can tax a 100 Mbps link easily if it is operating at peak load. Additionally, when AD RMS issues complex licenses, with many individual rights expressed, each license can grow as large as 1 megabyte (MB). In these cases, a large number of licenses can overwhelm the capacity of a 1 gigabyte per second network. This means that even if AD RMS can scale to a large number of processor cores per computer, the network interface can become a bottleneck for very large systems. We recommend that you do not exceed the recommended configuration of four to eight cores per server. You can add more individual servers when necessary.

Planning AD RMS Management

Windows Server 2012 specifies three new administrative roles that you can use to delegate control of your AD RMS environment. You can add these roles as local security groups when you install the AD RMS role. Each of the roles has a different level of access to AD RMS.

You can use the administrative roles to delegate AD RMS tasks without giving full administrative control over the entire AD RMS cluster. We recommend that you create an Active Directory security group for each administrative role, and then add the role to the corresponding local security group. This way, you can scale your AD RMS deployment across several servers without adding specific user accounts to each AD RMS server.

These new administrative roles are:

- AD RMS Enterprise Administrators. Members can manage all AD RMS policies and settings. During provisioning, AD RMS adds the user account that is installing the AD RMS server role to the AD RMS Enterprise Administrators role. As a best practice, you should assign this role only to user accounts that need full AD RMS administrative control.

AD RMS administrative roles:
 AD RMS Enterprise Administrators
 AD RMS Template Administrators
 AD RMS Auditor

Use these roles as a best practice:
 Do not assign Enterprise role to multiple people; usually used only once when installing AD RMS
 You can temporarily assign the auditors' role to governing agencies conducting reviews

- AD RMS Template Administrators. Members can manage rights policy templates. Specifically, members can read cluster information and list, create, modify, and export rights policy templates.
- AD RMS Auditors. Members can manage logs and reports. This is a read-only role. Members can only read cluster information and logging settings, and run reports that are available on the AD RMS cluster.

Microsoft recommends that you use these roles, and that you do not assign the AD RMS Enterprise Administrators role to multiple people. You do not use this role frequently, but only when you install AD RMS. You should assign the AD RMS Template Administrators role to users who manage rights and rights policy templates, and you should assign the AD RMS Auditors role to users who monitor AD RMS.

Implementing an AD RMS Backup and Recovery Strategy

When you set up AD RMS, you should consider how you would handle a loss of service. You can lose service unexpectedly if any AD RMS component fails: the intranet connection that the AD RMS server uses; the AD RMS root-certification cluster; the AD RMS licensing-only cluster; or even the database servers that host the AD RMS configuration databases.

You should develop a backup strategy for each of the following components, so that AD RMS is as highly available and reliable as possible:

- Intranet and Internet connectivity
- Certification and licensing pipelines
- Database servers
- Directory services

If you need to restore AD RMS, you can:

- Restore a database
- Restore a previous AD RMS installation



Additional Reading: To restore a database, follow the regular procedures for a database restoration on SQL Server.

See <http://go.microsoft.com/fwlink/?LinkID=285335> for a review of these procedures.

To restore a previous AD RMS installation, first you must determine whether the failed server was a member of a cluster and whether it was the root configuration server. If the server was a member of a cluster, you can remove and repair the node, reinstall the software (as necessary), and then rejoin the node to the cluster as though it were a new node. If the failed server was not a member of a cluster, you can recover the server by fixing the underlying problem (such as by obtaining new hardware), reinstalling the operating system and AD RMS software (as necessary), and then reprovisioning the server.

When you design a backup and restore strategy for AD RMS, you should:

- Identify the AD RMS components that you should back up
- Make a backup strategy for each component
- Choose a restore scenario (database or full AD RMS)
- Consider the ServiceConnectionPoint object
- Test your backup and restore strategies

If the server was the AD RMS root configuration server, and AD DS has a ServiceConnectionPoint object for AD RMS, the server tries to provision itself as a subenrolled, licensing-only server, which does not work. To troubleshoot this problem, you can:

1. Delete the ServiceConnectionPoint object manually from AD DS, and then restart the provisioning process.
2. Reprovision the server.
3. When the ProvisionLicensing.aspx page appears, replace the page name in the URL with ProvisionCertification.aspx, and then press Enter.

When you design a backup and restore strategy for AD RMS, you should:

- Back up and manage the AD RMS private key as part of the security plan, because AD RMS uses private keys to encrypt content.
- Maintain a current backup of the AD RMS database at all times.
- Provide a redundant Internet link for AD RMS if you are servicing external clients.
- Back up all certificates that are on AD RMS.
- Back up custom templates that are on AD RMS.
- Test your backup.

Decommissioning and Removing AD RMS

If you decide to remove the AD RMS role from a server, first you must decommission AD RMS. When you decommission AD RMS, the AD RMS cluster provides a key that decrypts the rights-protected content that it published previously. You can use this key to save the content without AD RMS protection. This is useful if you decide to stop using AD RMS in your organization, but you still need the protected information.

You should enable decommissioning on each server in the cluster long enough that users can save the content without AD RMS protection, and that network and system administrators can disable any AD RMS-enabled clients to stop them from using the service.

After you enable decommissioning, the AD RMS console shows only the Decommissioning server information page in the results pane. You cannot do anything else to administer AD RMS.



Note: You cannot reverse this process. After you decommission AD RMS, you must remove it completely by using Server Manager. Then, reinstall AD RMS.

There are instances when you need to retire an AD RMS server or remove an AD RMS cluster entirely. Before you do, you should back up all AD RMS databases that the server or cluster uses, especially the configuration database.

- Remember that decommissioning a server can decrypt all protected content

When you decommission and remove AD RMS, remember to:

- Decommission AD RMS before you remove the AD RMS role from a server
- Plan the time of the decommission period
- Keep in mind that you cannot return a decommissioned server to its previous state
- Ensure that you have a current AD RMS backup

Options

- Remove one server from a cluster
- Retire a stand-alone server
- Replace an AD RMS cluster with another existing cluster

After you back up the databases, you can remove the server. How you do this depends on the role of the server and the topology of the AD RMS installation. You can do one of the following:

- Remove one server from a cluster. If you want to retire an AD RMS server from a cluster that has other servers that are still active and required, deprovision and uninstall AD RMS from the server that you want to retire, and then remove the server from the load balancing rotation. Consult the load balancer documentation for more information about how to remove a server.
- Retire a stand-alone server. If you want to retire a server that is the only server in a cluster, first decommission, deprovision, and uninstall the server. Then, remove it from the network, and immediately install and provision AD RMS on a replacement server. Configure the new AD RMS server, which creates a new single-server cluster. Use the same URL and configuration database as the retired server. Until you install and provision the replacement server, users cannot read rights-protected content that the original single-server cluster published.
- Replace an AD RMS cluster with another, existing AD RMS cluster. You might need to retire an AD RMS cluster and replace it with another, existing AD RMS cluster, such as when two companies merge and both are running AD RMS. In this case, export the trusted user domain and trusted publishing domain from the AD RMS cluster that you are retiring. Then, import the trusted user domain and the trusted publishing domain into the AD RMS cluster that is still active. This enables you to ensure that users in the active cluster can read content that the retired AD RMS cluster was protecting.
- Consider moving AD RMS over to Windows Azure™ Rights Management. You can get a free tenant account with Windows Azure and migrate your AD RMS infrastructure over to Windows Azure Rights Management. A later lesson in this module will provide details about Windows Azure Rights Management.

Upgrading an AD RMS Cluster

You may have an existing AD RMS cluster running on an older server operating system, such as Windows Server 2003, Windows Server 2008, or Windows Server 2008 R2, and you wish to upgrade the operating system to Windows Server 2012 or Windows Server 2012 R2. You may wish to upgrade a Windows Server 2012 AD RMS cluster to newer hardware. In each case, you can do so, with careful planning. Your production AD RMS cluster may be providing rights management for thousands of documents and end users.

Upgrading incorrectly, or having the upgrade abort for any reason, can restrict these documents suddenly and completely for all authors and recipients. Also, an incorrect upgrade can reduce your organization's ability to perform rights management on new documents.

You should take three key preliminary steps before attempting an upgrade, depending on whether you are performing an in-place upgrade or a migration. In an in-place upgrade, you simply upgrade an existing AD RMS cluster server running the Windows Server 2008 or Windows Server 2008 R2 operating system to Windows Server 2012. In a migration, you add an additional Windows Server 2012 server to the existing Windows Server 2008 AD RMS cluster and then move the certificates from the older cluster servers to that server with the new operating system. In either scenario, you should back up the AD RMS configuration database first, and then, for a migration, export the server licenser certificate and, if you are using a software-based cryptographic service provider (CSP) to protect your AD RMS private key, export the key container and install it on the new server.

- You can upgrade from Windows Server 2008 or Windows Server 2008 R2 to Windows Server 2012 or Windows Server 2012 R2, and from Windows Server 2012 to Windows Server 2012 R2
- To upgrade from Windows Server 2003, you must first perform an upgrade to Windows Server 2008 or Windows Server 2008 R2
- Before performing any upgrades, complete the following tasks:
 - Back up the AD RMS configuration database
 - Export the SLC
 - Export and install the software-based CSP key container, if used

Back Up the AD RMS Configuration Database

All the AD RMS cluster configuration information, in addition to the private key that signs all rights-protected content, is stored in the AD RMS configuration database. Backing up this database before moving to Windows Server 2012 or Windows Server 2012 R2 will ensure that you will not lose your AD RMS configuration data and the private key permanently. In the event of any catastrophe while upgrading, you could restore the database on another server. To back up the AD RMS configuration database, perform a SQL Server backup task from the SQL Server Management Studio on the database named `DRMS_Config_<RMS_cluster_URL>_80`, where `<RMS_cluster_URL>` is the URL of your AD RMS cluster.

Export the Server Licensor Certificate

Legitimate users use the server licensor certificate of the AD RMS cluster to decrypt all AD RMS cluster-protected content. If the server licensor certificate is lost, users cannot decrypt rights-protected content that the AD RMS cluster protects. Note that if you are using a hardware security module (HSM) to store the server licensor certificate, you cannot backup the server licensor certificate through the AD RMS console. You need to contact the HSM hardware manufacturer and get instructions on how to back up the key. If you are using a private key password to protect the server licensor certificate, you can back up the certificate by using the AD RMS console. To do so, in the AD RMS console, right-click the AD RMS cluster name and go to Properties. In the AD RMS Cluster Name Properties sheet, select the Server Certificate tab, and then select Export Certificate. Then you should see the Export Certificate dialog box. Modify the bin file name to include the name of your AD RMS cluster. Specify the storage location, and click Save.

Export and Install the Software-Based CSP Key Container

One of the AD RMS install options allows you to select private key protection managed by AD RMS, or hardware-based or software-based CSP key protection. Because the AD RMS configuration database stores the AD RMS-managed private key, there is minimal administrative overhead. Servers that you add to the cluster share this key. A hardware-based CSP provides more security because the private key is not stored anywhere in the software. A software-based CSP stores the AD RMS private key locally on each AD RMS server. This option is not ideal, but you may need to implement it, depending on requirements. If you do so, take extra care to secure any server with a software-based CSP key. If you have a software-based CSP, you must export and then install the AD RMS private key on a new server that is joining the AD RMS cluster. Alternatively, you can upgrade the operating system on an existing AD RMS cluster after installing the AD RMS private key. If you are using a hardware-based CSP, you should consult the manufacturer about specific steps for migrating the key device. The following is a basic procedure for exporting and installing the software-based CSP key container:

1. Find the private key name:
 - a. Open the SQL Server Management Studio on the SQL Server computer hosting the AD RMS cluster configuration database.
 - b. Drill down to the AD RMS configuration database, and then expand **Tables**.
 - c. Find the table named `DRMS_LicensorPrivateKey`, and then click **Open Table**.
 - d. Note the name found in the **KeyContainerName** column.
2. Export the software-based CSP RMS private key:
 - a. Open an administrator-level command prompt.
 - b. Type `cd %windir%\Microsoft.NET\Framework\v2.0.50727`, and then press Enter.
 - c. Type `aspnet_regiis.exe -pi "<keycontainername>" privatekey.xml -pri` (where `<keycontainername>` is the key container name that you retrieved from step 1), and then press Enter.

3. Install the software-based CSP RMS private key on the new server, or if this is the server that you will upgrade, perform these steps after upgrading the operating system:
 - a. Open an administrator-level command prompt.
 - b. Type **cd %windir%\Microsoft.NET\Framework\v2.0.50727**, and then press Enter.
 - c. Type **aspnet_regiis.exe -pi "<keycontainername>" privatekey.xml -exp** (where *<keycontainername>* is the key container name that you retrieved from step 1), and then press Enter.

You may now proceed with a normal upgrade from Windows Server 2008 or Windows Server 2008 R2 to Windows Server 2012 or Windows Server 2012 R2, or an upgrade from Windows Server 2012 to Windows Server 2012 R2. To upgrade from Windows Server 2003 to Windows Server 2012, you must first upgrade to Windows Server 2008 or Windows Server 2008 R2.

Lesson 3

Planning and Implementing AD RMS Templates and Policies

AD RMS enhances an organization's security strategy by protecting information through persistent usage policies, or usage rights and conditions, which remain in place no matter where the information moves. AD RMS persistently protects data, so that its usage rights remain with the information, instead of residing on an organization's network. This allows usage rights to be enforced even after an authorized recipient accesses the information, whether online or offline, and whether inside and outside the organization. A rights policy template is a document that contains a predefined usage policy. You can apply the template to protect content, avoiding the need to define and apply policy manually. It is easier to apply a template to generate an issuance license than to recreate an entire policy. Typically, an administrator on the AD RMS server creates the templates, stores them in the configuration database, copies them to a shared folder, and distributes them by using Group Policy, Microsoft System Center 2012 Configuration Manager, or another mechanism. This lesson details how to plan and implement these templates and policies.

Lesson Objectives

After completing this lesson, you will be able to:

- Understand options for configuring AD RMS rights and policy templates.
- Plan for AD RMS template distribution.
- Understand options for configuring AD RMS exclusion policies.
- Add user entities to an exclusions policy.
- Plan the AD RMS Super Users group.
- Plan for AD RMS client applications.

Options for Configuring AD RMS Rights Policy Templates

AD RMS rights allow you to control how a user can access, use, and redistribute rights-protected content. AD RMS-enabled applications or browsers enforce some rights exclusively, and the AD RMS client enforces other rights primarily, although applications can apply their own interpretation of the rights. The rights that the AD RMS client enforces control how you can use license information, such as whether you can use the license to reencrypt previously decrypted content. AD RMS-enabled applications, such as those in Microsoft Office, interpret and enforce rights that control how users utilize content. For example, the Microsoft Office applications enforce the View right by allowing a user to decrypt and view the contents of a protected document if the user has the View right. When a user wants to protect content, AD RMS enables some rights by default, including Full Control, View, Edit, Save, Print, Forward, and Reply. Additional rights available include Export (Save as), Reply All, Extract, Allow Macros, View Rights, and Edit Rights. Furthermore, you can create custom rights.

AD RMS rights policy templates configuration options:
Specify which users or groups have rights to work with content that the template helps protect
Rights include Full Control, View, Edit, Save, Print, Forward, and Reply
Templates can contain the following information:
A name and description
Users and groups that can be granted content licenses
The rights and associated conditions granted to the users
The content expiration policy
A set of extended policies
The revocation policy
A revocation list
A revocation list refresh interval
A public key file for the revocation list

Rights policy templates control access by dictating the rights and conditions of use. You can distribute the templates in a few ways, depending on the operating systems an organization uses. Rights policy templates control the rights of users and groups for a particular piece of rights-protected content. The templates can include various conditions, such as specific recipients or AD DS groups, how long a use license for the content remains valid, and how long after publication a user can read the content. AD RMS rights are part of the rights policy template.

Templates can contain the following information:

- A name and description.
- Users and groups that can be granted content licenses.
- The rights and associated conditions granted to the users.
- The content expiration policy.
- A set of extended policies.
- The revocation policy.
- A revocation list.
- A revocation list refresh interval.
- A public key file for the revocation list.

Some examples of rights policy templates are:

- Company Confidential. You can use this template to allow employees to view content, but not to forward, copy, or save it.
- Expires in 30 days. You can use this to ensure that content is not valid after 30 days. Users can read a letter of offer, a request for proposal, or a draft version of a document for a limited time.
- Must be Connected to Consume. This ensures that recipients connect to a licensing server and are not using cached copies of a use license to consume content. You can use this if a template is subject to change and you want the recipient have read-only access to the most recent version. Additionally, if a computer is lost or stolen, the rights-protected content is not accessible to the person who found or stole the computer.

To ease the administration of rights policy templates, AD RMS in Windows Server 2012 provides a Create Distributed Rights Policy Template Wizard.

Question: Can an organization use AD RMS without using rights policy templates? If so, what are the limitations?

Planning for AD RMS Template Distribution

When publishing protected content, the author selects the rights policy template from the templates that are available on the local computer. To make rights policy templates available for offline publishing, you must deploy them to client computers from a shared folder. In Windows 8, Windows 7, Windows Vista with Service Pack 1 (SP1), Windows Server 2012, Windows Server 2008 R2, and Windows Server 2008, the AD RMS client manages rights policy templates automatically. The AD RMS polls a template distribution pipeline for updates. If you add, change, or delete a rights policy template, the AD RMS client detects the change and updates the local rights policy templates during the next refresh.

Method	How to install	Advantages	Disadvantages
Automatic propagation	<ul style="list-style-type: none"> No need to install Runs as a scheduled task 	Results in minimal user intervention	Supported only on Windows Vista with SP1 or newer
Local computer	<ul style="list-style-type: none"> Use a script to copy Use a custom Office 2003 Installation Pack to configure Use Group Policy, Configuration Manager 2007/2012, or Systems Management Server 2003, together with a template distribution job in WMI 	Enables users to protect documents by using the rights policy templates on the machine on which they are signed in	Requires that you update and maintain the template for each user on each computer
Shared folder	Define the internal path to locate templates	Provides a single point to access the rights-policy templates	Lack of connectivity on AD RMS-enabled clients can disable the templates, which removes them from the user's menu
Shared folder using offline folders	Define the internal path to locate templates	Provides a single point to access the rights policy templates and to configure offline folder use	<ul style="list-style-type: none"> Requires additional configuration settings Requires end-user training

Computers that are running Windows Vista with SP1 or newer versions store the rights policy templates in the %localappdata%\Microsoft\DRM\templates folder. Computers that are running Windows XP, Windows 2000, or Windows Server 2003 store the templates in the %appdata%\Microsoft\DRM\templates folder.

You can store the rights policy templates in the configuration database or in a shared folder. Users must be able to access the templates so that they can access rights-protected content. Many administrators prefer to place the rights policy templates on each client computer, which allows users to use the templates to publish rights-protected content both offline and online.

If you modify a rights policy template on the AD RMS server and configure the AD RMS cluster to specify a file location for storing copies of rights policy templates, the server updates the template in both the configuration database and the shared folder. If you use an AD RMS client other than Windows Vista with SP1 or newer versions and you modify a rights policy template, you must redeploy the template to the client computers, so that users have the most current version available.

If the computer runs a version of the AD RMS client that is older than Windows Server 2008 and Windows Vista with SP1, then you can use Group Policy, Microsoft System Center Configuration Manager 2012 or Configuration Manager 2007, or Microsoft Systems Management Server (SMS) 2003 to distribute rights policy templates.

Starting with Windows Server 2008, you can distribute rights policy templates to client computers, and you can archive any rights policy templates that you do not want to distribute. By default, all rights policy templates are distributed. You should not delete a rights policy template, because if you do, any content that the rights policy template protects will not be accessible.

The following table lists the three methods for distributing rights policy templates, and the advantages and disadvantages of each.

Method	How to install	Advantages	Disadvantages
Automatic propagation	No need to install. Runs as a scheduled task.	<ul style="list-style-type: none"> Results in minimal user intervention. 	<ul style="list-style-type: none"> Supported only on Windows Vista with SP1 or newer.

MCT USE ONLY. STUDENT USE PROHIBITED

Method	How to install	Advantages	Disadvantages
Local computer	<ul style="list-style-type: none"> • Use a script to copy. • Use a custom Office 2003 installation pack to configure. • Use Group Policy, Configuration Manager 2012, Configuration Manager 2007, or Systems Management Server 2003, together with a template distribution job in Windows Management Instrumentation (WMI). 	<ul style="list-style-type: none"> • Enables users to protect documents by using the rights policy templates on the machine on which they sign in. 	<ul style="list-style-type: none"> • Requires that you update and maintain the template for each user on each computer.
Shared folder	<ul style="list-style-type: none"> • Define the internal path to locate templates. 	<ul style="list-style-type: none"> • Provides a single point to access the rights policy templates. 	<ul style="list-style-type: none"> • Lack of connectivity on AD RMS-enabled clients can disable the templates, which removes them from the user's menu.
Shared folder using offline folders	<ul style="list-style-type: none"> • Define the internal path to locate templates. 	<ul style="list-style-type: none"> • Provides a single point from which to access the rights policy templates and to configure offline folder use. 	<ul style="list-style-type: none"> • Requires additional configuration settings. • Requires end-user training.

Options for Configuring AD RMS Exclusion Policies

AD RMS uses a lockbox to store a user's private key. If a vulnerability is found in a certain version of a lockbox, Microsoft will release a new lockbox based on the minimum version of the AD RMS client software and exclude lockboxes created by the previous versions of the AD RMS client software. Once you have enabled this feature, you specify the latest minimum lockbox version that the Microsoft Activation Service signed. You then enable lockbox exclusion on each AD RMS cluster on which you want it to take effect. All certification and licensing requests are checked to make sure that the lockbox meets the minimum version criteria. You can implement exclusion policies so that you prevent certain entities from getting certificates and requesting licenses. Windows Server 2012 AD RMS has three ways to set exclusion policies: by user, by application, and by lockbox version.

Options for configuring AD RMS exclusion policies:

Add entities to exclusion policies so that they cannot get certificates or licenses
 Remove entities from the exclusion list when the rationale to exclude no longer applies
 Exclude the following entities:
 Users: usually by email address
 Applications: by file name and extension
 Lockbox versions: by Microsoft version number

Once you implement an exclusion policy, any use license for that specific entity will be restricted by its presence in the exclusion list. If you decide to remove the restriction from the entity because you are establishing a trust or because of further direction from management, you can delete that entity from the exclusion list. As you add new certificates or receive requests for new licenses, the exclusion will no longer apply to that entity.

When you remove an entity from an exclusion list, be certain that all of the certificates that you issued before the exclusion policy listing for that entity have expired. If you do not, both the old and new certificates will allow the content to be decrypted, which your organization's management may not have intended.

Exclude Users

You can exclude a user account from obtaining use licenses from an AD RMS cluster by specifying either the user's email address or the public key string of the RAC associated with the user's RAC.

The following are examples of users who you may add to an exclusion list:

- You should exclude users by their email addresses if they are not allowed to consume rights-protected content but they do have email accounts in your AD DS forest.
- You can exclude a user who is trusted but whose AD RMS credentials become compromised. You exclude such a user by excluding only the compromised RAC's public key. When you do this, AD RMS denies new use-license requests that involve that RAC. After you exclude a RAC, the next time that user attempts to acquire a use license for new content, AD RMS will deny the request. To acquire a use license, the user will have to retrieve a new RAC with a new key pair.
- You can exclude external users who are not part of your AD DS forest, such as Windows Account users, federated users, and users identified by a trusted user domain. You can also specify a RAC to exclude their public keys.
- You should exclude a user on all licensing-only clusters if you add a user to the exclusion list of the AD RMS root cluster. Each AD RMS cluster has independent exclusion lists.

Exclude Applications

The following conditions apply to applications that you add to an exclusion list:

- You can specify the version of an AD RMS-enabled application against which AD RMS checks all licensing requests. Application exclusion stamps use a license with the condition that the license can bind only to the rights-protected content for which it is issued, if the application that is requesting the license is not on the excluded list.
- System administrators can use their usual mechanism to cause client computers to install the update, when deploying an update for an AD RMS-enabled application. Then administrators can set application exclusion policies that they define by using the version information of the application. This exclusion policy restricts AD RMS from issuing licenses to clients that are running older versions of the software.
- You must configure application exclusion on each cluster for which you want it to take effect.
- You can apply this exclusion policy on your cluster, but you should remember that clients cannot use the excluded application to request and bind new use licenses to rights-protected content. However, clients can continue to use the excluded application to consume previously licensed files.

Exclude Lockbox Versions

The following conditions apply to lockbox versions that you add to an exclusion list:

- Lockboxes store a user's private key. If a certain version of a lockbox is vulnerable, Microsoft will release a new lockbox. You can ensure that clients use a minimum version of the AD RMS client software by using the lockbox version associated with the client to exclude the older AD RMS client-software versions. When you enable this feature, you specify the latest minimum lockbox version that the Microsoft Activation Service signed. Then you enable lockbox exclusion on each AD RMS cluster on which you want it to take effect. AD RMS checks all certification and licensing requests to ensure that the lockbox meets the minimum version criteria.
- If you enable an exclusion based on lockbox version, clients that are using a version of the lockbox software that is older than the specified version cannot acquire RACs or use licenses, because AD RMS will deny their requests. These clients must install a new version of the AD RMS client software to acquire a new lockbox that uses the current software version.
- If a user with an excluded lockbox was issued licenses for content previously, the user can consume that content without acquiring a new lockbox.

Demonstration: Adding User Entities to an Exclusion Policy

You also can perform the task that this procedure describes by using Windows PowerShell:

1. To enable user RAC exclusion, type the following at a Windows PowerShell command prompt, and then press Enter:

```
Set-ItemProperty -Path <drive>:\ExclusionPolicy\User -Name IsEnabled -Value $true
```

2. To exclude a user's RAC by specifying the user's email address, type the following at a Windows PowerShell command prompt, and then press Enter:

```
New-Item -Path <drive>:\ExclusionPolicy\User -Name <user_name>@<domain>
```

Where *<drive>* is the name of the Windows PowerShell drive, *<user_name>* is the user name of the user whose RAC you are excluding, and *<domain>* is the domain portion of the email address of the user whose RAC you are excluding.

3. To exclude a user's RAC by specifying the user's RAC public key, type the following at a Windows PowerShell command prompt, and then press Enter:

```
New-Item -Path <drive>:\ExclusionPolicy\User -PublicKey "<key>"
```

Where *<drive>* is the name of the Windows PowerShell drive, and *<key>* is the user's RAC public key.

4. To stop excluding a user's RAC, type the following at a Windows PowerShell command prompt, and then press Enter:

```
Get-ChildItem -Path <drive>:\ExclusionPolicy\User
```

Where *<drive>* is the name of the Windows PowerShell drive. Note the ID of the user whose RAC you want to stop excluding.

5. At a Windows PowerShell command prompt, type the following, and then press Enter:

```
Remove-Item -Path <drive>:\ExclusionPolicy\User\<user_ID>
```

Where *<drive>* is the name of the Windows PowerShell drive, and *<user_ID>* is the ID of the user you found in the previous step.

Demonstration Steps

Add Domain Users Group to local Remote Desktop Users group

1. Sign in to LON-CL1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. Add the **Domain Users** group to the local **Remote Desktop Users** group in the **System Properties**, and then sign out.

Exclude RACs

1. Sign in to LON-CL1 as **Adatum\Hani** with the password **Pa\$\$w0rd**.
2. In Internet Explorer, add the URL **<http://lon-dc1.adatum.com>** to the **Local Intranet** sites.
3. Open and protect a Microsoft Word document named **ADRMS-TST.docx**. Add the following text: **Research employees can read this document, but they cannot change, print, or copy it on the blank document page.**
4. Protect the document with the following: **Research@adatum.com** has Read permissions.
5. Save the document to **\\LON-DC1\ConfidentialResearch**.

Create an RAC exclusion

1. Open the AD RMS Administration Console.
2. In the **Exclusion Policies** page in the **Users** node, click **Enable User Exclusion**, and then exclude **hani@adatum.com**.

Stop excluding RACs for specific users

- In **Exclusion Policies, Users** node, delete hani@adatum.com.

Planning the AD RMS Super Users Group

The AD RMS Super Users group is a special group that has full control over all the rights-protected content that the cluster manages. Group members of the Super Users group have full owner rights in all use licenses that the AD RMS cluster manages. This means that members of the group can decrypt any rights-protected content file and remove its rights protection.

You can use the Super Users group to help recover data. For example, you can use it if the author of a document leaves the organization or if all access to a document expires. We recommend

that if you enable the Super Users group, you also enable success and failure auditing for audit account management and audit directory services access. This enables you to track membership changes. In addition, we recommend that you enable the Super Users group only if you need to recover data.

By default, the Super Users group is not enabled and has no members. If you enable it by configuring the Super Users setting in the AD RMS console, you can specify an AD DS universal group to use as the Super Users group for AD RMS. The group must be in the same forest as the AD RMS installation. Any user accounts that belong to the group that you specify as the AD RMS Super Users group gain the permissions of the Super Users group automatically.

When you plan a Super Users group, remember:

- The AD RMS Super Users group is a special group that has full control over all content that AD RMS manages
- You can use the Super Users group to recover data
- By default, the Super Users group is not enabled
- If you enable the group, specify a universal group to be the AD RMS Super Users group
- An Exchange Server can belong to the Super Users group
- Members of the Super Users group have full owner rights in all use licenses

If you integrate AD RMS with Microsoft Exchange Server 2010 or Exchange Server 2013, you must perform additional steps. If you want to enable Exchange Server to decrypt messages and attachments that the AD RMS cluster protects, you must add a user account that represents the Exchange Server to the AD RMS Super Users group. However, first you must create a mail-enabled universal group (distribution list) that contains the Federated Delivery Mailbox user account (named FederatedEmail.4c1f4d8b-8179-4148-93bf-00a95fa1e042). Alternatively, you can add that account to an existing mail-enabled universal group that you are using as the Super Users group in AD RMS. Because the Federated Delivery Mailbox user account is a system mailbox, it is not visible in the Exchange Management Console or the Exchange Administration console. However, you can access the Federated Delivery Mailbox by using the Exchange Management Shell.

Windows PowerShell Equivalent Commands

To set up a Super Users group:

1. At a Windows PowerShell command prompt, type the following cmdlets, and then press Enter after each:

```
Set-ItemProperty -Path <drive>:\SecurityPolicy\SuperUser -Name IsEnabled -Value $true
Set-ItemProperty -Path <drive>:\SecurityPolicy\SuperUser -Name SuperUserGroup -Value
<group_e-mail>
```

Where *<drive>* is the name of the Windows PowerShell drive, and *<group_email>* is the email address of the universal group that you want to designate as the AD RMS Super Users group. For example, to designate a group named SecurityAdmins at fabrikam.com as the AD RMS Super Users group for a Windows PowerShell drive named Z, you would type the following Windows PowerShell cmdlets, and then press Enter after each:

```
Set-ItemProperty -Path Z:\SecurityPolicy\SuperUser -Name IsEnabled -Value $true
Set-ItemProperty -Path Z:\SecurityPolicy\SuperUser -Name SuperUserGroup -Value
securityadmins@fabrikam.com
```

2. You should ensure that you enable the Super Users group only when you require it. Disable it when you no longer need it.

To disable the Super Users group:

- At a Windows PowerShell command prompt, type the following, and then press Enter:

```
Set-ItemProperty -Path <drive>:\SecurityPolicy\SuperUser -Name IsEnabled -Value
$false
```

Where *<drive>* is the name of the Windows PowerShell drive.

Planning for AD RMS Client Applications

You can integrate AD RMS with other Microsoft products to provide automated and semiautomated protection of various types of documents. Most frequently, enterprises integrate AD RMS with Exchange Server and SharePoint.

Integrating AD RMS with Exchange Server

You can integrate Exchange Server 2010 and Exchange Server 2013 with IRM technology, which can help your organization avoid disclosing sensitive information through email. When combined with AD RMS in Windows Server 2012, Exchange Server 2010 and Exchange Server 2013 can help automatically secure messages that contain sensitive information, while enabling you to scan and archive those same messages unencrypted.

AD RMS enables the following features of Exchange Server 2010 and Exchange Server 2013:

- Transport Protection Rules. These rules help protect email messages automatically with IRM.
- Transport Decryption. This feature gives trusted agents plain text access to messages that IRM helps protect. You can use this to archive messages and attachments, and to scan them for malware.
- Journal Report Decryption. With this feature, you can journal IRM-protected messages.
- IRM Decryption for Search. Exchange Search can use this feature to index content in messages that IRM helps protect.
- IRM-enabled Microsoft Outlook Web App. This feature enables you to send and open messages that IRM helps protect, by using Microsoft Outlook Web App in any supporting browser, without needing to install client software.
- IRM-enabled Unified Messaging (UM). With this feature, users can listen to protected voicemail messages in Outlook Web App, in Outlook, and on the telephone. This feature provides a Do Not Forward policy for private voicemail.
- Prelicensing attaches a prelicense to protected messages. This way, the client does not need to make repeated trips to the AD RMS server to retrieve a use license. In addition, it enables users to view IRM-protected messages and attachments offline, which in turn enables them to view the messages in Outlook Web Application.

Integrating AD RMS with SharePoint

If you want to integrate SharePoint 2010 or SharePoint 2013 with AD RMS, you must install AD RMS and SharePoint on separate servers. You cannot install AD RMS and SharePoint on the same server, because they need to control the same web server resources.

When you use SharePoint, users can collaborate easily on documents by posting them to a SharePoint site, from which they can access the documents over the corporate network. The goal of integrating SharePoint with AD RMS is to protect documents that users download from the SharePoint server.

However, it is important to note that integrating SharePoint with AD RMS does not protect the documents while they are on the server. When a user uploads a document to a SharePoint site, the server removes all protection until it receives a download request. At that time, the server applies the appropriate restrictions to the document before the client computer downloads it.

- AD RMS can integrate with:
 - Exchange Server 2010 and newer
 - Office SharePoint Server 2010 and newer
- Integration between Exchange Server and AD RMS:
 - Provides protection with transport rules
 - Enables journal report decryption
 - Enables IRM decryption for Search
 - Enables IRM Outlook Web Application
- Integration between SharePoint and AD RMS:
 - Helps protect documents on a SharePoint site
 - Provides strong cryptography for SharePoint

Utilizing Strong Cryptography for AD RMS

AD RMS enables you to increase cryptographic strength by invoking strong cryptography through an advanced mode known as cryptographic mode 2.

AD RMS in strong cryptographic mode provides an updated and enhanced implementation to support much stronger encryption and cryptographic keys. For example, while in mode 2 operation, Rivest-Shamir-Adleman (RSA) encryption is enhanced from 1024-bit encryption to 2048-bit encryption. Cryptographic keys are also longer. The old mode used 160-bit keys, while strong cryptographic mode uses 256-bit keys. Additionally, there is an option to use Secure Hash Algorithm 2 (SHA-2) hashing algorithm (SHA-2/SHA-256) standards.

Strong cryptography in AD RMS enables your organization to satisfy regulatory compliance with current security standards that the National Institute of Standards and Technology (NIST) sets. NIST issued Special Publication 800-57, which recommends the use of 2048-bit RSA keys. United States federal agencies are required to comply with NIST recommendations, per U.S. governmental regulations. Additionally, many private enterprises and other countries implement this recommendation.

Lesson 4

Planning and Implementing External Access to AD RMS Services

Most organizations form partnerships and other commercial relationships with suppliers, customers, and vendors. These relationships often require that they share confidential information via email and other digital delivery mechanisms. This confidential information can include project plans and sensitive intellectual property. The organizations in these relationships rarely have a common identity and security infrastructure. This makes it difficult to guarantee secure communication when a person shares confidential information with someone in the partner organization. AD RMS provides several alternatives to help partner organizations share confidential information while maintaining independent infrastructures.

In this lesson, you will learn about options for planning and implementing AD RMS access for external users. You will also learn how to integrate applications with AD RMS.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe options for enabling external users to access AD RMS.
- Describe options for enabling applications for AD RMS clients.
- Plan trusted user domains.
- Plan trusted publishing domains.
- Integrate AD RMS with Microsoft Federation Gateway support.
- Integrate AD RMS–protected documents by using Microsoft Account.
- Integrate AD RMS with AD FS.
- Describe further considerations for enabling external users to access AD RMS.

Options for Enabling External Users to Access AD RMS

External users are either employees who do not connect to your corporate network or users from other organizations. You should ensure that both types of users can access AD RMS. Users can share rights-protected content across internal or external AD DS forests. However, when an organization with multiple AD DS forests deploys AD RMS, and when it deploys certification servers in each forest that is hosting user accounts, configuration of trust policies will be necessary. You can support multiple forests and provide access to external users by using any of the following:

Use trust policies to allow an AD RMS cluster to process licensing requests for content that another AD RMS cluster is protecting

You can define trust policies for the following:

- Trusted user domains
- Trusted publishing domains
- Federated Gateway support
- Windows Live ID
- Federated trust

- Trusted user domains. This type of trust policy enables an AD RMS server to trust the RACs that another AD RMS root certification server generates. It also issues use licenses to users who have RACs from another organization's AD RMS server.

- Trusted publishing domains. A trusted publishing domain allows one AD RMS cluster to issue use licenses against publishing licenses that another AD RMS cluster issues. To add a trusted publishing domain, import the server licensor certificate and the private key of the server to trust.
- Microsoft Federated Gateway support. Microsoft Federation Gateway support allows an AD RMS cluster to federate to the Microsoft Federation Gateway. The Federation Gateway acts as a trusted broker between organizations. By establishing these federation relationships, organizations can enable applications such as Exchange Server 2010 with SP1 and Exchange Server 2013 to create AD RMS–secured messages that users who belong to an external organization can access.
- Microsoft Account credentials. People who receive rights-protected content but who do not have their own AD RMS infrastructure can use Microsoft Account credentials to obtain a use license. The content creator can send an email to the recipient that specifies the rights that are assigned to the recipient's Microsoft Account credentials.
- A federated trust. You can use a federated trust to authenticate external users through AD FS. After AD FS authenticates the users, AD RMS enforces policies and automatically provides the external users with the appropriate content licenses for the protected content.

Question: In what scenarios do you need to provide external access to AD RMS?

Options for Enabling Application Access for AD RMS Clients

Supported AD RMS–Enabled Office Applications

Several applications in the various Microsoft Office suites support rights management.

The following table describes the AD RMS–enabled applications that specific versions and editions of Microsoft Office support.

Options for enabling application access include:

- Supported AD RMS–enabled Office applications
- XML Paper Specification
- Office Viewers, XPS Viewers, and Rights Management Add-on

AD RMS–enabled applications	Microsoft Office 2003	Microsoft Office 2013, Microsoft Office 2010, Microsoft Office 2007	Microsoft Office Mobile
Microsoft Word	√	√	√
Microsoft Excel®	√	√	√
Microsoft PowerPoint®	√	√	√
Microsoft Outlook	√	√	√
Microsoft InfoPath®	Not provided	√	Not provided

Not all Office editions can produce rights management–protected documents, but all editions of the rights management–enabled versions of Office (Office 2013, Office 2010, Office 2007, and Office 2003) can consume rights management–protected documents.

This table shows the rights management capabilities of the different Office editions and associated viewers that are enabled with rights management.

Office version	Edition	Create IRM-protected Word, Excel, and PowerPoint documents	Read and edit IRM-protected Word and Excel documents	Read and edit IRM-protected PowerPoint documents	Send Rights protected email messages	Receive and reply to rights-protected email messages	Create or use protected InfoPath forms
Office 2007	Home and Student		√	√			
	Standard		√	√		√	
	Small Business		√	√		√	
	Professional		√	√		√	
	Professional Plus	√	√	√	√	√	√
	Enterprise	√	√	√	√	√	√
	Ultimate	√	√	√	√	√	√
	Downloadable Office Readers		√ (read-only)				
Office 2010, Office 2013	Professional Plus	√	√	√	√	√	√
	Professional	√	√	√	√	√	√
	Home and Business		√			√	
	Home and Student		√				
	Professional Academic (Office 2010) University (Office 2013)		√			√	
	Word and Excel Viewers		√ (read-only)				

XML Paper Specification

XML Paper Specification (XPS) is a Microsoft specification that describes the design of the XPS Document file format, a representation of electronic paper based on XML. The XPS Document format is an open, cross-platform document format that lets users to create, share, print, and archive paginated documents effortlessly.

Applications that are running on Windows XP and newer versions can create XPS document. Users can view these documents if they have Windows XP or a newer operating system, with .NET Framework 3.0 SP1 or one of the stand-alone XPS viewers that are available for download. Windows Vista and newer operating systems have an XPS viewer installed by default.

The file name extension of XPS documents is .xps. Any application that can print documents to the Microsoft XPS Document Writer virtual printer can create XPS documents. Alternatively, you can create them by using the Save As option in Office 2007 or newer, and then choosing XPS Document. This enables you to extend rights management to the other applications in Office 2007 or newer versions. For example, you can save a Microsoft Visio® 2010 design as an XPS document, which ensures that the document is rights-managed. This also applies to Office 2007 or newer versions of Microsoft Access®, Microsoft Office Publisher, and Microsoft OneNote®. If you do not have Microsoft Office 2007 or a newer version, you can create rights-managed XPS documents by using the free Microsoft XPS Viewer.

Office Viewers, XPS Viewers, and Rights Management Add-On

Enforcement of rights occurs at the application level. You must use an AD RMS-enabled application, such as Office 2010, Office 2007, or Office 2003, to create and consume rights-protected information. For users who are not running a compatible version of Office, Microsoft created the Microsoft Office Viewer, the XPS Viewer, and a free Rights Management Add-on for Internet Explorer that enables users to view protected information while still enforcing rights.



Additional Reading: You can download these free from the Microsoft website, at <http://go.microsoft.com/fwlink/?LinkID=285336>.

Rights Management Add-on for Internet Explorer

The Rights Management Add-on for Internet Explorer offers a way for users of supported Windows operating systems to view, but not alter, files with restricted permission. These restrictions, as with all RMS-protected content, prevent unauthorized entities from forwarding, editing, or copying sensitive documents, web-based information, and email messages. These restrictions provide protection while the information is in transit and after the recipient receives it.

Planning Trusted User Domains

By default, AD RMS does not service requests from a user if another AD RMS cluster issued the user's RAC. However, at times you need to enable AD RMS to service these kinds of requests, so that users can share AD RMS-protected documents between organizations that have their own AD RMS clusters. To do this, you can add AD RMS domains to a list of trusted user domains in an AD RMS cluster.

A trusted user domain is a trust between AD RMS clusters that instructs a licensing server to accept RACs from an AD RMS server in a different Active

Directory forest. An AD RMS trust is not the same as an Active Directory trust, but it is similar because it allows one environment to accept another environment's identities as valid.

To add a trusted user domain, import the server licenser certificate from an AD RMS cluster that you want to trust to the AD RMS cluster that trusts. To do this:

1. Ensure that each AD DS forest has its own AD RMS cluster.
2. Ask the administrator of the domain that you want to trust to export that domain's service licenser certificate from his or her AD RMS cluster. Then ask him or her to send the certificate file to you, as the administrator of the trusting AD RMS cluster.
3. Import the certificate by specifying the file location. You do not transfer the private key information when you set up a trusted user domain.
4. Specify which email domains you want to trust for each trusted user domain. This is an important security step because it prevents a user from a trusted domain from impersonating an internal user.

Trusted user domains are unidirectional, so if users in each forest need to access documents in the other forest, you need to perform this process twice (once per forest).

You establish a trust offline, so you do not need to connect directly to the AD RMS clusters or their Active Directory forests. However, if you want to access content in the other forest, a user needs to access the licensing cluster that issued the publishing license. Therefore, the AD RMS cluster that trusts and publishes the content must be accessible either from the Internet or through a private network.

You can establish a trusted user domain between forests that do not have an Active Directory forest trust between them. However, if both forests are part of the same organization, you might want to establish an Active Directory forest trust also. This allows you to authenticate remote users with their own credentials, instead of enabling anonymous access to the corresponding pipeline URL in the trusting AD RMS cluster.

AD RMS uses trusted user domains to process requests for use licenses from users whose RACs are issued by another AD RMS cluster:

- A trusted user domain is a trust between AD RMS clusters
- To add a trusted user domain, import the server licenser certificate
- Trusted user domains are unidirectional
- You do not need to connect the AD RMS clusters directly to establish a trusted user domain
- In different forests, a trusted user domain is a one-way trust
- If you create an AD DS Forest Trust, you will simplify AD RMS trusts across forests

A forest trust makes group expansion possible, which in turn makes it much easier for you to apply rights to documents that have large audiences. If you want to enable group expansion and you do not have a forest trust, synchronizing group membership between forests will create significant additional work for you.

In an organization that has multiple forests, you can synchronize the global address lists (GALs) between the forests to make it easier for content creators to identify recipients.

Planning Trusted Publishing Domains

With a trusted publishing domain, an AD RMS cluster can issue use licenses that correspond to publishing licenses that another AD RMS cluster issues.

By default, an AD RMS cluster can issue use licenses that correspond only to the publishing licenses from that same cluster. However, if you configure a trusted publishing domain, the AD RMS cluster can grant use licenses for content that another AD RMS cluster protects, in the same organization or a separate one. You can configure any number of trusted publishing domains for an AD RMS cluster.

A trusted publishing domain is useful if a company acquires another company that already has an AD RMS implementation that it must deprovision, or if an organization eliminates a cluster from an AD RMS implementation. With a trusted publishing domain, AD RMS can issue end-user licenses and client licensor certificates from a single point. Users can still read documents that both AD RMS clusters protect, even if you deprovision one cluster.

When you use trusted publishing domains, you:

- Can use a trusted publishing domain to allow an AD RMS cluster to issue use licenses that correspond to publishing licenses that another AD RMS cluster issues:
 - Most companies use a trusted publishing domain only if they acquire another company that has AD RMS already
- Must share the private key for the server
- Must configure the registry in the client
- Must synchronize the rights policy templates



Note: Configure a trusted publishing domain carefully. Sharing the private key requires that you trust the recipient strongly.

To add a trusted publishing domain:

1. Export the service licensor certificate, the private key, and all rights policy templates from the original AD RMS cluster to a password-protected file.
2. Import these items to the AD RMS cluster in the domain that you want to trust by specifying the file location and password.



Note: If the private key is in a hardware security module (HSM), you must use a special procedure to export the private key. However, it might not be possible on some HSMs.

If a client needs to obtain a use license from the local AD RMS cluster, you must configure the registry in the client to map the original path of the licensing server to the new path. This way, users can obtain use licenses without accessing the original licensing server, which may reside in a separate network or may no longer exist.

To configure the registry to map correctly, follow this procedure:

1. Open **regedit.exe**, and then add the **DRM** key and the **AdminTemplatePath** subkey with the value set to the local user's template folder, **%LocalAppData%\Microsoft\DRM\Templates**.
2. Use this registry key: **HKEY_CURRENT_USER\Software\Microsoft\Office\14.0\Common\DRM**.
3. Close **regedit.exe**.

Because the server that hosts the subjects of the use licenses performs group expansion, you do not need a forest trust if you use a trusted publishing domain. Therefore, trusted publishing domains have less administrative and traffic overhead than trusted user domains. However, you might want to establish a forest trust, so that you can synchronize the GALs, which makes it easier for a content creator to identify recipients.

If you have a trusted publishing domain, you must synchronize rights policy templates periodically, so that they are up to date in the destination system. You do not need to synchronize the templates if you use a trusted publishing domain to replace a cluster that you remove from service.

If you implement a trusted publishing domain, you must share the private key for a cluster. However, this is a significant risk because a malicious user can use the key to spoof an AD RMS cluster and then decrypt content that the cluster manages and licenses. Typically, you implement trusted publishing domains within a single organization, due to this security risk. Most organizations implement a trusted publishing domain to issues licenses for documents that a deprovisioned cluster protects.

Integrating AD RMS with Microsoft Federation Gateway Support

Microsoft Federation Gateway Support is a feature of AD RMS introduced in SP1 for Windows Server 2008 R2. Microsoft Federation Gateway Support enables an AD RMS cluster to unite with the Microsoft Federation Gateway, which acts as a trusted broker between organizations. By establishing these federation relationships, organizations can configure applications, such as Exchange Server 2010 with SP1 and Exchange Server 2013, to be able to create AD RMS-secured messages that users from an external organization can access.

When integrating Microsoft Federation Gateway Support, remember that it:

- Is an identity service **that runs** on the Internet
 - Mediates between an organization and external services
- Uses an SSL certificate that** must be from a Microsoft-approved CA

If you are using multiple AD RMS clusters using Microsoft Federation Gateway for different purposes, ensure that you use a SAN with a different URL for each

Virtual directories use http, not https, for addresses to Microsoft Federation Gateway Support

Microsoft Federation Gateway is an identity service that runs over the Internet. It mediates between an organization or business and the external services that the organization wants to use. The gateway connects users and other identities to the services with which it works. In this way, administrators must manage only a single identity-federation relationship to enable identities to access all Microsoft services that they allow.

Microsoft Federation Gateway uses SSL certificates to prove domain ownership. It provides applications with a simple, standards-based method of establishing trust between separate organizations, because the organizations federate with the gateway instead of each other. Microsoft Federation Gateway makes it easier for an organization to establish trust relationships with multiple partners than when it uses conventional one-on-one federation or other trust relationships. You can control the scope of AD RMS by creating allow or deny lists of users and domains for licensing, in addition to specifying the domains that can receive publishing licenses. This ensures that only the organizations that you approve receive access to protected information.

Consider the following before you install AD RMS with Microsoft Federation Gateway:

- Configure the AD RMS cluster to use an SSL-encrypted connection that uses a certificate that the Microsoft Federation Gateway trusts. To prove ownership of the domain that you want to federate, you must own the X.509 SSL certificate for that domain. This certificate must be from one of the trusted root certification authorities (CAs) that are configured in the Microsoft Federation Gateway.
- Ensure that the SSL certificate that you enroll with Microsoft Federation Gateway shows ownership of the AD RMS cluster's extranet URL. If the AD RMS cluster is configured with an intranet URL that is different from the extranet URL, and if the intranet URL is not a domain name that users can access from the Internet, you must install the SSL certificate associated with the extranet URL on this AD RMS server. Then you must select that certificate when enrolling with the Microsoft Federation Gateway.
- Ensure that if the SSL certificate contains a subject alternative name (SAN), the last entry in the SAN is the fully qualified domain name of the domain you want to enroll.
- To avoid conflicts, do not enroll your AD RMS cluster with the Microsoft Federation Gateway by using the same URL that you use to federate another resource. Microsoft Federation Gateway can always have federated relationships to Microsoft Online and Microsoft Exchange Server. If you have used the URL that your AD RMS cluster uses as its external URL to federate with the Microsoft Federation Gateway for another purpose, you must enroll the AD RMS cluster with the Microsoft Federation Gateway. Do this by creating and using a certificate that contains the AD RMS URL as the last entry in the SAN and a common name (CN) that is not the same as the registered resource. For example, if the DNS name of your AD RMS server is resource.contoso.com, and if you have used that name to federate another resource to the Microsoft Federation Gateway, you can create a certificate in the following format to avoid federation conflicts:

```
Copy Subject: CN=adrmsservice.adatum.com
SAN:
DNS Name=adrmsservice. adatum.com
DNS Name=resource. adatum.com
```

- If you are not using SSL, configure your firewall to enable http data to pass through, because the virtual directories that Microsoft Federation Gateway Support use http://. However, the http:// transactions for Microsoft Federation Gateway Support use message-level security.

Integrating AD RMS with a Microsoft Account

AD RMS utilizes each user's email address for identification. Normally, only users within the same email domain can access protected content, which often means that only users within the same company can consume protected content. Previous topics explain how to configure AD RMS to work outside your organization's boundaries, if the other party has a separate AD RMS infrastructure.

However, if you want to share AD RMS-protected content with users who do not have their local AD RMS infrastructure, you cannot take advantage of technologies such as trusted user domains or trusted publishing domains.

Note that Windows Live ID is now called a Microsoft Account. When recipients of IRM documents use Microsoft Account to read AD RMS-protected content, remember:

- You need to establish a trust policy between AD RMS and Microsoft Account
- The user must have a Microsoft Account
- Anonymous access to the AD RMS IIS licensing service is required
- This service is not guaranteed to be permanent
- The recipient can only consume content, not protect it

If you want to share content with a user at an organization that does not have AD RMS in place, the user can access the content by using a Microsoft Account, which is based on a Windows Live® ID, Outlook, or Hotmail® email address. Microsoft Account is a free, cloud-based email and identity service from Microsoft that allows anyone to open and use an account for email and other services. Microsoft provides an AD RMS service for Microsoft Accounts, which you can integrate with AD RMS by setting up a trust. Microsoft Account is the new name for what used to be called *Windows Live ID*.



Note: Microsoft provides the Microsoft Account certification and licensing service as a temporary service, and reserves the right to stop providing the service at any time.

If you set up a trust with Microsoft Account, an AD RMS user can send rights-protected content to a recipient who has a Microsoft Account. The recipient can read the content, and the sender can protect it and apply specific permissions for Microsoft Account users.

This solution works only if the external users are few in number and need to read content only, not protect it. Recipients cannot protect the content themselves.

Microsoft provides an account certification service that uses Microsoft Account to establish the user's RAC. If you want users with RACs from that service to be able to obtain use licenses from your AD RMS cluster, set up a trusted user domain that accepts user credentials from the online Microsoft AD RMS service.

Users with Microsoft Account–based RACs will not be able to acquire licenses unless you configure anonymous access in Internet Information Services (IIS) to the AD RMS licensing service. This is essential because, by default, the licensing service is configured to use Integrated Windows authentication.

After configuring Microsoft Account access, you can exclude users of this service based on their email addresses, if necessary.

AD RMS utilizes each user's email address for identification. Normally, only users within the same email domain can access protected content, which typically means only users within the same company can consume protected content. Previous topics explain how to configure AD RMS to work outside your organization's boundaries, if the other party has a separate AD RMS infrastructure.

However, if you want to share AD RMS–protected content with users that do not have their local AD RMS infrastructure; you cannot leverage technologies such as trusted user domains or trusted publishing domain.

If you want to share content with a user at an organization that does not have AD RMS in place, the user can access the content by using a Microsoft Account email address. Microsoft Account is a free, cloud-based email and identity service from Microsoft that allows anyone to open an account, and then use it for email and other services. Microsoft provides an AD RMS service for Microsoft Account accounts, and you can integrate this service with AD RMS by setting up a trust.



Note: Microsoft provides the Microsoft Account certification and licensing service as a temporary service, and reserves the right to stop providing the service at any time.

If you set up a trust with Microsoft Account, a user of AD RMS can send rights-protected content to a recipient who has a Windows Live ID. The recipient can read the content, and the sender can protect it, and then apply specific permissions for Microsoft Account users.

This solution works only if the number of external users is small and they only need to read content, not protect it. Recipients cannot protect the content themselves.

Microsoft provides an account-certification service that uses Microsoft Account to establish the user's RAC. If you want users with RACs from that service to be able to obtain use licenses from your AD RMS cluster, set up a trusted user domain that accepts user credentials from the online Microsoft AD RMS service.

Users with Microsoft Account-based RACs will not be able to acquire licenses unless you configure anonymous access in IIS to the AD RMS licensing service. This is essential because, by default, the licensing service is configured to use Windows Integrated authentication.

After configuring Microsoft Account access, you can exclude users of this service based on their email addresses, if necessary.

To enable anonymous access to the AD RMS licensing service:

1. Sign in to a server in the AD RMS cluster.
2. Open the IIS Manager console, and then expand the server that is hosting AD RMS.
3. In the console tree, expand **sites**, and then expand the website on which you have configured AD RMS. By default, this is the default website.
4. In the console tree, expand the **_wmcs** virtual directory, right-click the licensing virtual directory, and then click **Switch to Content View**.
5. In the results pane, right-click **license.asmx**, and then click **Switch to Features View**.
6. In the results pane, double-click **Authentication** to open the **Authentication** page.
7. Click **Anonymous Authentication**, and then, under **Tasks**, click **Enabled**.
8. Repeat steps 1 through 7 for each server in the AD RMS cluster.

To trust Microsoft Account-based rights-account certificates (note that the AD RMS console still uses the term Windows Live ID, but you will use the actual Microsoft Account certificate instead):

9. Sign in to a server in the AD RMS cluster.
10. Open the AD RMS console, and then expand the AD RMS cluster.
11. In the console tree, expand **Trust Policies**, and then click **Trusted User Domains**.
12. In the Actions pane, click **Trust Windows Live ID**. The Windows Live ID certificate appears in the trusted user domain list in the results pane.

To specify Microsoft Account email domains to exclude:

1. Sign in to a server in the AD RMS cluster.
2. Open the AD RMS snap-in, and then expand the AD RMS cluster.
3. In the console tree, expand **Trust Policies**, and then click **Trusted User Domains**.
4. Select the **Windows Live ID** certificate in the results pane, and then in the Actions pane, click **Properties**.
5. Click the **Excluded Windows Live IDs** tab.
6. Type the email domain that you want to exclude.
7. Click **Add** to add the specified object to the exclusion list.
8. Repeat steps 5 through 7 for all email domains that you want to exclude.
9. Click **OK** to apply the exclusion list to the cluster.

Integrating AD RMS with AD FS

You can assign AD RMS to AD FS users when they connect through a federated trust. This allows an organization to share access to rights-protected content with another organization, without having to establish a separate Active Directory trust or AD RMS infrastructure.

AD FS is an Active Directory role that is a standards-based service, providing identity federation by implementing claims-based authentication between forests. Claims-based authentication is the process of user authentication based on a set of claims about the user's identity that a trusted token contains. An entity issues the token that can authenticate the user by other means, and which is trusted by the entity that performs the claims-based authentication.

You can establish identity federation between two organizations by establishing trust between two security realms. An AD FS server on one side, called the Accounts side, or AD FS-A, authenticates the user through the standard means in AD DS, and then issues a token that contains a series of claims about the user, including his or her identity. On the other side, the Resources side, or AD FS-R, an AD FS server validates the token, and then issues another token that authorizes the local servers to accept the claimed identity. This way a system can provide controlled access to its resources or services to users who belong to another security realm, without requiring users to authenticate again to the other system, and without the two systems sharing a database of user identities or passwords.

To benefit from identity federation, a service has to accept federated identities. AD RMS is one such system. It accepts requests for licenses from remote users through a single sign-on (SSO) agent or web SSO. Then it redirects them to the local federation server, which is the server in the federation's resource side, or AD FS resource domain. This server in turn asks the user to authenticate to its own domain, which is the AD FS account domain, which means the Accounts server in the user's network. This server requests authentication to AD DS and issues the corresponding security token. Then the user's security subsystem presents this token to the Web SSO, which validates the token, and then provides the identity to the AD RMS server. AD RMS then issues the licenses that the user requests.

AD FS integration for AD RMS has some limitations when compared to the other alternatives, such as trusted user domains and trusted publishing domains. Its current implementation includes one significant limitation. AD RMS with AD FS does not provide group expansion capabilities for remote groups, which means that remote users belonging to groups that are assigned rights to a document cannot exercise those rights unless the rights have also been assigned to the users individually.

Another limitation is that AD FS integration depends on the capabilities of the client accessing the protected documents. Clients using Windows Mobile versions that are older than Windows Phone® 7 are not able to authenticate through AD FS. These clients cannot consume AD RMS-protected documents unless they are in the same forest as the AD RMS server that issued the publishing license, or the organizations that use trusted user domains or trusted publishing domains to integrate the AD RMS infrastructures in each forest. Another limitation is that the Rights Management Add-on document viewer, which the receiving users use when they do not have IRM-capable applications, does not support AD FS authentication.

Using AD FS for AD RMS creates some significant infrastructure requirements, such as access to the AD RMS servers from the Internet and specific client configurations. You must configure the URLs for the remote federation servers in the trusted zone in Internet Explorer. You must also add the URLs for the local federation servers to the intranet zone.

When you integrate AD RMS with AD FS, remember to:

- Deploy and install AD FS properly in both organizations
- Add and configure AD RMS as a claims-aware application
- Grant security **audit** privileges to the AD RMS service **account**
- Add an extranet URL

Install and enable the **Identity Federation Support** role service for AD RMS

- Assign home realm to AD FS-R computers via registry changes



Additional Reading: Windows Phone clients can use certain extended features, which the article at <http://go.microsoft.com/fwlink/?LinkID=285333> describes. All of these disparate client settings can increase support costs.

Even with these limitations, AD FS can provide major benefits, especially in those environments where the partner organizations cannot implement their own AD RMS servers. AD FS offers solutions that require minimal trust between the organizations.

Integrating AD RMS with AD FS requires you to establish an AD FS infrastructure between an AD FS-A domain and an AD FS-R domain. After configuring AD FS in both forests to federate with each other, you must also configure AD RMS to work with AD FS. First, add and configure AD RMS as a claims-aware application by using the AD FS management console. Then enable the AD RMS service account to generate security audit events when it uses AD FS. To grant security audit privileges to the AD RMS service account, use the User Rights Assignment of the local security policy of the AD RMS servers. Next, specify the AD RMS extranet cluster URLs in the AD RMS management console. After specifying the cluster extranet URLs, you can add the Identify Federation Support role service to the AD RMS server.

After adding the Identity Federation Support role service, enable federated identity support for the AD RMS cluster. Once you do this, federated identity support allows user accounts to use credentials established by a federated trust relationship through AD FS to obtain a Rights Account Certificate from an AD RMS cluster. Lastly, you must configure client computers in the AD FS-R domain for federation support with AD RMS. The registry entry, HKLM\Software\Microsoft\MSDRM\Federation\FederationHomeRealm, assigns the AD FS home realm for AD RMS.



Additional Reading: For more information, go to <http://go.microsoft.com/fwlink/?LinkID=285334>

Integrating AD RMS with Windows Azure RMS

Like AD RMS, Windows Azure Active Directory Rights Management lets you safeguard sensitive information created by and when using Office applications and services such as email or correspondence that requires confidential treatment. AD RMS assigns rights to content when users create it. During distribution, an encrypted form of persistent protection follows the content wherever it travels. By using template-based assignment, AD RMS can assign rights such as the ability to allow or deny viewing, printing, copying of messages or documents as needed.

- Users can sign up as a tenant in Windows Azure AD for Windows Azure Rights Management
- When you use the RMS app to send a message to an organization you wish to employ Rights Management to, remember:
 - Messages contain simple instructions to obtain free tenant status
 - You can then start using Rights Management across a B2B partnership
- You could replace your AD RMS infrastructure with Windows Azure Rights Management

By connecting to AD DS via the Microsoft Directory Synchronization component, the Windows Azure Rights Management service benefits from a cross-organization trust. When your organization sets up a Windows Azure Active Directory (Windows Azure AD) tenant and, optionally, configures DirSync and federation services, Windows Azure Rights Management can rely on those trusts to enable business-to-business (B2B) collaboration in a more easily configured and administrated manner. Before Windows Azure AD, to ensure secure collaboration, organizations had to use paired federation trusts with every

organization with which they wanted to collaborate. Now, each organization with which you want to share protected documents must be a tenant in Windows Azure AD. Becoming a Windows Azure tenant costs \$2 per user per month. You can use the RMS app to send someone an email containing a protected file that provides simple instructions for how to sign up as a tenant (if he or she is not a Windows Azure tenant already). Then he or she can start accessing the file you shared.

How does Windows Azure Rights Management compare to AD RMS? The following table provides a side-by-side comparison of the features and benefits of Rights Management and AD RMS.

Windows Azure Rights Management	AD RMS
Supports IRM capabilities in Microsoft online services and other online offerings, such as Exchange Online and SharePoint Online, in addition to conventional on-premises products, such as Microsoft Exchange Server and SharePoint servers.	Works primarily with on-premises Microsoft server products including Microsoft SharePoint Server, Exchange Server, and File Classification Infrastructure (FCI).
Enables implicit trust between organizations and users who are current Microsoft Office 365™ subscribers. Users can share content easily and safely with other internal users or with valid users in other Office 365 tenant accounts.	Defines trusts explicitly in a direct point-to-point relationship between two organizations using either trusted user domains or federated trusts that you create by using AD FS.
Offers a predefined set of rights policy templates for use. Two templates are included; one delivering read-only viewing of protected content and the other delivering write or modify permissions over the protected content.	Provides the ability to create, define, and use your own rights policy templates. You can configure templates for AD RMS on a more granular level.
Supports users of Microsoft Office 2010 and Microsoft Office 2013 products.	Supports users of Microsoft Office 2013, Microsoft Office 2010, and Microsoft Office 2007.
Does not provide support for Windows Vista or Windows XP users. Supports users who are running Windows 7 and Windows 8.	Provides support for users running Windows 8, Windows 7, Windows Vista, and Windows XP.
Supports Cryptographic Mode 2 only.	Supports both Cryptographic Mode 1 and as Cryptographic Mode 2.
Supports outbound migration only from Windows Azure Rights Management to AD RMS.	Supports migration from Windows Azure Rights Management and migration from AD RMS on Windows Server 2003.

Windows Azure Rights Management is a new service that you could use to replace your entire AD RMS cluster infrastructure. You could do so over an extended period, allowing new content to use the Windows Azure Rights Management services while your existing AD RMS continues to support the older content.

 **For more information on Windows Azure AD Rights Management, see the following link:**

<http://go.microsoft.com/fwlink/?LinkID=386641>

Considerations for Enabling External User Access to AD RMS

When designing access to AD RMS, follow these guidelines to improve reliability and security:

- Enable both internal and external clients to access the publishing license's URL by setting the URL of the root-certification cluster to an address that clients can access over the Internet. Additionally, the address must resolve in the intranet to the root-certification cluster.
- Enable SSL and require an SSL connection between the AD RMS clients and the AD RMS server.
- Dedicate a license server to extranet users and configure the URL of the extranet cluster.
- Configure web proxy settings, if necessary.

When enabling external users to access AD RMS, follow these guidelines:

- Set the root certification cluster URL to an address that:
 - Is accessible over the Internet
 - Resolves in the intranet to AD RMS servers in the same cluster
- Enable SSL, and require an SSL connection between the AD RMS clients and the AD RMS server
- Dedicate a license server to extranet users, and configure the URL of the extranet cluster
- Configure web proxy settings, if necessary

If you control access to resources outside of the local forest by using a web-proxy server, you might need to configure AD RMS to use the proxy. This might be necessary if:

- Your organization has clients of Information Rights Management Services (RMS) 1.0 that do not have Internet connectivity. Windows RMS 1.0 is a precursor to AD RMS. Microsoft must activate clients of Windows RMS 1.0 over the Internet before the first use. The client tries to use the AD RMS cluster proxy for the activation request, but if the AD RMS cluster does not respond, the client tries to obtain the credentials directly through an Internet connection on the local computer. Windows RMS 1.0 with SP1 or SP2 does not require activation of AD RMS clients over the Internet.
- You trust RACs from Microsoft Account users. The AD RMS cluster needs to be able to validate the user against the Microsoft Account sites and services.
- Your organization has multiple forests that a web-proxy server separates, and you establish trusts between them. This is similar to the scenario with Microsoft Account, in that the registry of the user accounts that the AD RMS cluster must validate against is behind the web-proxy server.

The external AD RMS clients use the URLs of the extranet cluster to connect to the AD RMS cluster for licensing and certification. Make sure that you register the URLs in the DNS and verify that DNS is available from the Internet. If you add URLs for an extranet cluster to an existing AD RMS cluster, the current AD RMS clients must obtain new client licenser certificates. The extranet cluster URLs are part of the Extranet-License-Acquisition-URL field in the issuance license. Additionally, clients use these URLs during the service discovery process. To add URLs for an extranet cluster to an existing AD RMS cluster, you must be a member of at least the local AD RMS Enterprise Administrators group or its equivalent.

Lesson 5

Planning and Implementing AD RMS Integration with Dynamic Access Control

Dynamic Access Control is one of the new features that Windows Server 2012 introduces. Dynamic Access Control enables a new method for document authors to apply IRM protection at the storage level, instead of relying on an administrator setting security in the share and New Technology File System (NTFS) file service permissions. Integrating AD RMS and Dynamic Access Control enhances IRM protection further, to secure the transport and storage of sensitive documents.

Lesson Objectives

After completing this lesson, you will be able to:

- Integrate AD RMS and Dynamic Access Control.
- Understand how AD RMS integrates with Dynamic Access Control.
- Deploy encryption of office files.

Integrating AD RMS and Dynamic Access Control

Mitigating organizational risk is an important goal for all IT departments. These departments must try to ensure the safety of sensitive information throughout its sharing and dissemination via email and the Internet. In addition, they must meet the various compliance regulations, such as the Health Insurance Portability and Accountability (HIPPA) Act or Payment Card Industry Data Security Standard. These regulations dictate how to encrypt information. Additionally, there are numerous business reasons to encrypt sensitive business information. However,

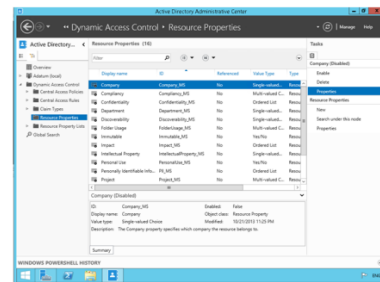
encrypting information is expensive and can affect business productivity. Accordingly, organizations tend to have different approaches and priorities for encrypting their information.

You can configure Windows Server 2012 Dynamic Access Control to encrypt sensitive Microsoft Office files automatically, based on their classification. You can do this by invoking AD RMS protection for sensitive documents a few seconds after AD RMS identifies the file as sensitive on the file server. You can also do it when you run continuous file management tasks on the file server.

AD RMS encryption provides an additional layer of protection for files. Even if a person with access to a sensitive file sends that file through email inadvertently, the AD RMS encryption persists with the file. Users who want to access the file must first authenticate themselves to an AD RMS server to receive the decryption key.

Support for non-Microsoft file formats may be available through independent software vendors (ISVs) who write code based on Microsoft attributes. After an author protects a file through AD RMS encryption, or encryption occurs automatically via Dynamic Access Control deployment, data management features, such as search functionality or content-based classification, are no longer available for that file.

- Dynamic Access Control applies encryption by using AD RMS
- Dynamic Access Control protects documents even if inadvertently saved, sent, or processed incorrectly
- Dynamic Access Control extends AD RMS to the file server



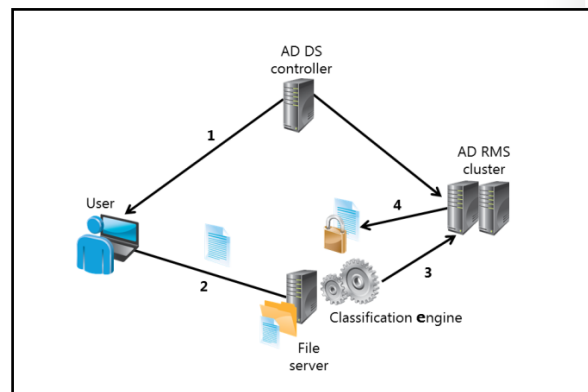
AD RMS enables both individuals and administrators, through IRM policies, to specify access permissions to documents, workbooks, and presentations. This helps prevent unauthorized people from printing, forwarding, or copying sensitive information. Once you or Dynamic Access Control restrict a file's permission with IRM, AD RMS enforces access and usage restrictions no matter where the information is, because the file's permission is stored in the document file.

File and Storage Services Dynamic Access Control lets you set up and manage file servers that provide central locations on your network, where you can store files and share them with users. To provide continuous file management, file server administrators can configure file management tasks that invoke AD RMS protection for sensitive documents a few seconds after AD RMS identifies the file as sensitive on the file server.

How AD RMS Integrates with Dynamic Access Control

Dynamic Access Control enables you to protect sensitive information automatically by using AD RMS. For example, you could apply AD RMS if a file had the word *confidential* in it, as follows:

1. You create a rule to automatically apply RMS protection to any file that contains the word *confidential*.
2. A user creates a file with the word *confidential* in the text, and then saves it.
3. The AD RMS Dynamic Access Control classification engine, following rules set in the central access policy, discovers the document with the word *confidential*, and then initiates AD RMS protection accordingly.
4. AD RMS applies an encryption template to the document on the file server, and then encrypts and classifies it.



Demonstration: Deploy Encryption of Office Files

Instead of clicking and expanding items in the console, you can perform many graphical user interface actions with Windows PowerShell cmdlets. Often, using Windows PowerShell is faster than using the graphical user interface. You can save a series of cmdlets into a script. Then you can edit the script for unique scenarios each time the script runs.

Windows PowerShell Equivalent Commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Type each cmdlet on a single line, even though they may appear word-wrapped across several lines because of formatting constraints. To enable resource properties, type the following at a command prompt, and then press Enter:

```
Set-ADResourceProperty -Enabled:$true -Identity:"CN=Impact_MS,CN=Resource
Properties,CN=Claims Configuration,CN=Services,CN=Configuration,DC=adatum,DC=com"
Set-ADResourceProperty -Enabled:$true -Identity:"CN=PII_MS,CN=Resource
Properties,CN=Claims Configuration,CN=Services,CN=Configuration,DC=adatum,DC=com"
```

To create the high-impact classification rule, type the following, then press Enter:

```
Update-FSRMClassificationPropertyDefinition
$date = Get-Date
$AutomaticClassificationScheduledTask = New-FsrmScheduledTask -Time $date -Weekly
@(3, 2, 4, 5,1,6,0) -RunDuration 0;
Set-FsrmClassification -Continuous -schedule $AutomaticClassificationScheduledTask
New-FSRMClassificationRule -Name "High Business Impact" -Property "Impact_MS" -
Description "Determines if the document has a high business impact based on the
presence of the string 'Contoso Confidential'" -PropertyValue "3000" -Namespace
@("C:\Finance Documents") -ClassificationMechanism "Content Classifier" -Parameters
@("StringEx=Min=1;Expr=Contoso Confidential") -ReevaluateProperty Overwrite
```

To protect documents with AD RMS, type the following cmdlets, and then press Enter after each one:

```
$fmjRmsEncryption = New-FSRMFmjAction -Type 'Rms' -RmsTemplate 'Adatum Finance Admin
Only'
$fmjCondition1 = New-FSRMFmjCondition -Property 'PII_MS' -Condition 'Equal' -Value
'5000'
$date = get-date
$schedule = New-FsrmScheduledTask -Time $date -Weekly @( 'Sunday' )
$fmj1=New-FSRMFileManagementJob -Name "High Impact" -Description "Automatic RMS
protection for high PII documents" -Namespace @('C:\Finance Documents') -Action
$fmjRmsEncryption -Schedule $schedule -Continuous -Condition @($fmjCondition1
```

Demonstration Steps

Enable resource properties

1. Open the Active Directory Administrative Center. Under **Resource properties** in **Dynamic Access Control**, assign the **Impact** and **Personally Identifiable Information** properties. Publish both policies in the **Global Resource list**.

Create classification rules

1. Add the **File and Storage Services File Server Resource Manager** (FSRM) service role.
2. In Windows PowerShell, update the **Global Resource Properties** to allow FSRM access.
3. Open FSRM, navigate to **Classification Management**, and then click the **Classification Rules** node. Enable the **Classification Schedule** to allow continuous classification.
4. Create a new rule as follows: Assign **High** to the **Impact** property when a document contains the **Adatum Confidential** text statement.

Protect documents with AD RMS

1. In FSRM, navigate to **File Management Tasks**, and then create a task that will provide automatic AD RMS protection to high impact documents.
2. Select the schedule to run continuously for all files.

View the results

1. On LON-SVR1, examine the properties sheet of the document in the **Finance Documents** folder. Note that there are no values in the **Classification** tab.
2. Sign in to LON-CL1 as **Adatum\Hani**, with a password of **Pa\$\$w0rd**, and then open the **Finance Memo** document. Add **Adatum Confidential** to the text, and then save it.
3. Switch back to LON-SVR1, and then note the value on the **Classification** tab of the document properties.

Lab: Planning and Implementing an AD RMS Infrastructure

Scenario

Due to the highly confidential nature of the research team's work at A. Datum Corporation, the security team wants to implement additional security for some of the research team's documents. In particular, the security team wants to ensure that users inside and outside the organization cannot share confidential documents with any unauthorized users.

You need to plan and implement an AD RMS solution that will provide the level of protection that the security team has requested.

Objectives

After completing this lab, you will be able to:

- Create an AD RMS plan.
- Implement AD RMS for internal users.
- Implement AD RMS integration with Dynamic Access Control.
- Implement AD RMS integration with a partner organization.

Lab Setup

Estimated Time: 60 Minutes

- Virtual machines:
- 20414C-LON-HOST1
- 20414C-LON-DC1
- 20414C-LON-SVR1
- 20414C-LON-CL1
- Do not start 20414C-LON-SVR1 or 20414C-LON-CL1 until 20414C-LON-DC1 is at the Sign-in prompt.
- User name: **Adatum\Administrator**
- Password: **Pa\$\$w0rd**
- 20414C-TREY-DC1
- 20414C-TREY-CL1
- Start 20414C-TREY-DC1. Do not start 20414C-TREY-CL1 until TREY-DC1 is at the Sign-in prompt.
- Sign in to the TreyResearch domain as: **TreyResearch\Administrator**
- Password: **Pa\$\$w0rd**

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

1. On LON-HOST1, start **Hyper-V Manager**.
2. In Hyper-V® Manager, click **20414C-LON-DC1**, and then in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.

4. Sign in by using the following credentials:
 - User name: **Administrator**
 - Password: **Pa\$\$w0rd**
 - Domain: **Adatum**
5. Repeat steps 2 through 4 for **20414C-LON-SVR1** and **20414C-TREY-DC1**. For TREY-DC1, sign in to the TreyResearch domain.
6. Repeat steps 2 through 3 for **20414C-LON-CL1** and **20414C-TREY-CL1**. Do not sign in until directed to do so.

Exercise 1: Planning the AD RMS Deployment

Scenario

The security department has defined the following requirements for the AD RMS deployment:

- You must use AD RMS to designate all documents in the ConfidentialResearch folder as confidential and protected.
- Only members of the Research and Managers groups should be able to access the documents.
- A. Datum employees must be able to share the AD RMS–protected content with Trey Research users.
- You must not secure any AD RMS–protected content in such a way that no one can access it.

The main tasks for this exercise are as follows:

- Read the supporting documentation.
- Update the proposal document with your planned course of action.
- Examine the suggested proposals in the Lab Answer Key.

► Task 1: Read the supporting documentation

Supporting Documentation

----- Original Message -----

From: Charlotte Weiss [Charlotte@contoso.com]
 Sent: 03 Feb 2013 08:45
 To: Ed@contoso.com
 Subject: AD RMS

Ed,

Because of the highly confidential nature of the work that the A. Datum research team performs, the security team at A. Datum Corporation wants to implement additional security for some of the research team's documents. In particular, the security team wants to ensure that users inside and outside the organization cannot share confidential documents with any unauthorized users.

You must plan and implement an AD RMS solution that will provide the level of protection requested by the security team.

Please create a plan to install AD RMS at both A. Datum and Trey Research that incorporates the rest of the action items.

Thank you,
 Charlotte

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

Answer the questions in the A. Datum IRM Plan: AD RMS document shown below.

A. Datum IRM Plan: AD RMS	
Document Reference Number: GW00612	
Document Author	Charlotte Weiss
Date	6 th February
<ul style="list-style-type: none"> • Requirements Overview • Design an AD RMS deployment for the A. Datum and the Trey Research companies. • The goal of the AD RMS deployment is to protect information, no matter where it goes. Once you add AD RMS protection to a digital file, the protection stays with the file. By default, only the content owner is able to remove the protection from the file. The owner grants rights to other users to perform actions on the content, such as the ability to view, copy, or print the file: • AD RMS must automatically designate as confidential and then protect all documents in the ConfidentialResearch folder. • Only members from the Research and Managers groups should be able to access the documents. • Configure the integration of this AD RMS deployment with Dynamic Access Control. • A. Datum employees must be able to share the AD RMS–protected content with users at Trey Research. • The A. Datum security team wants to ensure that you do not secure any AD RMS–protected content in such a way so that no one can access the information. 	
<p>Additional Information</p> <p>We need to cut down on the total number of servers required. Consider using the internal database instead of a full SQL Server deployment. We have authorization to deploy AD RMS on the existing domain controllers.</p>	

Proposal Questions:

Question: How many AD RMS clusters do you need to deploy to satisfy the security requirements for both companies?

Question: What database solution will you deploy?

Question: What service accounts, if any, do you need to create?

Question: What Secure Sockets Layer (SSL) certificate requirements do we have? How do we satisfy them in both forests?

► **Task 3: Examine the suggested proposals in the Lab Answer Key**

Compare your solution to the lab answer key

1. How many AD RMS clusters do you need to deploy to satisfy the security requirements for both companies?
2. What database solution will you deploy?
3. What service accounts, if any, do you need to create?
4. What SSL certificate requirements do you have? How do you satisfy those requirements in both forests?

Results: At the end of this lab exercise, you will have planned the deployment of an Active Directory® Rights Management Services (AD RMS) infrastructure for the Adatum.com and TreyResearch.Net forests, based on business requirements and management specifications.

Exercise 2: Deploying an AD RMS Infrastructure for Internal Users

Scenario

The first step in deploying AD RMS at A. Datum is to deploy AD RMS for internal users. You will begin by configuring the appropriate DNS records and the AD RMS service account, and then you will continue with installing and configuring the first AD RMS server. You will deploy the second AD RMS server in the Toronto office and enable the AD RMS Super Users group. After deploying the AD RMS servers, the next step is to configure the rights policy templates and exclusion policies for the organization.

The main tasks for this exercise are as follows:

1. Configure the AD RMS prerequisites.
2. Deploy the first server in an AD RMS cluster.
3. Deploy the AD RMS cluster in the Trey Research forest.
4. Configure the AD RMS Templates.
5. Configure AD RMS Exclusion Policies.
6. Validate the internal deployment.

► Task 1: Configure the AD RMS prerequisites

Create and configure accounts

1. On LON-DC1, from Server Manager, open Active Directory Users and Computers, and then create the following user account in the **Users** container:
 - Name: **ADRMSSRVC**
 - Password: **Pa\$\$w0rd**
 - Clear the **User must change password at next logon** check box.
 - Group Membership: **Domain Admins**
2. In the general tab of the users' properties sheets, create email addresses for **Hani Loza**, **Limor Henig**, **Stuart Glasson**, and **Toni Poe** by using the format **user@adatum.com**.
3. In the **Users** container, create the following universal security groups: **Employees**.
4. Create an email address for the following groups: **Employees**, **Managers**, and **Research**. Use the format **group@adatum.com**.
5. Assign the following users to groups:
 - Limor Henig-Managers, Employees
 - Toni Poe-Managers
 - Stuart Glasson-Employees

Create shared folders on LON-DC1

1. On LON-DC1, open File Explorer, and then create and share **C:\ConfidentialResearch**.
2. Set share permissions for the **Research** and **Managers** groups as **Read/Write**.
3. Create and share **C:\Public**.
4. Set share permissions for the **Everyone** group as **Read**.

► **Task 2: Deploy the first server in an AD RMS cluster**

Deploy the AD RMS cluster in Adatum.com

1. In the Add Roles and Features Wizard on LON-DC1, add the **Active Directory Rights Management Services** role.
2. Accept the required features. When finished, click **Close**.

To configure a new AD RMS root cluster

1. In Server Manager, perform the Additional Configuration as stated under **Notifications**. Use the following values:
 - AD RMS Cluster: **Create a new AD RMS root cluster**
 - Configuration Database: **Windows Internal Database**
 - Service account:
 - Name
 - Password: **Pa\$\$w0rd**
 - Cryptographic mode: **Cryptographic Mode 2**
 - Cluster Key Storage: **centrally managed**
 - Cluster Key Password: **Pa\$\$w0rd**
 - Cluster Web Site: Default web site
 - Cluster address:
 - <http://LON-DC1.Adatum.com>
 - Licensor certificate: **LON-DC1**
 - SCP Registration: Register the SCP now.
2. Sign out of LON-DC1, and then sign in as **ADATUM\Administrator** with the password **Pa\$\$w0rd**.

Open the Active Directory Rights Management Services console

1. On LON-DC1, from Server Manager, open the Active Directory Rights Management Services console.
2. Verify that the console opens with no errors.
3. Close the Active Directory Rights Management Services console.

► **Task 3: Deploy the AD RMS cluster in the Trey Research forest**

Add the AD RMS server role

1. On TREY-DC1, use the Add Roles and Features Wizard to add the **Active Directory Rights Management Services** role.
2. Accept the required features. When finished, click **Close**.

3. In Server Manager, perform the Additional Configuration as stated under **Notifications**. Use the following values:
 - AD RMS Cluster: **Create a new AD RMS root cluster**
 - Configuration Database: **Windows Internal Database**
 - Service account:
 - Name: **TreyResearch\ADMSSVC**
 - Password: **Pa\$\$w0rd**
 - Cryptographic mode: **Cryptographic Mode 2**.
 - Cluster Key Storage: **Centrally managed**
 - Cluster Key Password: **Pa\$\$w0rd**
 - Cluster Web Site: **Default web site**
 - Cluster address:
 - <http://TREY-DC1.TreyResearch.net>
 - Licensor certificate: **TREY-DC1**
 - SCP Registration: Register the SCP now.
4. Sign out of TREY-DC1, and then sign in as **TreyResearch\Administrator** with the password **Pa\$\$w0rd**.

Open the Active Directory Rights Management Services console

1. On TREY-DC1, verify that the Active Directory Rights Management Services console opens without errors.
2. Close the Active Directory Rights Management Services console.

► Task 4: Configure the AD RMS Templates

Configure the AD RMS templates

1. Switch to LON-DC1, and then open the Active Directory Rights Management Services console.
2. Navigate to the **Rights Policy Templates** node, and then configure the **Templates** file location to be \\LON-DC1\public.
3. Create a rights policy template named **Adatum.com RC** that gives the **Managers** group **View** access and the **Research** group **Full Control**.
4. Click **Finish**, and then close the Active Directory Rights Management Services console.
5. Sign in to LON-CL1 as **adatum\administrator**, with the password **Pa\$\$w0rd**.
6. Open Scheduled Tasks, and then enable **AD RMS Rights Policy Template Management (Automated)** in the local **Task Scheduler**.
7. Open **regedit.exe**, and then add the **DRM** key and the **AdminTemplatePath** subkey with the value set to the local user's template folder: **%LocalAppData%\Microsoft\DRM\Templates**.
8. Use the registry key: **HKEY_CURRENT_USER\Software\Microsoft\Office\14.0\Common\DRM**. Close **regedit.exe**.
9. Add the **Domain Users** group to the **Local Remote Users** group.
10. Sign out of LON-CL1.

► Task 5: Configure AD RMS Exclusion Policies

A. Datum has a contract with a management consulting group. A. Datum has given Toni Poe, a management consultant from that group, office space and an A. Datum computer system to perform office work. She has an AD DS account, and A. Datum has placed her in the Management universal group. However, she should not have Read permissions on restricted Research Confidential files to which the Management group has view access. To that end, we will create an exclusion policy on the AD RMS cluster for her.

1. Sign in to LON-CL1 as **adatum\toni** with the password **Pa\$\$w0rd**.
2. Open Microsoft Word, and then create a new document.
3. Assign protection to the new document for Toni only, and save it as **Text.docx** in the Documents library. At the authentication request, provide **Adatum\Toni** with the password **Pa\$\$w0rd**.
4. Sign out of LON-CL1.
5. Switch back to LON-DC1. Open the Active Directory Rights Management Services console.
6. Navigate to the **Exclusion Policies- Users** node.
7. Add the **RAC** for **toni@adatum.com** to the **Exclusion policy**.
8. Close the Active Directory Rights Management Services console.

► Task 6: Validate the internal deployment

1. Sign in to LON-CL1 as **Adatum\Hani** with the password **Pa\$\$w0rd**.
2. Open Internet Explorer, modify the Local intranet zone, and then add **http://LON-DC1.adatum.com**.
3. Sign out of LON-CL1.
4. Repeat steps 1 to 3 for **Adatum\Limor**.
5. Sign in to LON-CL1 as **Adatum\Hani**.
6. Open Microsoft Word, and then type **Management employees can read this document, but they cannot change, print, or copy it. Research group members have Full control.**
7. Protect the document by using the following settings: **Management** group has **Read** access and **Research** group has **Change** access.
8. Save the document as **\\LON-DC1\ConfidentialResearch\ADRMS-TST.docx**.
9. Sign out of LON-CL1.
10. Sign in to LON-CL1 as **Adatum\Limor**.
11. Open File Explorer, and then browse to the **\\LON-DC1\ConfidentialResearch** share. Open **ADRMS-TST**, and then test permissions.
12. Sign out of LON-CL1.

Results: You should have a working AD RMS cluster in both the Adatum.com and TreyResearch.net forests. In addition, you should be able to protect Microsoft Office documents with IRM, and see the results of that protection.

Exercise 3: Implementing AD RMS Integration with Dynamic Access Control

Scenario

To ensure that AD RMS protects all documents in the ConfidentialResearch folder automatically, you need to configure the integration of AD RMS and Dynamic Access Control. In addition, you must configure the Dynamic Access Control claims, file classifications, and central access policies.

The main tasks for this exercise are as follows:

1. Enable resource properties.
2. Create classification rules.
3. Automatically protect documents with AD RMS.
4. Verify the deployment.

► Task 1: Enable resource properties

To enable resource properties

1. On LON-DC1, from Server Manager, open the Active Directory Administrative Center, and then switch to **Tree View**.
2. In **Dynamic Access Control**, click **Resource properties**. Enable the **Impact** and **Personally Identifiable Information** properties, and then publish both properties to the **Global Resource Property List**.
3. If the policies are in the **Global Resource** list already, click **Cancel**.

► Task 2: Create classification rules

This task explains how to create the **High Impact** classification rule. This rule will search the content of documents, and it finds the string "Adatum Confidential", it will classify this document as having high business impact. This classification will override any previously assigned classification of low business impact.

You will also create a **High PII** rule. This rule searches the content of documents, and if it finds a Social Security number, it classifies the document as having high personally identifiable information (PII).

1. Sign in to LON-SVR1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. In Server Manager, start the Add Roles and Features Wizard, and then add the **File and Storage Services File Server Resource Manager (FSRM)** service role.
3. Open Windows PowerShell, and then update the **Global Resource Properties** to allow FSRM access. Type the following command at the command prompt, and then press Enter:

```
Update-FSRMClassificationPropertyDefinition
```

4. Open File Server Resource Manager, and then navigate to the **Classification Management, Classification Rules** node. Enable the **Classification Schedule** to allow continuous classification for new files.

5. Create two rules, as follows:
 - **High Business Impact:** When **Adatum Confidential** text is included in a document, classify the **Impact** property as **High**. This should manage the C:\Research Documents folder. Also, overwrite any existing property values.
 - **High PII:** If a Social Security Account Number is added, classify the **Personally Identifiable Information** property as **High**. This should manage the C:\Research Documents folder. Also, overwrite any existing property values. Enter the following, ensuring that you do include any line breaks in the expression:

```
^(?!000)([0-7]\d{2}|7([0-7]\d|7[012]))([ -]?)(!00)\d\d\3(!0000)\d{4}$
```

► Task 3: Automatically protect documents with AD RMS

Now that you have created rules to classify documents automatically based on content, you must create a file management task that uses AD RMS to protect certain documents automatically based on their classification. In this step, you will create a file management task that protects any documents with a high PII automatically. Only members of the Research group will have access to documents that contain high PII.

1. On LON-SVR1, open Internet Explorer, modify the Local intranet zone, and then add <https://LON-DC1.adatum.com>.
2. In File Server Resource Manager, navigate to **File Management Tasks**, and then create a task named **High PII** that will provide automatic AD RMS protection for the **C:\Research Documents** folder when the **Personally Identifiable Information** property is set to **High**.
3. Configure the schedule to run continuously for all files.

► Task 4: Verify the deployment

1. On LON-SVR1, open the File Explorer, and then navigate to **C:\Research Documents**.
2. Examine the properties of the two documents in the **Research Documents** folder. Note that they have no values configured on the **Classification** tab.
3. Switch to LON-CL1, and sign in as **Adatum\Hani** with the password **Pa\$\$w0rd**.
4. From the desktop, type **\\LON-SVR1\Research Documents**, and then press Enter.
5. Open the **Finance Memo** document. Type **Adatum Confidential**. Save the document, and then close Microsoft Word.
6. Open the **Request for Approval to Hire** document. Type **Social Security #:**, press the Enter key, and then on a new line, type **777-77-7777**. This must be on a new line for Dynamic Access Control to notice the expression quickly. Save the document, and then close Microsoft Word.
7. Switch to LON-SVR1, and then browse to **C:\Research Documents**. Note the values in the **Classification** tab of both documents' properties.
8. Switch back to LON-CL1. Open the **Request for Approval to Hire** document. You should see the **Connect to AD RMS server** window appear on the screen. Close the document, and then sign out.
9. On all machines, close all open windows, and then sign out.

Results: You should have applied Dynamic Access Control classification rules to IRM-protected content.

Exercise 4: Implementing AD RMS Integration for External Users

Scenario

As part of the AD RMS deployment, you must ensure that AD RMS functionality extends 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 domain policy.
2. Export the trusted publishing domain 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.
6. Verify user access to the protected document.
7. To prepare for the next module.

► Task 1: Export the trusted user domain policy

Export the trusted user domains policy

1. On LON-DC1, open the Active Directory Rights Management Services console, navigate to **Trusted Policies – Trusted User Domains**, and then export the policy to a file named **C:\ADRMS_LON-DC1_LicenserCert.bin**.
2. Repeat step 1 above on TREY-DC1, saving the file as **C:\ADRMS_TREY-DC1_LicenserCert.bin**.

► Task 2: Export the trusted publishing domain policy

Export the trusted publishing domain policy

1. From LON-DC1, open the Active Directory Rights Management Services console, and then navigate to **Trusted Policies – Trusted Publishing Domain**.
2. Select the **LON-DC1** certificate, and then export it, saving the file as **C:\AdatumTrustedPubDomain.xml** with the password **Pa\$\$w0rd**.
3. Repeat steps 1 and 2 above on TREY-DC1 by selecting the **TREY-DC1** certificate and using the file name **C:\TreyTrustedPublDomain.xml**, with the same password.

► Task 3: Import the trusted user domain policy from the partner domain

1. On LON-DC1, in Server Manager, click **Tools**, and in the drop-down list box, click **DNS**. Expand **LON-DC1**, select and right-click **Conditional Forwarders**, and then click **New Conditional Forwarder**.
2. Under **DNS Domain**, type **TreyResearch.net**. In the **IP addresses of the master servers** box, type **172.16.10.10**. Press Enter, and then click **OK**. Close DNS Manager.
3. On TREY-DC1, repeat steps 1 to 3. Use the DNS domain **Adatum.com** and set the **IP Address of the master servers** as **172.16.0.10**.
4. On LON-DC1, open the Active Directory Rights Management Services console, navigate to **Trust Policies – Trusted User Domains**, and then import the file **\\TREY-DC1\C\$\ADRMS_TREY-DC1_LicenserCert.bin**. Name the file **TreyResearch**.
5. Repeat step 4 above on TREY-DC1 by using the file **\\LON-DC1\C\$\ADRMS_LON-DC1_LicenserCert.bin**. Name the file **Adatum**.

► **Task 4: Import the trusted publishing domains policy from the partner domain**

Add a trusted publishing domain

1. On LON-DC1, open the Active Directory Rights Management Services console, navigate to **Trust Policies – Trusted Publishing Domains**, and then import `\\TREYDC1\c$\TreyTrustedPubDomain.xml`. Use the password **Pa\$\$w0rd** and set the **Display name** as **TreyResearch Domain**.
2. Repeat step 1 above on TREY-DC1 by using the file `\\LON-DC1\C$\AdatumTrustedPubDomain.xml`. Use the password **Pa\$\$w0rd** and a **Display name** of **Adatum Domain**.

► **Task 5: Configure anonymous access to the AD RMS licensing server**

- On LON-DC1, open Internet Information Services (IIS), and then enable **Anonymous Authentication** and disable **Windows Authentication** on the following two files under **Default Web Site\wmcs\Licensing**:
 - **license.asmx**
 - **ServiceLocator.asmx**

► **Task 6: Verify user access to the protected document**

1. Sign in to LON-CL1 as **Adatum\Hani**, in the `\\LON-DC1\ConfidentialResearch` folder, open the **ADRMS-TST.docx** document, add **liberty@treyresearch.net** with the **Read** permission set to **Restricted Access**, and then save the document. Close all open windows and sign out.
2. At the Microsoft Word prompt, click **Change User**, and then on the **Select User** prompt, click **OK**.
On LON-DC1 as **Adatum\Administrator**, copy the **ADRMS-TST.docx** from `C:\ConfidentialResearch` to the `\\TREY-DC1\Public` share.
3. Sign in to TREY-CL1 as **TreyResearch\Liberty**, and add `http://TREY-DC1.TretResearch.net` as a **Local Intranet** site in **Internet Explorer**. Navigate to the `\\TREY-DC1\Public` share, and then open the **ADRMS-TST.docx** file. Observe that the document is rights-protected and cannot be saved or printed. Sign out of all virtual machines.

► **Task 7: 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 **20414C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20414C-LON-SVR1**, **20414C-LON-CL1**, **20414C-TREY-DC1**, and **20414C-TREY-CL1**.

Results: You should have both a working trusted user and trusted publishing domain policy between the Adatum.com and TreyResearch.net forests. In addition, you should be able to protect Microsoft Office documents with IRM for external users across the domains.

Question: How many AD RMS clusters do you need to deploy to satisfy the security requirements for both companies?

Question: What database solution should you deploy?

Question: What service accounts, if any, do you need to create?

Question: What SSL certificate requirements do you have? How do you satisfy these requirements in both forests?

Module Review and Takeaways

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
Be aware that you can use self-signed certificates, but you will need to add the server certificates from the AD RMS cluster servers to the Trusted Certification Authorities Store of any computer involved in IRM protection.	
When you installed the AD RMS cluster in the lab, you used the Windows Internal database. However, if you intend to have multiple AD RMS servers in the cluster, you will need to use a full version of SQL Server to be able to add database items from multiple servers.	

Real-world Issues and Scenarios

If you enable the Super Users Group, you also enable success and failure auditing for Audit account management and Audit directory services access.

Tools

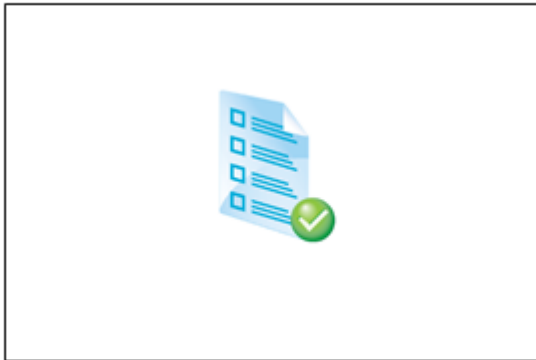
Tool	Where is it?
Active Directory Rights Management Service Administration Console	Server Manager, Tools
Active Directory Administrative Center	Server Manager, Tools
Windows PowerShell	Start Screen, Desktop Task Bar
File Services Resource Manager	Server Manager, Tools
Regedit.exe	Type in Run text box or command prompt

Course Evaluation

Your evaluation of this course will help Microsoft understand the quality of your learning experience.

Please work with your training provider to access the course evaluation form.

Microsoft will keep your answers to this survey private and confidential and will use your responses to improve your future learning experience. Your open and honest feedback is valuable and appreciated.



Module 1: Overview of Management in an Enterprise Data Center

Lab: Considerations for Implementing an Enterprise Data Center

Exercise 1: Planning the Secure Implementation of Services Within an Enterprise Data Center

► Task 1: Review the IT and business requirements

- Read the lab scenario

► Task 2: Identify how to address the IT and business requirements

1. How will you meet the requirement to provide a single tool for managing virtualization hosts and virtual machines?

Answer: Microsoft® System Center 2012 R2 Virtual Machine Manager (VMM) provides a single management interface for managing virtualization hosts and virtual machines.

2. How will you meet the requirement for business unit administrators to manage their applications and virtual machines?

Answer: There are several ways to provide this functionality. In VMM, you can delegate access to parts of the infrastructure. You can find these permissions in System Center 2012 R2 App Controller, when you use it as the self-service site for VMM. App Controller, System Center 2012 R2 Service Manager, and System Center 2012 R2 Orchestrator all provide options for delegating different levels of permissions.

3. How will you meet the requirement to provide detailed information about the performance of all data center components?

Answer: System Center 2012 R2 Operations Manager provides this functionality.

4. How will you meet the requirement to automate processes within the data center?

Answer: The System Center tools provide a variety of options for automation. You can use Service Manager to automate responses to issues. You can use Orchestrator runbooks to automate a complex sequence of tasks.

5. How will you meet the requirement to provide high availability for the required applications?

Answer: Your approach will vary depending on the application or service that you are making highly available. Microsoft Exchange Server provides a built-in high availability solution. A Microsoft SQL Server® website provides high availability options, if you use it as a back end for front-end websites. Probably, you will use some sort of load balancing to make websites highly available. You can make the applications that run on a Windows® server highly available by virtualizing the server. Do this by using a P2V conversion, and then using failover clustering to make the virtual machine highly available.

6. How will you meet the security department requirements?

Answer: All of the requirements will require certificates, so you will need to deploy an internal public key infrastructure (PKI), such as Active Directory® Certificate Services (AD CS).

7. How will you meet the requirement to improve the performance of the sales website?

Answer: You can use the System Center tools to automate the performance management for the website. You could use Operations Manager to monitor the performance on the website. When performance reaches a specified threshold, you could use Orchestrator to trigger a runbook that deploys an additional

web server or assigns more resources to the computer running SQL Server. As the number of requests decrease, the System Center components can return the environment to the previous state automatically.

8. How will you meet the requirement to enable integration with Trey Research?

Answer: By deploying Active Directory Rights Management Services (AD RMS), you can enable protection of documents after they leave your organization.

9. How will you meet the requirement to provide the required access to the partner website?

Answer: You can use Active Directory Federation Services (AD FS) to provide different levels of access on a site to user accounts from remote domains.

► **Task 3: Discuss your proposed solution with the class, as guided by your instructor**

- Be prepared to discuss your answers with the class

Results: After completing this exercise, you will have identified the components that you will need to include in the data center design.

Module 2: Planning and Implementing a Server Virtualization Strategy

Lab: Planning and Implementing a Server Virtualization Strategy

Exercise 1: Planning the Hyper-V Host Deployment

► Task 1: Read the supporting documentation

- Read the documentation that the student handbook provides.

► Task 2: Update the proposal document

- Answer the questions in the proposals section of the A. Datum Server Virtualization Strategy document.

Question: Is it possible to have a single Virtual Machine Manager (VMM) server to accommodate the requirements or should you implement three VMM management servers?

Answer: Yes, you can distribute the VMM components across the branches, and then configure the appropriate host groups.

Question: How many library servers should you deploy?

Answer: At least one per site. However, in London, there is mention of both a primary and a secondary data center. Therefore, it would be good to deploy four library servers to provide greater availability for the head office clusters.

Question: Can you prevent virtual machines from powering down during the day?

Answer: Yes, you can configure a schedule for power optimization.

Question: How can you keep the number of templates for server deployment low?

Answer: You can establish equivalent objects that ensure the same International Organization for Standardization (ISO) images and virtual hard disks (VHDs) are in each library.

Question: What measures will you take to avoid bandwidth utilization?

Answer: You can set up equivalent objects, and then on each site, deploy Windows Server Update Services (WSUS), deploy Windows Deployment Services (Windows DS), and establish a library.

Question: Which hosts servers may require different host reserve settings? Where do you choose these settings?

Answer: The branch office and the maintenance host servers may require different reserves. You can set them at the host group level or individually per host.

Question: What is the benefit of a maintenance host? It seems like a waste of resources. Can it do anything else?

Answer: Deploying a maintenance host is a good practice, as you can deploy virtual machines to it without affecting production systems. Additionally, you can conduct performance and load testing of workloads. Furthermore, maintenance hosts can act as your branch office library servers.

Question: How can you get reports on power savings?

Answer: Reports on power settings are available by integrating VMM with Microsoft® System Center 2012 R2 Operations Manager.

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

- Compare your proposals with the ones shown above.

Exercise 2: Configuring Hyper-V Host Groups

► Task 1: Add Hyper-V host to VMM

Add Hyper-V hosts

1. On LON-VMM1, on the desktop, double-click **Virtual Machine Manager Console**.
2. On the **Connect to Server** dialog box, ensure that the **Use current Microsoft Windows session identity** check box is selected, and then click **Connect**. The Virtual Machine Manager console opens.
3. Click the **Fabric** workspace.
4. From the ribbon, click **Add Resources**, and then select **Hyper-V Hosts and Clusters**. The Add Resource Wizard starts.
5. On the **Resource Location** page, click **Windows Server computers in a trusted Active Directory domain**, and then click **Next**.
6. On the **Credentials** page, click **Manually enter the credentials**, and then enter the following credentials:
 - User name: **Adatum\Administrator**
 - Password: **Pa\$\$wOrd**
7. Click **Next**.
8. On the **Discovery scope** page, click **Specify Windows Server computers by names**, in the **Computer names** field, type **LON-HOST1**, and then click **Next**.
9. On the **Target resources** page, click **lon-host1.adatum.com**, and then click **Next**. On the warning about the Hyper-V role, click **OK**.
10. On the **Host Settings** page, click **All Hosts**, and then click **Next**.
11. On the **Summary** page, click **Finish**.
12. Review the status, and then close the Jobs window when the task finishes. A warning message may appear, related to multipath I/O not being enabled. You can ignore this message for this task.
13. Click the **Fabric** workspace.
14. From the ribbon, click **Add Resources**, and then select **Hyper-V Hosts and Clusters**. The Add Resource Wizard starts.
15. On the **Resource location** page, click **Windows Server computers in a trusted Active Directory domain**, and then click **Next**.
16. On the **Credentials** page, click **Manually enter the credentials**, and then enter the following credentials:
 - User name: **Adatum\Administrator**
 - Password: **Pa\$\$wOrd**
17. Click **Next**.
18. On the **Discovery Scope** page, click **Specify Windows Server computers by names**, in the **Computer names** field, type **LON-HOST2**, and then click **Next**.
19. On the **Target resources** page, click **lon-host2.adatum.com**, and then click **Next**. On the warning about the Hyper-V role, click **OK**.
20. On the **Host settings** page, click **All Hosts**, and then click **Next**.
21. On the **Summary** page, click **Finish**.
22. Review the status, and then close the Jobs window when the task finishes. A warning message may appear, related to multipath I/O not being enabled. You can ignore this message for this task.

► **Task 2: Create the host groups**

Create the host groups

1. On LON-VMM1, from the Virtual Machine Manager console, click **VMs and Services**.
2. In the navigation pane, click the **All Hosts** node.
3. On the ribbon, click **Create Host Group**, and then replace the default name with **London Hosts**.
4. Repeat steps 2 and 3 to create the following:
 - **Toronto Hosts**
 - **Sydney Hosts**

► **Task 3: Configure the host groups**

Configure host groups

1. On LON-VMM1, from the Virtual Machine Manager console, click **VMs and Services**.
2. From the VMs and Services workspace, right-click **London Hosts**, and then click **Properties**.
3. Click **Host Reserves**, clear the **Use the host reserves settings from the parent host group** check box, click in the **Memory** field, and then change the value to **1024 MB**.
4. Click **Dynamic Optimization**, and then clear the **Use dynamic optimization settings from the parent host group** check box.
5. Select the **Automatically migrate virtual machines to balance load at this frequency** check box, and then change the value to **60 Minutes**.
6. Select the **Enable power optimization** check box, click **Settings**, review the settings, and then click **OK**.
7. Click **OK** to apply the changes.
8. From the **Fabric** workspace, expand **Servers**, and then click **All Hosts**.
9. Right-click lon-host1.adatum.com and click Move to Host Group. From the Parent host group selection, choose London Hosts. Click **OK**.
10. Right-click lon-host2.adatum.com and click Move to Host Group. From the Parent host group selection, choose London Hosts. Click **OK**.
11. From the **VMs and Services** workspace, right-click **Toronto Hosts**, and then click **Properties**.
12. Click **Dynamic Optimization**, and then clear the **Use dynamic optimization settings from the parent host group** check box.
13. Select the **Automatically migrate virtual machines to balance load at this frequency** check box, and then change the value to **60 Minutes**.
14. Select the **Enable power optimization** check box, and then click **Settings**.
15. Configure 7 a.m. to 7 p.m. Monday to Friday for no power optimization. Click **OK**.
16. Click **OK** to close the **Toronto Hosts Properties** box.

Results: After completing this exercise, you will have added Microsoft Hyper-V® hosts to VMM and created and configured host groups.

Exercise 3: Configuring VMM Libraries

► Task 1: Configure a library server and share

Configure a library server and share

1. On TOR-SVR1, on the Server Manager Dashboard, click **File and Storage Services**, and then click **Shares**.
2. In the Shares workspace, click **Tasks**, and then click **New Share**.
3. On the **Select Profile** page, click **SMB Share – Quick**, click **Next**, and on the **Share Location** page, click **Next**. On the **Share Name** page, in the **Share name** field, type **TORVMMLibrary**, and then click **Next**.
4. On the **Other Settings** page, click **Next**, and on the **Permissions** page, click **Next**. On the **Confirmation** page, click **Create**, and then click **Close**.
5. On LON-VMM1, in the Virtual Machine Manager console, click the **Library** workspace.
6. From the ribbon, click **Add Library Server**.
7. On the **Enter Credentials** page, click

Enter a user name and password, enter the following credentials, and then click **Next**:

- User name: **Adatum\Administrator**
 - Password: **Pa\$\$w0rd**
8. On the **Select Library Servers** page, click **Search**.
 9. On the **Computer Search** page, in the **Computer name** field, type **TOR-SVR1**, and then click **Search**.
 10. In the Search results area, click **tor-svr1.adatum.com**, click **Add**, click **OK**, and then click **Next**.
 11. On the **Add Library Shares** page, select the **TORVMMLibrary** check box, and then select the **Add Default Resources** check box. Click **Next**, and then on the **Summary** page, click **Add Library Servers**.
 12. Review the job status, and then close the Jobs window when the job is complete.

► Task 2: Copy the required files to the branch office's data center

Manually copy files to library shares

1. On LON-VMM1, in the Virtual Machine Manager console, click the **Library** workspace.
2. Expand **Library Servers**, and then click **LON-VMM1.Adatum.com**.
3. In the details pane, right-click **Blank Disk – Small.vhdx**, and then click **Open File Location**.
4. In the VHDs window, right-click **Blank Disk – Small.vhdx**, and then click **Copy**. Close the VHDs window, in the navigation pane, expand **Library Servers**, and then expand **tor-svr1.adatum.com**.
5. Right-click **TORVMMLibrary**, and then click **Explore**. In the TORVMMLibrary window, right-click an empty space, and then click **Paste**.
6. Close the TORVMMLibrary window. Right-click **TOR-SVR1.adatum.com**, and then click **Refresh**. Confirm that the **Blank Disk - small.vhdx** file appears in the objects list.

► **Task 3: Assign a library to a host group**

Assign a library to a host group

1. On LON-VMM1, in the Virtual Machine Manager console, click the **Library** workspace.
2. Expand **Library Servers**, right-click **tor-svr1.adatum.com**, and then click **Properties**.
3. Click the **Host group** drop-down list box, and then select the **Toronto Hosts** host group. Click **OK**.
4. Close the Virtual Machine Manager console.

► **Task 4: To prepare for the next module**

Do not revert the virtual machines, as you will need them during the next module.

Results: After completing this lab, you will have configured a VMM library server, copied files to the branch office's data center, and assigned the library to a host group.

Module 3: Planning and Implementing Networks and Storage for Virtualization

Lab: Planning and Implementing Virtualization Networks and Storage

Exercise 1: Planning a Storage Infrastructure for Virtualization

► Task 1: Read the supporting documentation

- Read the documentation that the student handbook provides.

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

- Answer the questions in the proposals section of the A. Datum Corporation Toronto storage strategy document.

Question: The servers provided have a single 10 GB onboard network adapter with all the latest Internet Small Computer System Interface (iSCSI) features. For redundancy, you want to purchase more. There are three PCI Express slots available. How many should you buy?

Answer: A single 10-gigabyte (GB) Ethernet card should be sufficient for the number of virtual machines. One more network adapter will provide high availability for your storage path.

Question: The Network team advises you that there will be a significant cost and delay should you wish to implement a separate iSCSI network. The current network is relatively new and has some 10 GB capability, though its number of connections is limited. You must advise the team on the number of 10 GB connections required. How many are required?

Answer: Each Hyper-V® host requires two 10 GB Ethernet cards, and each iSCSI target server requires two 10 GB Ethernet cards. Eight 10 GB Ethernet ports are required to provide fully redundant storage paths.

Question: You have considered using another host server as a maintenance host, but your budget does not permit this. What is an alternative solution?

Answer: Given the high cost of 10 GB Ethernet cards, you could choose to use multiple 1 GB cards. By doing so, you may be able to afford another smaller host server that can act as a deployment and maintenance server. You can determine the requirements by monitoring existing disk and network utilization.

Question: You identified data protection as both a risk and potential bottleneck. What do you think is the cause, and how can you resolve the issue?

Answer: The scenario does not mention backups. However, given the approximately 20 terabytes of data you have to back up, you need to ensure that you can continue to adhere to the company's existing backup policy. Additionally, you must ensure that backups will complete without affecting production performance. You can include an upgrade to your backup servers. However, before buying 10 GB Ethernet cards, confirm that the number of source disks in your storage area network (SAN) and the number of target disks in your backup servers can provide sufficient I/O to warrant 10 GB.

Question: What Windows Server® feature will help keep your virtual machines running when a network storage component (such as a network adapter or network switch) fails?

Answer: Multipath I/O (MPIO) is the feature that you must enable to allow Windows® to communicate with storage over multiple paths.

- ▶ **Task 3: Examine the suggested proposals in the Lab Answer Key**
 - Compare your proposals with those shown above.
- ▶ **Task 4: Discuss your proposed solution with the class, as guided by your instructor**
 - Be prepared to discuss your proposals with the class.

Exercise 2: Planning a Network Infrastructure for Virtualization

- ▶ **Task 1: Read the supporting documentation**
 - Read the documentation that the student handbook provides.
- ▶ **Task 2: Update the proposal document with your planned course of action**
 - Answer the questions in the proposals section of the A. Datum Toronto Network Virtualization Strategy document.

Question: Which of the listed networks might you need to virtualize to support the objectives?

Answer: You might need to virtualize the infrastructure servers, application servers, and possibly the backup network.

Question: Where and how can you configure virtualized networks?

Answer: You can use Virtual Machine Manager (VMM) or Windows PowerShell® to set up and configure network virtualization.

Question: What are the steps to configure the virtualized networks?

Answer: This depends on the requirements, but in summary:

- Identify the virtual machines and networks.
- Configure the virtual machine networks and gateways.
- Assign the virtual machine IP addresses.

Question: Are there other ways to host these overlapping virtual machines?

Answer: You may be able to use virtual local area networks (VLANs). However, you need to determine whether you are going to move the virtual machines between hosts, or if they must communicate with other virtual machines that are in the isolated network on different hosts. If the latter, then you must configure routing and VLANs on all network infrastructure that passes their traffic.

Question: How can you ensure that clients in the A. Datum and Wingtip Toys branch offices are able to access the correct servers?

Answer: Deploy Windows Server Gateway to allow communication between the clients on the A. Datum and Wingtip Toys branch office networks and the virtual machines on virtualized networks.

- ▶ **Task 3: Examine the suggested proposals in the Lab Answer Key**
 - Compare your proposals with those shown above.
- ▶ **Task 4: Discuss your proposed solution with the class, as guided by your instructor**
 - Be prepared to discuss your proposals with the class.

Exercise 3: Implementing a Storage Infrastructure for Virtualization

► Task 1: Configure iSCSI targets for the virtual machine deployment

Add virtual disks to LON-SVR1

1. On LON-HOST1, on the taskbar, click **File Explorer**.
2. Double-click **Local Disk (C:)**.
3. On the title bar, click the **New folder** icon.
4. Name the folder **StoragePool**.
5. Close File Explorer
6. In Hyper-V Manager, right-click **20414C-LON-SVR1**, and then click **Settings**.
7. In the **Settings for 20414C-LON-SVR1 on LON-HOST1** dialog box, click **SCSI Controller**.
8. Click **Hard Drive**, and then click **Add**.
9. In the Media area, click **New**.
10. On the **Before You Begin** page of the New Virtual Hard Disk Wizard, click **Next**.
11. On the **Choose Disk Format** page, click **VHDX**, and then click **Next**.
12. On the **Choose Disk Type** page, click **Dynamically expanding**, and then click **Next**.
13. On the **Specify Name and Location** page, type **iSCSI1.vhdx**, and then click **Browse**.
14. Navigate to **C:\StoragePool**, and then click **Select Folder**.
15. On the **Specify Name and Location** page, click **Next**.
16. On the **Configure Disk** page enter the size as **50 GB**, and then click **Finish**.
17. On the **Hard Drive** page, click **Apply**.
18. Repeat steps 7 to 17 to create disks **iSCSI2.vhdx** and **iSCSI3.vhdx**.
19. Click **OK** to close the **Settings for 20414C-LON-SVR1 on LON-HOST1** dialog box.
20. In the Actions pane, click **Start**, and then click **Connect**.
21. Sign in to **LON-SVR1** as **Adatum\Administrator** with the password **Pa\$\$w0rd**.

Add the iSCSI Target Server Role Service

1. On LON-SVR1, on the Manage menu 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 a 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**, select the **iSCSI Target Server** check box, and then click **Next**.
6. On the **Select features** page, click **Next**.
7. On the **Confirm installation selections** page, click **Install**.
8. When the installation completes, click **Close**.

Create a storage pool

1. In Server Manager, in the navigation pane, click **File and Storage Services**, and then click **Storage Pools**.
2. Confirm that the disks you have added are visible in the Physical Disks section. If they are not, in the Storage Pool section, click **Tasks**, and then click **Refresh**.
3. In the Storage Pool section, click **Tasks**, and then click **New Storage Pool**.
4. On the **Before You Begin** page, click **Next**.
5. In the **Storage Pool Name** page, In the Name field, type **VMPool**, and then click **Next**.
6. On the **Physical Disks** page, select all three disks, and then click **Automatic** in Allocation. Note that you can assign a hot spare. Leave the selection as **Automatic**, and then click **Next**.
7. On the **Confirmation** page, click **Create**.
8. On the **Results** page, click **Create a virtual disk when this wizard closes**, and then click **Close**. The New Virtual Disk Wizard launches.
9. On the **Before you begin** page, click **Next**, and on the **Storage Pool** page, click **Next**. On the **Virtual Disk Name** page, in the **Name** field, type **VMStorage**, and then click **Next**.
10. On the **Storage Layout** page, click **Parity**, and then click **Next**. On the **Provisioning** page, click **Thin**, and then click **Next**.
11. On the **Size** page, in the **Specify size**, type **100**, and then click **Next**.
12. On the **Confirmation** page, review the settings, and then click **Create**.
13. On the **View results** page, click **Close**.
14. The New Volume Wizard launches. On the **Before You Begin** page, click **Next**. On the **Server and Disk** page, in the Disk area, click **VMStorage** Virtual Disk, and then click **Next**. On the **Size** page, leave the default (**99.9**) **GB**, and then click **Next**.
15. On the **Drive Letter or Folder** page, set the drive letter to **F**, and then click **Next**. On the **File System Settings** page, click the **Volume Label** field, and then type **VMStorage**. Click **Next**, review the settings, click **Create**, and then on the **Results** page, click **Close**.
16. In the File and Storage Services pane, click **iSCSI**.
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 **F:**, and then click **Next**.
19. On the **Specify iSCSI virtual disk name** page, in the **Name** field type **LONHOST1-iSCSIDisk1**, and then click **Next**.
20. On the **Specify iSCSI virtual disk size** page, in the **Size** box, type **90**, in the drop-down list box, ensure **GB** is selected, click **Dynamically expanding**, and then click **Next**.
21. On the **Assign iSCSI target** page, click **New iSCSI target**, and then click **Next**.
22. On the **Specify target name** page, in the **Name** box, type **LON-HOST1**, and then click **Next**.
23. On the **Specify access servers** page, click **Add**.
24. In the **Add Initiator ID** dialog box, click **Browse**.
25. In the **Select Computer** dialog box, type **LON-HOST1**, click **Check Names**, and then click **OK**.

26. In the **Add initiator ID** dialog box, click **OK**.
27. On the **Specify access servers** page, click **Next**.
28. On the **Enable Authentication** page, click **Next**.
29. On the **Confirm selections** 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 location**, click **C:**, and then click **Next**.
33. On the **Specify iSCSI virtual disk name** page, in the **Name** field type **iSCSIDisk2**, and then click **Next**.
34. 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**.
35. On the **Assign iSCSI target** page, click **LON-HOST1**, and then click **Next**.
36. On the **Confirm selections** page, click **Create**.
37. On the **View results** page, wait until creation completes, and then click **Close**.

► Task 2: Configure iSCSI initiators

Configure iSCSI initiators

1. On LON-HOST1, in Server Manager, click **Tools**, and then click **iSCSI Initiator**. When prompted to start the Microsoft® iSCSI service, click **Yes**.
2. On the **Targets** page, in the **Target** field, type **172.16.0.12**. Click **Quick connect**, click **Done**, and then click **OK** to close the page.
3. On LON-HOST1, right-click the Start menu, and then click **Disk Management**.
4. Find the new 90 GB disk, right-click the disk, and then click **Online**. Right-click again, click **Initialize Disk**, and then on the **Initialize disk** page, click **OK**.
5. Right-click the unallocated space, and then click **New Simple Volume**.
6. On the **Welcome** page, click **Next**, on the **Specify Volume Size** page, leave the default value, and then click **Next**.
7. On the **Assign Driver letter or Path** page, click the drop-down list box, and then select the letter **V**. Click **Next**.
8. On the **Format Partition** page, in the **Volume label** field, type **VMStorage**. Click **Next**, review the settings, and then click **Finish**.
9. Close **Disk Management**.

Results: After completing this exercise, you will have configured iSCSI targets and iSCSI initiators and implemented iSCSI storage.

Exercise 4: Implementing a Network Infrastructure for Virtualization

► Task 1: Configure logical networks that your design requires

Create the logical network

1. On LON-VMM1, on the task bar, click **Virtual Machine Manager Console**. In the **Connect to Server** dialog box, click **Connect**.
2. Click the **Fabric** workspace, on the ribbon, click **Create**, and then click **Logical Network**.
3. In the **Name** field, type **Toronto Production Network**, and in the **Description** field, type **Adatum Toronto – Production Logical Network**. Click **Allow new VM networks created on this logical network to use network virtualization**, and then click **Next**.
4. On the **Network site** page, click **Add**, and then in the Host groups section, select **All Hosts**.
5. Click **Insert row**, and click the **VLAN** field. Type **0**, click the **IP subnet** field, type **172.16.3.0/24**, click **Next**, and then click **Finish**.
6. Close the Jobs window.

Create an IP pool

1. Click the **Fabric** workspace, on the ribbon, click **Create**, and then click **IP Pool**. On the **Name** page, in the **Name** field, type **Toronto Apps**, and in the description field, type **Toronto Production Application IP Pool**. Click the **Logical network** drop-down list box, click **Toronto Production Network**, and then click **Next**.
2. On the **Network site** page, click **Use an existing network site**, select **Toronto Production Network_0**, click the **IP subnet** drop-down list box, select **172.16.3.0/24**, and then click **Next**.
3. On the **IP address range** page, in the **VIPs and Reserved IP addresses** field, type **172.16.3.100-172.16.3.120**, and then click **Next**.
4. On the **Gateway** page, click **Insert**, click the **Gateway address** field, type **172.16.3.1**, and then click **Next**.
5. On the **DNS** page, click **Next**, and then on the **WINS** page, click **Next**.
6. On the **Summary** page, click **Finish**.
7. Close the Jobs window.

Create a native port profile

1. On the ribbon, click **Create**, and then click **Hyper-V Port Profile**.
2. On the **General** page, in the **Name** field, type **Toronto Default Network Adapter**, click **Uplink port profile**, leave the default load balancing and teaming settings, and then click **Next**.
3. On the **Network configuration** page, under **Network sites**, click **Toronto Production Network_0**, click **Enable Hyper-V Network Virtualization**, and then click **Next**.
4. On the **Summary** page, click **Finish**.
5. Close the Jobs window.

Create a logical switch

1. On the ribbon, click **Create**, and then click **Logical Switch**.
2. On the **Getting Started** page, click **Next**.
3. On the **General** page, in the **Name** field, type **Toronto Logical Switch**, in the **Description** field, type **Toronto Production Hyper-V Switch**, and then click **Next**.
4. On the **Extensions** page, leave the default extensions, and then click **Next**.

5. On the **Uplink** page, click **Add**, click **Port profile**, and select the **Toronto Default Network Adapter**. Click **OK**, and then click **Next**.
6. On the **Virtual Port** page, click **Add**, click **Browse**, click **Host management**, click **OK**, click **Include a virtual network adapter port profile in this virtual port**, click **Native virtual network adapter port profile**, click **Host management**, and then click **OK**.
7. On the **Virtual Port** page, click **Add**, click **Browse**, click **Live migration workload**, click **OK**, click **Include a virtual network adapter port profile in this virtual port**, click **Native virtual network adapter port profile**, click **Live migration**, and then click **OK**.
8. On the **Virtual Port** page, click **Add**, click **Browse**, click **Host Cluster Workload**, click **OK**, click **Include a virtual network adapter port profile in this virtual port**, click **Native virtual network adapter port profile**, click **Cluster**, and then click **OK**.
9. On the **Virtual Port** page, click **Add**, click **Browse**, click **High bandwidth**, click **OK**, click **Include a virtual network adapter port profile in this virtual port**, click **Native virtual network adapter port profile**, click **High bandwidth adapter**, and then click **OK**.
10. On the **Virtual Port** page, click **Next**, and then on the **Summary** page, click **Finish**.
11. Close the Jobs windows.

Add the Microsoft Loopback adapter

1. On LON-HOST1, click **Start**, on the Start screen, click **Control Panel**, in the Control Panel, click **Hardware**, and then on the **Hardware** page, click **Device Manager**.
2. In Device Manager, right-click **LON-HOST1**, and then click **Add legacy hardware**. The Add Hardware Wizard launches.
3. On the Welcome screen, click **Next**, click **Install the hardware that I manually select from a list**, click **Next**, scroll down until you see Network adapters, click **Network adapters**, click **Next**, under Manufacturer, click **Microsoft**, under Network Adapter, scroll down until you see Microsoft KM-TEST Loopback Adapter, click **Microsoft KM-TEST Loopback Adapter**, and click **Next**. The Add Hardware Wizard should now show that the hardware is ready to install. Click **Next**, and then on the **Completing the Add Hardware Wizard** page, click **Finish**.
4. Close the Device Manager and Control Panel.

Refresh LON-HOST1

1. On LON-VMM1, in the Virtual Machine Manager console, click the **Fabric** workspace, expand **Servers**, expand **All Hosts**, click **London Hosts**, right-click **LON-HOST1**, and then click **Refresh**.
2. On the ribbon, click Home, click **Jobs**, wait for the host job to finish refreshing, and then close the Jobs window.

Update Hyper-V hosts to use logical networks

1. From the Fabric workspace, click **LON-HOST1.adatum.com**, on the ribbon, click **Properties**, click **Virtual Switches**, click **New Virtual Switch**, and then click **New Logical Switch**.
2. With **Toronto Logical Switch** selected, under **Physical adapters**, click the **Adapter** drop-down list box, and then select **Microsoft KM-TEST Loopback Adapter**.
3. On the **Properties** page, click **Hardware**, scroll down, and then expand **Network adapters**. Click **Microsoft KM-TEST Loopback Adapter**. You can see that the adapter is available for placement.
4. Click the **Logical network** that is listed under **Microsoft KM-TEST Loopback Adapter**, and then on the right side of the **Hardware** page, under **Logical network connectivity**, click **Toronto Production Network**, click **OK**, read the warning, and then click **OK**.

► Task 2: Configure network virtualization

Configure a virtual machine network for the application servers by using network virtualization

1. In the Virtual Machine Manager console, click the **VMs and Services** workspace. On the ribbon, click **Create VM Network**.
2. On the **Name** page, in the **Name** field, type **Toronto Applications VM Network**, click the **Description** field, and then type **Toronto Application Servers**. Click the **Logical network** drop-down list box, select **Toronto Production Network**, and then click **Next**.
3. On the **Isolation** page, click **Isolate using Hyper-V network-virtualization**, and then click **Next**.
4. On the **VM Subnets** page, click **Add**, and in the **Name** field, type **Toronto Application Servers**. In the **Subnet** field, type **172.16.3.0/24**, and then click **Next**.
5. On the **Connectivity** page, click **Next**. Review the summary, and then click **Finish**.
6. Close the Jobs window.
7. Repeat steps 1 through 6 to create a VM network named **Toronto Partner_Applications VM Network**, and use the same subnet.

Create the VM Network IP pools

1. In the VMs and Services workspace, click **VM Networks**, click **Toronto Applications VM Network**, and then right-click and click **Create IP Pool**.
2. On the **Name** page, click the **Name** field, and then type **Toronto Applications VM IP Pool**. Make sure that the VM Network is set to **Toronto Applications VM Network** and that the VM subnet is set to **Toronto Application Servers (172.16.3.0/24)**, and then click **Next**.
3. On the **IP address range** page, leave the default values, and then click **Next**.
4. On the **Gateway** page, click **Next**, and on the **DNS** page, click **Next**. On the **WINS** page, click **Next**, and then on the **Summary** page, click **Finish**. Close the Jobs window.

► Task 3: Assign virtual machines to VM networks

Assign virtual machines to VM Networks

1. In the Virtual Machine Manager console, click the **VMs and Services** workspace. On the ribbon, on the **Home** tab, click the upper half of the **Create Virtual Machine** button.
2. On the **Select Source** page, click **Create the new virtual machine with a blank virtual hard disk**, and then click **Next**.
3. On the **Specify Virtual Machine Identity** page, in the **Virtual machine name** field, type **TOR-CRM1**, and then click **Next**.
4. On the **Configure Hardware** page, click **Network Adapter 1**, and then under **Connectivity**, click **Connected to a VM Network**. Click **Browse**, select **Toronto_Applications VM Network**, and then click **OK**. Under the Port Profile section, click the **Classification** drop-down list box. This is where you can assign the native port virtual adapter profiles. Select **High bandwidth**, and then click **Next**.
5. On the **Select Destination** page, click **Next**.
6. On the **Select Host** page, in the details section, click the **Rating Explanation** for LON-HOST1. You should see **This destination meets all the requirements of this virtual machine**. Click **Next**.
7. On the **Configure Settings** page, leave the default settings, and then click **Next**. On the **Add Properties** page, click **Next**, and then on the **Summary** page, click **Create**.
8. Close the Jobs window.

Review the virtual machine

1. In the VMs and Service workspace, on the ribbon, on the **Home** tab, click **VMs**, right-click **TOR-CRM1**, and then click **Properties**.
2. Click **Hardware Configuration**, and then click **Network Adapter 1**. You will see the VM Network and Logical switch. Click **Connection details**, and you should see the IP details assigned from the pool. Click **OK**, and then click **Cancel**.

► Task 4: To prepare for the next module

- Do not revert the virtual machines, as you will need them for the next module.

Module 4: Planning and Deploying Virtual Machines

Lab: Planning and Implementing a Virtual Machine Deployment and Management Strategy

Exercise 1: Planning Virtual Machine and Service Templates

► **Task 1: Read the supporting documentation**

- Read the documentation provided in the student workbook.

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

- Answer the questions in the proposals section of the Development Service Template Reference Document.

► **Task 3: Examine the suggested proposals in the Lab Answer Key**

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

1. How many hardware, guest operating system, and Microsoft® SQL Server® profiles will you require?

Answer: Because you have five different systems with different hardware requirements, you will require five hardware profiles.

2. Should you use service templates or virtual machine templates?

Answer: You will require a virtual machine template before you can create a service template. You can create a service template for the application using three servers. You can choose to either include the clients or deploy them separately.

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

- Be prepared to discuss your proposals with the class.

Results: After completing this exercise, you should have planned a basic service template based on an existing infrastructure.

Exercise 2: Configuring VMM Profiles and Templates

► Task 1: Configure a guest operating system profile

Configure a guest operating system profile

1. On LON-VMM1, on the taskbar, click **Virtual Machine Manager Console**.
2. In the **Connect to Server** dialog box, ensure that the **Use current Microsoft Windows session identity** check box is selected, and then click **Connect**.
3. In the VMM console, click the **Library** workspace, expand **Profiles**, and then click **Guest OS Profiles**.
4. Right-click **Guest OS Profiles**, and then click **Create Guest OS Profile**.
5. In the **New Guest OS Profile** dialog box, on the **General** page, in the **Name** text box, type **TOR-WEB OS Profile**, and then in the **Description** text box, type **Guest OS Profile for new development Web Server**.
6. Click **Guest OS Profile**.
7. On the **Guest OS Profile** page, under **General Settings**, click **Identity Information**.
8. In the **Computer name** text box, type **TOR-WEB#**.
9. Click **Admin Password**, and then click **Specify the password of the local administrator account**. In the **Password** and **Confirm** text boxes, type **Pa\$\$w0rd**.
10. Click **Operating System**, and then ensure **Windows Server 2012 R2 Standard** is selected.
11. Under **Networking**, click **Domain/Workgroup**.
12. Click **Domain**, and then in the **Domain** text box, type **adatum.com**.
13. Under **Domain credentials**, click **Specify credentials to use for joining the domain**.
14. In the **Domain user** text box, type **Adatum\Administrator**. In the **Password** and **Confirm** text boxes, type **Pa\$\$w0rd**.
15. Click **OK** to close the **New Guest OS Profile** dialog box.

► Task 2: Configure a hardware profile

Configure a hardware profile

1. In the VMM console, click the **Library** workspace, expand **Profiles**, and then click **Hardware Profiles**.
2. On the ribbon, click **Create**, and then click **Hardware Profile**.
3. In the **New Hardware Profile** dialog box, on the **General** page, in the **Name** text box, type **WsStd2012R2**, and then in the **Description** text box, type **Hardware Profile for new Windows Server 2012 R2 Servers**.
4. Click **Hardware Profile**.
5. On the **Hardware Profile** page, under **Compatibility**, click **Cloud Capability Profile**.
6. Select the **Hyper-V** check box.
7. In the **General** section, click **Processor**, and then select the **Allow migration to a virtual machine host with a different processor version** check box.
8. Click **Memory**, verify that **Static** is selected, and then change the **Virtual machine memory** option to **1024 MB**.
9. In the **Network Adapters** section, click **Network Adapter 1**. Under **Connectivity**, click **Connected to a VM network**, click **Browse**, click **External Network**, and then click **OK**.
10. Click **OK** to close the **New Hardware Profile** dialog box. The new profile displays in the results pane.

► Task 3: Configure a SQL Server profile

Configure a SQL Server profile

1. In the VMM console, click the **Library** workspace, expand **Profiles**, and then click **SQL Server Profiles**.
2. Right-click SQL Server Profiles, and then click Create SQL Server Profile.
3. In the **New SQL Server Profile** dialog box, on the **General** page, in the **Name** text box, type **SQLDev1**, and then in the **Description** text box, type **Template for new SQL servers**.
4. Click **SQL Server Configuration**.
5. On the **SQL Server Configuration** page, next to **Add**, click **SQL Server Deployment**.
6. Configure the following settings:
 - In the **Name** field, type **SQLDev1**
 - In the **Instance name** field, type **MSSQLSERVER**
 - In the **Instance ID** field, type **DefaultInstance**
7. Under **Installation Run As account**, click **Browse**. In the **Browse Run As Accounts** dialog box, click **Administrator**, and then click **OK**.
8. Click **Configuration**.
9. In the **Media source** text box, type **C:\SQLInstall**.
10. In the SQL Server administrators, text box, type **Adatum\Administrator**, and then click **Add**.
11. Next to **Security mode**, verify that **Windows Authentication** is selected.
12. Select the **Use TCP\IP for remote connections** check box, and then click **Service Accounts**.
13. Under **SQL Server service Run As Account**, click **Browse**.
14. In the **Browse Run As Accounts** dialog box, click **Administrator**, and then click **OK**.
15. Under **SQL Agent service Run As Account**, click **Browse**.
16. In the **Browse Run As Accounts** dialog box, click **Administrator**, and then click **OK**.
17. Under **Reporting Services service Run As Account**, click **Browse**.
18. In the **Browse Run As Accounts** dialog box, click **Administrator**, and then click **OK**.
19. Click **OK** to close the New SQL Server Profile dialog box.

► Task 4: Configure a virtual machine template

Configure a virtual machine template

1. On LON-VMM1, open File Explorer.
2. In the address bar, type **\\LON-HOST1\e\$\Program Files\Microsoft Learning\Base**, and then press Enter.



Note: You may need to substitute f\$ for e\$ in the above path depending on your environment.

3. Right-click **Base14A-WS12R2.vhd**, and then click **Copy**.
4. In the address bar, type **\\lon-vmm1\MSSCVMLibrary\vhd** and press Enter. In the File Explorer window, in an empty area, right-click and then click **Paste**. The file will take several minutes to copy.
5. In the VMM console, click the **Library** workspace, expand **Library Servers**, expand **LON-VMM1.Adatum.com**, expand **MSSCVMLibrary**, and then click **VHDs**.

6. Confirm that the file has finished copying, right-click **VHDs**, and then click **Refresh**.
7. In the VMM console, click the **Library** workspace, expand **Templates**, and then click **VM Templates**.
8. On the ribbon, click **Create VM Template**.
9. In the Create VM Template Wizard, on the **Select Source** page, click **Use an existing VM template or a virtual hard disk stored in the library**, and then click **Browse**.
10. In the **Select VM Template Source** dialog box, click Base14A-WS12R2.vhd, and then click **OK**.
11. On the **Select Source** page, click **Next**.
12. On the **VM Template Identity** page, in the VM Template name text box, type **Adatum Web Application Server**, in the Description text box, type **Web Server hosting the Adatum Web Application**, and then click **Next**.
13. On the **Configure Hardware** page, in the **Hardware profile** drop-down list box, click **WsStd2012R2**. Notice that the settings from the hardware profile import into the template, and then click **Next**.
14. On the **Configure Operating System** page, in the **Guest OS profile** drop-down list box, click **TOR-WEB OS Profile**.
15. Under **Roles and Features**, click **Roles**.
16. Select the **Web Server (IIS)** check box, and then click **Next**.
17. On the **Applications Configuration** page, in the **Application profile** drop-down list box, click **None – do not install any applications**, and then click **Next**.
18. On the **SQL Server Configuration** page, in the **SQL Server profile** drop-down list box, click **None - no SQL Server configuration settings**, and then click **Next**.
19. On the **Summary** page, click **Create**.

► **Task 5: Configure a service template**

Configure a service template

1. In the VMM console, click the **Library** workspace, expand **Templates**, and then click **Service Templates**.
2. On the ribbon, click **Create Service Template**.
3. In the VMM Service Template Designer, in the **New Service Template** dialog box, in the **Name** text box, type **Adatum Web Service**.
4. Under **Patterns**, click **Single Machine**, and then click **OK**. Wait while Adatum Web Service loads into the Template Designer.
5. Under **VM Templates**, click and drag the **Adatum Web Application Server** to the Add applications section of the tier.
6. On the ribbon, click **Save and Validate**.
7. Close the VMM Service Template Designer.

Results: After completing this exercise, you should have configured Microsoft System Center 2012 R2 Virtual Machine Manager (VMM) profiles and templates.

Exercise 3: Implementing Hyper-V Replica

► Task 1: Configure a replica on both host machines

1. On LON-HOST1, open the Hyper-V Manager console.
2. In Hyper-V Manager, right-click **LON-HOST1**, and then select **Hyper-V Settings**.
3. In Hyper-V Settings for LON-HOST1, 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 **This PC**, double-click **Local Disk (E)**, and then click **New folder**. Type **VMReplica** for folder name, and then press Enter.
8. Select the **E:\VMReplica** folder, and then click **Select Folder**. (Note: The drive letter might change depending on your host hardware configuration.)
9. In Hyper-V Settings for LON-HOST1, click **OK**.
10. In the Settings window, read the notice, and then click **OK**.
11. Click to the **Start** screen, and then click **Control Panel**.
12. In the Control Panel, click **System and Security**, and then click **Windows Firewall**.
13. Click **Advanced settings**.
14. Click **Inbound Rules**, in the right pane, in the rule list, right-click the rule **Hyper-V Replica HTTP Listener (TCP-In)**, and then click **Enable Rule**.
15. Close the Windows Firewall with Advanced Security console, and then close **Windows Firewall**.
16. Repeat steps one through 15 on LON-HOST2.

► Task 2: Configure replication for the virtual machine

1. On LON-HOST1, open Hyper-V Manager. Click **LON-HOST1**, and then right-click **20414C-LON-CORE**.
2. Click **Enable Replication**, and then 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 the settings, make sure that **Use Kerberos Authentication (HTTP)** is selected, and then click **Next**.
6. On the **Choose Replication VHDs** page, make sure that **20414C-LON-CORE.avhd** is selected, and then click **Next**.
7. On the **Configure Replication Frequency** page, click **Next**.
8. On the **Configure Additional Recovery Points** page, select **Maintain only the latest recovery point**, and then click **Next**.

9. On the **Choose Initial Replication Method** page, click **Send initial copy over the network**, select **Start replication immediately**, and then click **Next**.
10. On the **Completing the Enable Replication wizard** page, click **Finish**.



Note: Wait approximately five to 10 minutes. You can monitor the progress of the initial replication in the **Status** column in the Hyper-V Manager console on LON-HOST1. When it completes (progress reaches 100%), make sure that 20414C-LON-CORE has appeared on LON-HOST2 in Hyper-V Manager.

► **Task 3: Validate a planned failover to the Replica site**

1. On LON-HOST2 in Hyper-V Manager, right-click **20414C-LON-CORE**.
2. Select **Replication**, and then click **View Replication Health**.
3. Review the content of the window that appears, and then make sure that there are no errors.
4. Click **Close**.
5. On LON-HOST1, open Hyper-V Manager, and then verify that 20414C-LON-CORE is turned off.
6. Right-click **20414C-LON-CORE**, select **Replication**, and then click **Planned Failover**.
7. In the Planned Failover window, make sure that **Start the replica virtual machine after failover** is selected, and then click **Fail Over**.
8. On LON-HOST2, in Hyper-V Manager, make sure that **20414C-LON-CORE** is running.
9. On LON-HOST1, right-click **20414C-LON-CORE**, point to **Replication**, and then click **Remove Replication**.
10. In the **Remove Replication** dialog box, click **Remove Replication**.
11. On LON-HOST2, right-click **20414C-LON-CORE**, and then select **Shut Down**. In the **Shut Down Machine** dialog box, click **Shut Down**.

► **Task 4: Prepare for the next module**

1. In Hyper-V Manager on LON-HOST1, click **Virtual Switch Manager** in the Actions pane, click **Toronto Logical Network** and click **Remove**. Click **OK** to close Virtual Switch Manager. Click **Yes** on the warning dialog box.
2. Do not revert the virtual machines, as you will need them for the next module.

Results: After completing this exercise, students will have implemented Hyper-V® Replica.

Module 5: Planning and Implementing a Virtualization Administration Solution

Lab: Planning and Implementing an Administration Solution for Virtualization

Exercise 1: Configuring Process Automation in System Center

► Task 1: Install and configure System Center Integration Packs for VMM

Install the System Center Integration Pack for VMM

1. On LON-DC1, from the task bar, click **File Explorer**.
2. In File Explorer, browse to **E:\Labfiles**.
3. Double-click **System_Center_2012_R2_Integration_Packs.EXE**. Click **OK**, and then click **OK** again to close the Extraction complete message.
4. Close File Explorer.
5. On LON-OR1, click the **Start** icon, when the Start page appears, click the down arrow, and then click **Deployment Manager**.
6. In the Deployment Manager console, right-click **Integration Packs**, and then click **Register IP with the Orchestrator Management Server**.
7. On the **Welcome to the Integration Pack Registration Wizard** page, click **Next**, and then on the **Select Integration Packs or Hotfixes** page, click **Add**.
8. On the **Open** page, in the **File name** field, type **\\lon-dc1\e\$\Labfiles\SC2012R2_Integration_Pack_for_Virtual_Machine_Manager.oip**. Click **Open**, click **Next**, and then on the **Completing the Integration Pack Wizard** page, click **Finish**. In the Microsoft Software License Terms window, click **Accept**.
9. Wait until the registration is complete, click and expand **Orchestration Management Server**, and then click **Integration Packs**.
10. Right-click **System Center Integration Pack for System Center 2012 Virtual Machine Manager**, and then click **Deploy IP to Runbook Server or Runbook Designer**.
11. On the **Welcome to the Integration Pack Deployment Wizard** page, click **Next**, on the **Deploy Integration Packs or Hotfixes** page, select the check box next to **System Center Integration Pack for System Center 2012 Virtual Machine Manager**, and then click **Next**.
12. On the **Computer Selection** page, in the **Computer** field, type **LON-OR1**, click **Add**, and then click **Next**.
13. On the **Installation Options** page, click **Next**, and then on the **Completing the Integration Pack Deployment Wizard** page, click **Finish**.
14. Review the log entries, and then close the Orchestrator Deployment Manager.

Set the Windows PowerShell execution policy to RemoteSigned

1. On LON-OR1, on the task bar, right-click **Windows PowerShell**, and then under **Tasks**, click **Run as Administrator**.
2. At the Windows PowerShell® prompt, type **set-executionpolicy remotesigned**, press Enter, type **Y**, and then press Enter.

3. Close the Windows PowerShell window.
4. On LON-VMM1, on the task bar, right-click **Windows PowerShell**, and then under **Tasks**, click **Run as Administrator**.
5. At the Windows PowerShell prompt, type **set-executionpolicy remotesigned**, press Enter, type **Y**, and then press Enter.
6. Close the Windows PowerShell window.

Enable Remote Management Trusted Hosts

1. On LON-OR1, right-click the **Start** icon and click **Run**. On the **Run** page, in the **Open** field, type **gpedit.msc**, and then click **OK**.
2. On the left side, under **Local Computer Policy\Computer Configuration**, click to expand **Administrative Templates**, double-click **Windows Components**, and then double-click **Windows Remote Management (WinRM)**.
3. Click **WinRM Client**, and then on the right, under **WinRM Client**, double-click **Trusted Hosts**.
4. On the **Trusted Hosts** page, click **Enabled**, in the **TrustedHostList** field, type ***** and then click **OK**.
5. Close the Local Group Policy editor.
6. Repeat steps one through five on LON-VMM1.

Configure the System Center Integration Pack for VMM

1. On LON-OR1, click **Start**, click the down arrow, and then click **Runbook Designer**.
2. On the menu, click **Options**, and then click **SC 2012 Virtual Machine Manager**.
3. On the **Configurations** page, click **Add**, on the **Add Configuration** page, in the **Name** field, type **LON-VMM1**, and then click the browse (...) button.
4. On the **Item Selection** page, click **System Center Virtual Machine Manager**, and then click **OK**.
5. On the **Add Configuration** page, under **Properties**, in the **VMM Administrator Console** field, type **LON-VMM1**. In the **VMM Server** field, type **LON-VMM1**, in the **User** field, type **Adatum\Administrator**, delete the text in the **Domain** field, and then in the **Password** field, type **Pa\$\$w0rd**.
6. Click the **Authentication Type (Remote only)** field, click **Browse**, click **Negotiate**, and then click **OK**.
7. Click **OK** again, and then on the **Prerequisite Configuration** page, click **Finish**.

► Task 2: Configure automation in Orchestrator

Create a basic runbook

1. On LON-OR1, click the **Start** icon, click on the down arrow, and then click **Runbook Designer**.
2. In the Connections pane, right-click **Runbooks**, click **New**, click **Folder**, press the Delete key, type **20414 Runbooks**, and then press Enter.
3. Right-click the **20414 Runbooks** folder, click **New**, click **Runbook**, at the top of the central pane, right-click **New Runbook**, click **Rename**, click **Yes** to confirm the check-out of this runbook, type **VMM Library Monitor**, and then press Enter.
4. On the right pane under **Activities**, expand **File Management**, and then click and drag the **Monitor Folder** activity to the center of the central pane.
5. Right-click the **Monitor Folder** activity, click **Rename**, and then type **VMM Library Monitor**. Press Enter.

6. Right-click the **VMM Library Monitor**, click **Properties**, and then click **General**. On the **General Information** page, in the **Description** field, type **This Runbook monitors the VMM library for new virtual hard disks**.
7. Click **Details**, under folder to monitor, in the **Path** field, type **\\LON-VMM1\MSSCVMMLibrary**, and then click **Include sub-folders**.
8. In the File Filters section, click **Add**, on the **Filter Settings** page, click the **Name** drop-down list box, click **File Name**, in the **Value** field, type ***.vhd**, and then click **OK**.
9. On the left, click **Triggers**, in the **Trigger if** section, select the **Number of files is** check box, click the **Number of files is** drop-down list box, select **greater than**, and then in the **greater than** field, type **0**.
10. Click **Authentication**, in the **User name** field, type **Adatum\Administrator**, in the **Password** field, type **Pa\$\$w0rd**, and then click **Finish**.
11. Under **Activities**, click **Notification**, and then click and drag the **Send Event Log Message** activity to the central pane and to the right of the **VMM Library Monitor** activity.
12. Place the pointer over the **VMM Library Monitor** activity and then a small arrow should appear to the right. Place the pointer over the arrow and the pointer should change to a cross. Click the arrow, and then drag it to the **Send Event Log Message** activity. A link with an arrow should now appear between the two activities.
13. Right-click the arrow between the two activities, click **Properties**, when the **Link Properties** page appears, review the filter, and then click **Finish**.
14. Right-click **Send Event Log Message**, click **Properties**, and on the **Details** page, in the Properties section, in the **Computer** field, type **LON-OR1**. In the **Message** field, type **A virtual hard disk file was created or updated in the LON-VMM1 library**, in the **Severity** section, click **Warning**, and then click **Finish**.
15. On the ribbon, click **Check In**, and then click **Runbook Tester**.
16. On the **Confirm Check out** dialog box, click **Yes**.
17. In the Runbook Tester, click **Run**.
18. On the Windows task bar, click the **File Explorer** icon. In the address bar field, type **\\lon-vmm1\MSSCVMMLibrary\VHDs**, and then press Enter.
19. On the ribbon of the VHDs window, click **View**, and then select the checkbox next to **File name extensions**.
20. Right-click any of the **Blank Disk – Large.vhd** files, click **Copy**, right-click an empty space in the File Explorer window, and then click **Paste**. A new file is created called **Blank Disk – Large – Copy.vhd**.
21. Switch to the Runbook Tester. In the log section, wait until **Activity name Send Event Log Message** appears.
22. Switch to the **Start** page, type **Event**, and then click **Event Viewer**. In the center, in the Summary of Administrative Events pane, expand **Warning**, and then you should see an Event ID with the ID of **1** and a Source of **Orchestrator Runbook**. Double-click **Orchestrator event**, and then review the event.
23. Close the Event Viewer, and then close the File Explorer window.
24. Close the Runbook Tester.

Results: After this lab, you will have installed the Microsoft® System Center 2012 Integrations Pack for Virtual Machine Manager (VMM), created a basic runbook in the System Center 2012 Orchestrator Runbook Designer, and reviewed the Orchestrator web console.

Exercise 2: Planning Administrative Delegation and Self-Service in System Center 2012

► Task 1: Read the supporting documentation

- Read the documentation that the student workbook provides.

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

- Answer the questions in the proposals section of the Virtualization Administration Strategy document

1. Which administrative role would best meet the primary control objectives of the development administrative team?

Answer: The delegated fabric administrator would allow the development administrators to manage a scoped group of resources. This should consist of private clouds and host groups. They can add various resources, as necessary.

2. Most developers will require the ability to create self-service accounts. Which administrative role will be best for them?

Answer: The developers will require fewer permissions than the development administrators, so the Tenant or self-service roles will be suitable.

3. What can you do to help reduce the administrative burden on the development administrators?

Answer: Ensure that you configure self-service for those users who demand new systems most frequently. Additionally, implement Orchestrator and possibly even System Center 2012 Service Manager.

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

- Compare your proposals with those shown above.

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

- Be prepared to discuss your proposals with the class.

Exercise 3: Configuring Delegated Administration and Self-Service in VMM

► Task 1: Configure a delegated administrator role in VMM

Create a private cloud

1. On LON-VMM1, from the desktop, open the Virtual Machine Manager console.
2. In the **Connect to Server** dialog box, ensure that the **Use current Microsoft Windows session identity** check box is selected, and then click **Connect**. The Virtual Machine Manager console opens.
3. Click the **VMs and Service** workspace, and then on the ribbon, click **Create Cloud**.
4. On the **General** page, in the **Name** field, type **London Development**, in the **Description** field, type **London Development Cloud**, and then click **Next**.
5. On the **Resources** page, select **London Hosts**, and then click **Next**.
6. On the **Logical Networks** page, select **External Network**, and then click **Next**.
7. On the **Load Balancers** page, select **Microsoft Load Balancing (NLB)**, and then click **Next**.
8. On the **VIP Template** page, click **Next**.
9. On the **Port Classifications** page, select **Network load balancing, Medium Bandwidth**, and **High Bandwidth**, and then click **Next**.
10. On the **Storage** page, click **Next**.
11. On the **Library** page, in the Read-only library shares section, click **Add**. Select **MSSCVMMLibrary**, click **OK**, and then click **Next**.
12. On the **Capacity** page, review the capacity options. Clear the check box next to each selected resource, and then assign the following:
 - 8 virtual central processing units (CPUs)
 - 12 gigabyte (GB) memory
 - 250 GB storage
 - 15 quota points
 - 4 virtual machines
13. Click **Next**.
14. On the **Capability Profiles** page, select **Hyper-V**, and then click **Next**.
15. Review the **Summary** page, and then click **Finish**.
16. Close the Jobs window.

Configure delegated administration in VMM

1. In the Virtual Machine Manager console, click the **Settings** workspace, and then on the ribbon, click **Create User Role**.
2. On the **Name and description** page, in the **Name** field, type **DevAdmin**, in the **Description** field, type **Development team administrators**, and then click **Next**.
3. On the **Profile** page, click **Fabric Administrator (Delegated Administrator)**, and then click **Next**.
4. On the **Members** page, click **Add**, in the **Select Users, Computers, or Groups** dialog box, in the **Enter the object names to select** field, type **Rob Cason**, click **OK**, and then click **Next**.

5. On the **Scope** page, select the **London Development** cloud and the **London Hosts** host group, and then click **Next**.
6. On the **Library servers** page, click **Add**, select **LON-VMM1.Adatum.com**, click **OK**, and then click **Next**.
7. On the **Run As accounts** page, click **Add**, select **Administrator** account, click **OK**, and then click **Next**.
8. Review the summary, and then click **Finish**.
9. Close the Jobs window.

► **Task 2: Configure self-service administration in VMM**

Configure self-service in VMM

1. In the Virtual Machine Manager console, click the **Settings** workspace, and then on the ribbon, click **Create User Role**.
2. On the **Name and Description** page, in the **Name** field, type **DevContractors**, in the **Description** field, type **Development team contractors**, and then click **Next**.
3. On the **Profile** page, click **Application Administrator (Self-Service User)**, and then click **Next**.
4. On the **Members** page, click **Add**, type **Adam**, click **OK**, and then click **Next**.
5. On the **Scope** page, click **London Development**, and then click **Next**.
6. On the **Quotas for the London Development cloud** page, leave the default quotas, and then click **Next**.
7. On the **Networking** page, click **Add**, click **External Network**, click **OK**, and then click **Next**.
8. On the **Resources** page, click **Add**, on the **Add Resources** page, hold down the Ctrl key on your keyboard, and then click each item in the list. When you have selected all of them, click **OK**, and then click **Next**.
9. On the **Permissions** page, assign the following permitted actions:
 - **Deploy**
 - **Remote connection**
 - **Shut down**
 - **Start**
 - **Stop**
10. When you have selected the permitted actions, click **Next**.
11. On the **Run As accounts** page, click **Add**, select **Administrator** Account, click **OK**, and then click **Next**.
12. On the **Summary** page, review the settings, and then click **Finish**.
13. If you receive an error warning "**Unable to perform the job because one or more of the selected objects are locked by another job...**" click **OK** to dismiss the warning.
14. Close the Jobs window.

► Task 3: Validate the configuration by using VMM

Verify delegation of administration

1. From the top menu in the Virtual Machine Manager console, above the ribbon, click the arrow, and then click **Open New Connection**. The Connect to Server page opens.
2. Sign in to the Virtual Machine Manager console by using the following credentials:
 - User name: **Adatum\Rob**
 - Password: **Pa\$\$w0rd**
3. Click the **VMs and Services** workspace.
4. Expand **Clouds**, right-click **London Development**, and then click **Create Virtual Machine**.
5. On the **Select Source** page, click **Create the new virtual machine with a blank virtual hard disk**, and then click **Next**.
6. On the **Specify Virtual Machine Identity** page, click the **Virtual machine name** field, type **RobVM**, and then click **Next**.
7. On the **Configure Hardware** page, under **Compatibility**, click **Cloud Capability Profile**, select the **Hyper-V** capability profile, and then click **Next**.
8. On the **Select Destination** page, select **Deploy the virtual machine to a private cloud**, and then click **Next**.
9. On the **Select Cloud** page, click **LON-HOST1.adatum.com**, and then click **Next**.
10. On the **Add Properties** page, click **Next**.
11. On the **Summary** page, click **Create**. Close the Jobs window, and then close the DevAdmin instance of the Virtual Machine Manager console.
12. In the Administrator instance of the Virtual Machine Manager console, click the **VMs and Services** workspace, and then click the arrow next to **Clouds**. You should see the London Development cloud.
13. Click the **London Development** cloud, and then on the ribbon, click **Overview**. Note that you can see User roles and Virtual Machine Owners. Confirm that you can see the DevAdmin and DevContractors roles. Click to expand these roles and view their assigned users. Review the details on this page.

► Task 4: Validate the configuration by using App Controller

Connect App Controller to VMM

1. On LON-VMM1, click the **Start** icon, and when the Windows® user interface appears, click **App Controller**.
2. On the **App Controller Credentials** page, in the **User name** field, type **Adatum\Administrator**, in the **Password** field, type **Pa\$\$w0rd**, and then click **Sign In**.
3. On the **Overview** page, under **Status**, click **Connect a Virtual Machine Manager server and clouds**.
4. On the **Add a new VMM connection** page, click the **Connection name** field, and then type **LON-VMM1.adatum.com**. In the **Description** field, type **London VMM Server access**. In the **Server name** field, type **LON-VMM1.adatum.com**, and then click **OK**.
5. On the **Overview** page, in the top right corner of the browser, click **Sign Out**.

Verify self-service in App Controller

1. On the **App Controller Credentials** page, in the **User name** field, type **Adatum\Adam**, in the **Password** field, type **Pa\$\$w0rd**, and then click **Sign In**.
2. On the **Overview** page, under **Next Steps**, click **Deploy a new service or virtual machine**. The New Deployment window opens. In the center of the screen, click **Configure**.
3. On the **Select a cloud for this deployment** page, click **OK**.
4. In the **Template** section, click **Select a template**, on the **Choose a template** page, click **Adatum Web Service**, and then click **OK**.
5. In the Service section, click **Configure**.
6. On the **Properties of Adatum Web Service** page, in the **Service Name** field, type **Contractor Service**, in the **Cost center** field, type **London Development**, and then click **OK**.
7. In the Instance section, click **Configure**.
8. In the **Computer Name** field, type **LON-WEB1**, and then click **OK**.
9. Note that, by clicking **Deploy**, a new virtual machine would be deployed. This would take between 10 and 20 minutes.
10. Click **Cancel** to cancel the deployment.
11. Close the App Controller window.

Exercise 4: Implementing Host Updating in VMM

► Task 1: Configure VMM integration with WSUS

1. On LON-VMM1, switch to the Virtual Machine Manager console, where you are signed in as Administrator, and then click the **Fabric** workspace.
2. In the navigation pane, expand the **Servers** node, expand **Infrastructure**, and then click **Update Server**.
3. Right-click **Update Server**, and then click **Add Update Server**. The **Add Windows Server Update Services Server** dialog box opens.
4. In the **Add Windows Server Update Services Server** dialog box, in the **Computer name** field, type **LON-WSUS**, and then in the **TCP/IP port** field, type **8530**.
5. Click **Enter a user name and password**. In the **User name** field, type **Adatum\Administrator**, in the **Password** field, type **Pa\$\$w0rd**, and then click **Add**. The Jobs window opens.
6. In the Jobs window, select the **Add Update Server** job. On the **Summary** and **Details** tabs, monitor the status of the configuration job.
7. When the job displays as **Completed w/info**, close the Jobs window. Note: Status is expected to be Completed w/Info.
8. With the Update Server node selected, verify that **LON-WSUS.adatum.com** displays in the results pane and that the **Agent Status** column displays **Responding**.

► Task 2: Configure a software update baseline in VMM

1. On LON-VMM1, in the Virtual Machine Manager console, click the **Library** workspace.
2. In the navigation pane, expand **Update Catalog and Baselines**, and then click **Update Catalog**.
3. In the results pane, verify that various software updates display. These updates have been synchronized from the Windows Server Update Services (WSUS) server role. If you do not see any software updates, right-click **Update Catalog**, and then click **Synchronize Update Server**. When synchronization completes, close the Jobs window.
4. In the ribbon, click **Create**, and then click **Baseline**. The Update Baseline Wizard starts.
5. In the Update Baseline Wizard, on the **General** page, in the **Name** field, type **Server Baseline**, and then click **Next**.
6. On the **Updates** page, click **Add**.
7. In the **Add Updates to Baseline** dialog box, click the following update:
 - **Update for Windows Server 2012 R2 (KB2883200)**
8. Click **Add**, and then click **Next**.
9. On the **Assignment Scope** page, select the check boxes for the following items, and then click **Next**:
 - Library Servers: **LON-VMM1.Adatum.com**
 - Update Server: **LON-WSUS.Adatum.com**
 - VMM Server: **LON-VMM1.Adatum.com**
10. On the **Summary** page, click **Finish**.
11. In the Jobs window, verify that **Create new baseline** has completed successfully.
12. Close the Jobs window.

► **Task 3: Verify baseline compliance**

1. On LON-VMM1, click the **Fabric** workspace.
2. In the navigation pane, expand **Servers**, expand **Infrastructure**, and then click **Library Servers**.
3. On the ribbon, click **Compliance**.
4. In the results pane, note the compliance and operational status of **LON-VMM1.Adatum.com**. Compliance Status should display as **Unknown**, and Operational Status should display as **Pending Compliance Scan**.
5. Select **LON-VMM1.Adatum.com**, and then on the ribbon, click **Scan**. The Operational Status column changes to **Scanning**. After a couple minutes, Compliance Status should report as **Compliant**. This indicates that lon-vmm1.adatum.com is compliant with the baseline that you configured in the previous task. Note: If the Compliance status is reported as **Non-Compliant**, then right-click **LON-VMM1.Adatum.com** and select **Remediate**, in the Update Remediation window. Select the **Do not restart servers after remediation** check box, and then click **Remediate**. After a couple minutes, Compliance Status should change to **Compliant**.
6. Close the Virtual Machine Manager console.

► **Task 4: To prepare for the next module**

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

1. On the host computer, start Hyper-V® Manager.
2. On the Virtual Machines list, right-click **20414C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps two and three for 20414C-LON-VMM1, 20414C-LON-SVR1, 20414C-LON-SVR2, 20414C-LON-WSUS, and 20414C-TOR-SVR1.

Results: After completing this exercise, you will have configured VMM host updating.


Module 6: Planning and Implementing a Server Monitoring Strategy

Lab: Implementing a Server Monitoring Strategy

Exercise 1: Configuring Server Monitoring Using Windows Server 2012

► Task 1: Configure Server Manager to monitor multiple servers

1. On TOR-SVR1, in the Windows Server® 2012 Server Manager, on the Dashboard, click **Add other servers to manage**.
2. In the **Add Servers** dialog box, in the **Name (CN)** box, type **TOR-SS1**, and then click **Find Now**.
3. In the search results, click **TOR-SS1**, click on the arrow pointing to the right to add the server to the **Selected** box, and then click **OK**.
4. On the **Dashboard**, click **Create a server group**.
5. In the **Create Server Group** dialog box, in the **Server group name** box, type **Toronto Servers**.
6. In the list of servers, click **TOR-SVR1.Adatum.com**, then press the Ctrl key, and then click **TOR-SS1.Adatum.com**.
7. Click the arrow pointing to the right to add the servers to the **Selected** box, and then click **OK**.
8. In Server Manager, click **Toronto Servers**.
9. In the **SERVERS** box, click **TOR-SS1**.
10. In Server Manager, scroll down to **BEST PRACTICES ANALYZER**, to the right of BEST PRACTICES ANALYZER, click **TASKS**, and then click **Start BPA Scan**.
11. In the **Select Servers** dialog box, click **Start Scan**. It may take 2 to 5 minutes for the scan to finish.
12. In Server Manager, scroll down to the **PERFORMANCE** box, to the right of PERFORMANCE, click **TASKS**, and then click **Configure Performance Alerts**.
13. In the **Toronto Servers: Configure Performance Alerts** dialog box, in the **CPU (% usage)** box, type **75**.
14. In the **Memory (MB available)** box, type **100**, and then click **Save**.
15. In the list of servers under the performance graph, right-click **TOR-SS1**, and then click **Start Performance Counters**.

 **Note:** It may take 30 minutes or more for the data to start to display. Move on to the next exercise. At the end of the lab, you will check the data.

16. In Server Manager, scroll down to **ROLES AND FEATURES**, to the right of ROLES AND FEATURES, click **TASKS**, and then click **Add Roles and Features**.
17. In the **Add Roles and Features** wizard, on the **Before you begin** page, click **Next**.
18. On the **Select installation type** page, ensure that **Role-based or feature-based installation** is selected, and then click **Next**.
19. On the **Select destination server** page, click **TOR-SS1.Adatum.com**, and then click **Next**.

20. On the **Select server roles** page, click **Web Server (IIS)**, in the **Add Roles and Features Wizard** dialog box, click **Add Features**, and then click **Next**.
21. On the **Select features** page, click **Next**.
22. On the **Web Server Role (IIS)** page, click **Next**.
23. On the **Select role services** page, click **Next**.
24. On the **Confirmation installation selections** page, click **Install**.
25. On the **Results** page, wait for the successful completion of the **Feature installation** task, and then click **Close**.

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

1. On TOR-SVR1, in Server Manager, click **Toronto Servers**.
2. In the **SERVERS** box, right-click **TOR-SVR1**, and in the drop-down list box, click **Computer Management**.
3. In the Computer Management window, in the left pane, expand **Performance**, and then expand **Data Collector Sets**.
4. Under **Data Collector Sets**, expand **User Defined**, and then click **Server Manager Performance Monitor**.
5. Click **User Defined**.
6. In the Actions pane, click **More Actions**, in the drop-down list box, click **New**, and then click **Data Collector Set**.
7. In the **Create new Data Collector Set**, in the **Name** box, type **Main Resources**, and then click **Next**.
8. On the **Which template would you like to use?** page, click **System Performance**, and then click **Next**.
9. On the **Where would you like the data to be saved?** page, click **Next**.
10. On the **Create the data collector set** page, click **Finish**.
11. Right-click **Main Resources**, and then click **Properties**.
12. Click the **Schedule** tab, and then click **Add**.
13. In the **Folder Action** dialog box, in the **Start time** box, type **8:00:00 AM**.
14. Ensure only the following days are enabled, and then click **OK**:
15. **Monday**
16. **Tuesday**
- **Wednesday**
- **Thursday**
- **Friday**
17. Click the **Stop Condition** tab, and then, in the **Overall duration** box, type **10**.
18. To the right of **Overall duration**, in the **Units** drop-down list box, click **Hours**, and then click **OK**.
19. Click on **Main Resources**, then double-click **Performance Counter**.

20. In the **Performance Counter Properties** dialog box, review the counters being collected, and then, in the **Sample interval** box, verify that the interval is set to **1**.
21. In the **Units** drop-down list box, click **Minutes**, and then click **OK**.
22. Right-click **NT Kernel**, click **Delete** and then click **Yes**.
23. Under **User Defined**, right-click **Main Resources**, and then click **Start**.
24. Wait 10 minutes, then right-click **Main Resources**, and then in the drop-down list box, click **Stop**.
25. Right-click **Main Resources**, and then in the drop-down list box, click **Latest Report**.



Note: If the data collector is still running, you will see a report status message saying **Collecting data for 3600 seconds**.

26. In the **System Performance Report**, click **CPU**, click **Process**, and then view the counters collected.
27. Examine the report, and then close Computer Management.

► Task 3: Configure an event subscription

1. On TOR-SS1, while in the Desktop, click the Start screen icon, and then click **Administrative Tools**.
2. In Administrative Tools, double-click Windows Firewall with Advanced Security.
3. In Windows Firewall with Advanced Security, select **Inbound** rules, and then right-click **COM+ Network Access (DCOM-In)** and select **Enable Rule**.
4. Scroll down until you reach Remote Event Log Management items. Select all three by holding down the Shift key while clicking on the top Remote Event Log item and then the bottom Remote Event Log item (you should have three selected). Right-click and then select Enable Rule.
5. In the File menu of Windows® Firewall with Advanced Security and Administrative Tools, click **Exit**.
6. On TOR-SVR1, click the Start screen icon, and then click **Administrative Tools**.
7. In the Administrative Tools window, double-click **Event Viewer**.
8. In the Event Viewer window, click **Subscriptions**.
9. In the **Event Viewer** dialog box, click **Yes** to start the Event Collector Service.
10. In the Actions pane, click **Create Subscription**.
11. In the **Subscription Properties** dialog box, in the **Subscription name** box, type **TOR-SS1 Events**, and then click **Select Computers**.
12. In the **Computers** dialog box, click **Add Domain Computers**.
13. In the **Select Computer** dialog box, in the **Enter object name to select** box, type **TOR-SS1**, click **Check Names**, and then click **OK** twice.
14. Click **Select Events**.
15. In the **Query Filter** dialog box, click both **Critical** and **Error**.
16. In the **Event logs** drop-down list box, enable the check box to the left of **Windows Logs**, then click outside the drop-down list, and then click **OK**.
17. Click **OK** again.

18. In Server Manager, select **TOR-SS1** to manage the server remotely. Right-click **TOR-SS1**, and then in the drop-down list box, click **Computer Management**.
19. In Computer Management, expand **Local Users and Groups**, and then click **Groups**.
20. In the list of groups, right-click **Event Log Readers**, and then click **Add to Group**.
21. In the **Event Log Readers Properties** dialog box, click **Add**.
22. In the **Select Users, Computers, Service Accounts, or Groups** dialog box, click **Object Types**.
23. In the **Object Types** dialog box, click **Computers**, and then click **OK**.
24. In the **Enter the object names to select** box, type **TOR-SVR1**, click **Check Names**, and then click **OK**.
25. In the **Event Log Readers Properties** dialog box, click **OK**.
26. Switch back to Event Viewer.
27. Expand **Windows Logs** and then select **Forwarded Events**. It may take several minutes for events to be forwarded.

Results: After completing this exercise, you should have configured monitoring for the servers by using the tools available within the Windows Server 2012 operating system.


Exercise 2: Implementing Operations Manager Monitoring

► Task 1: Deploy the Operations Manager agent to virtual machines

1. On LON-OM1, on the task bar, click **Operations Console**.
2. In the Operations console, on the bottom left pane, click **Administration**.
3. In the Administration workspace, in the middle console tree pane, click **Discovery Wizard**.
4. In the Computer and Device Management Wizard, on the **Discovery Type** page, ensure that **Windows computers** is selected, and then click **Next**.
5. On the **Auto or Advanced?** page, ensure that **Advanced discovery** is selected, and then click **Next**.
6. On the **Discovery Method** page, click **Browse for, or type-in computer names**, in the box below, type **LON-SVR1, LON-SVR2**, and then click **Next**.
7. On the **Administrator Account** page, click **Discover**.
8. On the **Select Objects to Manage** page, click **Select All**, and then click **Next**.
9. On the **Summary** page, click **Finish**.
10. In the Agent Management Task Status window, wait for the task to complete, which may take several minutes. Close the window.
11. Switch to LON-SVR1.
12. In the Start screen, click **Administrative Tools**.
13. In the Administrative Tools window, double-click **Services**.
14. In the Services window, verify that the **Microsoft Monitoring Agent** service has started.


► Task 2: Configure agentless monitoring for host machines

1. On LON-OM1, in the Operations console, click **Administration**.
2. In the Administration workspace, in the Administration pane, click **Settings**, and then double-click **Security**.
3. In the **Global Management Server Settings – Security** dialog box, click **Review new manual agent installations in pending management view**, and then click **OK**.
4. Switch to LON-DC1.
5. Right-click the Start screen icon.
6. Click **Run**.
7. In the **Run** dialog box, in the **Open** box, type **\\LON-OM1\c\$**, and then click **OK**.
8. In the **c\$** window, double-click, in the following order: **Program Files, Microsoft System Center 2012 R2, Operations Manager, Server, AgentManagement**, and then **amd64**.
9. Double-click **MOMAgent.msi**.

 **Note:** If a security warning dialog box is displayed with the title **Open File – Security Warning**, click **Run**.

10. In the **Microsoft Monitoring Agent Setup** dialog box, click **Next**.
11. On the **Important Notice** page, click **I Agree**.

12. On the **Destination Folder** page, click **Next**.
13. On the **Agent Setup Options** page, click **Next**.
14. On the **Management Group Configuration** page, in the **Management Group Name** box, type **Adatum**. In the **Management Server** box, type **LON-OM1**, and then click **Next**.
15. On the **Agent Action Account** page, ensure that **Local System** is selected, and then click **Next**.
16. On the **Microsoft Update** page click **Next**.
17. On the **Ready to Install** page, click **Install**. Wait for the successful completion of the task, and then click **Finish**.
18. In Server Manager, click **Local Server**.
19. Scroll down to **SERVICES**, and locate the **Microsoft Monitoring Agent** service.

 **Note:** You may need to refresh the view to see the service, and you may need to order the list by service name to locate it more easily. If you do not see the service after a refresh, click on TOOLS in Server Manager, and then click Services. Locate the service in the Services management console.

20. Switch to LON-OM1.
21. In the Operations console, in the Administration workspace, click **Pending Management**.
22. In the list of pending manual agent installs, click **LON-DC1.Adatum.com**.
23. In the Tasks pane, click **Approve**.
24. In the **Manual Agent Install** dialog box, click **Approve**.
25. In the Administration pane, click **Agent Managed**. Notice that LON-DC1 now displays.

 **Note:** It may take 5 to 10 minutes for the status of LON-DC1 to update from Not monitored.

Results: After completing this exercise, you should have installed and verified the Operations Manager agent on computers in the London data center.

Exercise 3: Configuring the Operations Manager Monitoring Components

► Task 1: Install and configure management packs

1. On LON-OM1, in the Operations console, in the left pane, click **Administration**.
2. In the Administration workspace, under the Administration node, click **Management Packs**.
3. In the Tasks pane, click **Import Management Packs**.
4. In the **Import Management Packs** dialog box, click **Add**, and then, in the **Add** drop-down list box, click **Add from disk**.
5. In the **Online Catalog Connection** dialog box, click **No**.
6. In the **Select Management Packs to import** dialog box, expand drive **C**, expand **Program Files (x86)**, expand **System Center Management Packs**, expand **System Center Monitoring Pack for SQL Server**, select all SQL Server management pack files, and then click **Open**.
7. On the **Select Management Packs** page, click **Install**. Wait for the management packs to be imported, and then click **Close**.
8. Scroll down the list of management packs to locate the new management packs.
9. On the bottom left pane, click **Monitoring**.
10. In the Monitoring workspace, expand **Microsoft SQL Server**, and then click **Computers**. You may need to wait a few seconds before LON-OM1.Adatum.com displays.
11. In the Operations console, on the bottom left pane, click **Administration**.
12. In the Administration workspace, in the Administration pane, click **Management Packs**.
13. In the Tasks pane, click **Create Management Pack**.
14. In the **Create Management Pack** dialog box, on the **General Properties** page, in the **Name** box, type **SQL Server 2008 (Monitoring) – Overrides**, and then click **Next**.
15. On the **Knowledge** page, click **Create**.
16. In the Operations console, on the bottom left pane, click **Authoring** to open the Authoring workspace.
17. In the Authoring pane, expand **Management Pack Objects**, and then double-click **Monitors**.
18. On the yellow bar across the top of the Monitors pane, click **Change Scope**.
19. In the **Scope Management Pack Objects** dialog box, click **Clear All**, click **View all targets**, and then in the **Look for** box, type **SQL Server 2008**.
20. In the list of targets, click **SQL Server 2008 DB File**, and then click **OK**.
21. Expand **SQL Server 2008 DB File**, expand **Entity Health**, expand **Performance**, and then click **DB File Space**.
22. In the Tasks pane, click **Overrides**, in the drop-down list box, click **Override the Monitor**, and then click **For all objects of class: SQL Server 2008 DB File**.
23. In the **Overrides Properties** dialog box, select the **Lower Threshold** check box, and then change the **Override Value** to **20**.
24. Click the **Upper Threshold** check box, and then change the **Override Value** to **30**.
25. In the **Select destination management pack** drop-down list box, click **SQL Server 2008 (Monitoring) – Overrides**, and then click **OK**.

► Task 2: Configure notifications, subscribers, and subscriptions

Configure Notifications

1. In the Operations console, on the bottom left pane, click **Administration**.
2. In the Administration pane, under the Notifications node, click **Channels**.
3. In the Tasks pane, click **New**, and then in the drop-down list box, click **E-Mail (SMTP)**.
4. In the **E-Mail Notification Channel** dialog box, on the **Description** page, in the **Channel name** box, type **SMTP Notification Channel**, and then click **Next**.
5. On the **Settings** page, click **Add**.
6. In the **Add SMTP Server** dialog box, in the **SMTP server (FQDN)** box, type **LON-SVR1.Adatum.com**, and then click **OK**.
7. In the **Return address** box, type **om@datum.com**, and then click **Next**.
8. On the **Format** page, note the variables used in the subject and message for the email, and then click **Finish**.
9. Wait for the task to complete successfully, and then click **Close**.

Configure Subscribers

1. In the Administration pane, click **Subscribers**.
2. In the Tasks pane, click **New**.
3. In the **Notification Subscriber Wizard** dialog box, on the **Description** page, in the **Subscriber Name**, type **ADATUM\Administrator**, and then click **Next**.
4. On the **Schedule** page, click **Notify only during the specified time**, and then click **Add**.
5. In the **Specify Schedule** dialog box, under **Weekly recurrence**, click **From**, and then set the first time to **8:00 AM** and the second time to **8:00 PM**.
6. Under **On the selected days of the week**, click **Monday, Tuesday, Wednesday, Thursday, Friday, and Saturday**, and then click **OK**.
7. On the **Schedule** page, click **Next**.
8. On the **Addresses** page, click **Add**.
9. In the **Describe the Subscriber Address** dialog box, in the **Address name** box, type **ADATUM\Administrator**, and then click **Next**.
10. In the **Channel Type** drop-down list box, click **E-Mail (SMTP)**.
11. In the **Delivery address for selected channel** box, type **administrator@datum.com**, and then click **Next**.
12. On the **Schedule** page, click **Finish**.
13. On the **Addresses** page, click **Finish**, and then click **Close**.

Configure Subscriptions

1. In the Operations console, on the bottom left pane, click **Administration**.
2. In the Administration pane, click **Subscriptions**.
3. In the Tasks pane, click **New**.
4. In the **Create Notification Subscription** dialog box, in the **Subscription name** box, type **Critical SQL Alerts**, and then click **Next**.

5. On the **Criteria** page, in the **Conditions** list, click both **raised by any instance in a specific group**, and **of a specific severity**.
6. In the **Criteria description (click the underlined value to edit)** box, click the first occurrence of **specific**.
7. In the **Group Search** dialog box, in the **Filter by (optional)** box, type **SQL**, and then click **Search**.
8. In the **Available groups** list, click **SQL Server 2008 Computers**, click **Add**, and then click **OK**.
9. In the **Criteria description (click the underlined value to edit)** box, click the second occurrence of **specific**.
10. In the **Alert Type** dialog box, click **Critical**, and then click **OK**.
11. On the **Criteria** page, click **Next**.
12. On the **Subscribers** page, click **Add**.
13. In the **Subscriber Search** dialog box, click **Search**, click **ADATUM\Administrator**, click **Add**, and then click **OK**.
14. On the **Subscribers** page, click **Next**.
15. On the **Channels** page, click **Add**.
16. In the **Channel Search** dialog box, click **Search**, click **SMTP Notification Channel**, click **Add**, and then click **OK**.
17. On the **Channels** page, click **Next**.
18. On the **Summary** page, click **Finish**, and then click **Close**.

► Task 3: Configure reports

1. Switch to **TOR-SVR1**. In Server Manager, click **Toronto Servers**.
2. In the **SERVERS** box, click **TOR-SS1**.
3. Scroll down to the **PERFORMANCE** box and then view the performance data on the chart.



Note: It may take 30 minutes or more for data to be displayed. If data is not yet displayed at this point, leave the virtual machines running during lecture period and check again later.

► Task 4: To prepare for the next module

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

1. On the host computer, start Microsoft Hyper-V® Manager.
2. In the **Virtual Machines** list, right-click **20414C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20414C-LON-SVR1**, **20414C-LON-SVR2**, **20414C-LON-OM1**, **20414C-TOR-SS1**, and **20414C-TOR-SVR1**.

Results: After completing this exercise, you should have imported and configured the Microsoft SQL Server® Management pack in Operations Manager.

Module 7: Planning and Implementing High Availability for File Services and Applications

Lab: Planning and Implementing High Availability for File Services and Applications

Exercise 1: Planning a High Availability Strategy for File Services

► Task 1: Read the supporting documentation

- Read the documentation provided in the Student Workbook.

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

- Answer the questions in the proposals section of the Supporting Documentation.

Proposals

1. How should you allow access to the Internet small computer system interface (iSCSI) storage area network (SAN) in case of network switch failure?

Answer: You can use Multipath I/O (MPIO).

2. How should you allow access to the iSCSI SAN in case of network interface card failure at the server level?

Answer: You can use NIC teaming.

3. How should you configure your storage to allow access to data even if a physical disk fails?

Answer: You can use Storage Spaces with parity or mirror.

4. How should you configure your solution to allow users to access a file share in London when the Toronto servers are offline?

Answer: You can use Distributed File System (DFS) Replication and DFS namespace.

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

- Compare the proposals with the ones in the previous lab task.

Results: After completing this exercise, you should have planned a high availability strategy for file servers.

Students will design the storage space deployment. The design should include storage units and those servers and network components that will be deployed. The design also should contain notes about highly available features that the deployment will include, such as NIC teaming on the servers, MPIO for the network path, and RAID levels on the storage enclosure.

Students will design the DFS deployment. The design will include server locations, the deployed DFS namespaces, and DFS targets. The design should include notes on the configuration for the DFS Replication configuration.

Exercise 2: Planning a High Availability Strategy for Web Applications

► Task 1: Read the supporting documentation

- Read the documentation provided in the Student Workbook.

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

- Answer the questions in the Proposals section in the Student Handbook.

Proposals

1. How should you provide high availability for the web application?

Answer: You should use Network Load Balancing (NLB) between LON-SVR1 and LON-SVR2.

2. How should you manage session maintenance?

Answer: You should use client affinity.

3. How should you configure you servers to allow changes to be copied from TOR-SVR1 to the other servers?

Answer: You should use DFS Replication with TOR-SVR1 as a hub, and LON-SVR1 and LON-SVR2 as spokes.

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

- Compare your proposals with the answers in the previous section.

Results: After completing this exercise, you should have planned a high availability strategy for web applications that meets these criteria:

Your NLB deployment design should include the server deployment and the storage design. You should use the information from the previous exercise to plan the storage design.

Your design should include the NLB port rules and network settings design.

Exercise 3: Implementing a High Availability Solution for File Storage

► Task 1: Configure NIC Teaming

1. Sign in to TOR-SVR1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. In Server Manager, click **Local Server**.
3. In the Properties pane for TOR-SVR1, to the right of **Ethernet**, click the listed IP Address.
4. In the **Network Connections** dialog box, select and right-click **Ethernet 2**, and then click **Enable**. Repeat this step, enabling **Ethernet 3** and **Ethernet 4**.
5. Close the **Network Connections** dialog box.
6. Refresh the Properties pane for TOR-SVR1, and then to the right of **NIC Teaming**, click **Disabled**.
7. In the NIC Teaming window, in the Adapters and Interfaces pane, click **Ethernet**. Press and hold the **Ctrl** key, and then click **Ethernet 2**.
8. In the NIC Teaming window, in the Adapters and Interfaces pane, click **Tasks**. In the drop-down list box, click **Add to New Team**.
9. In the **Team name** field, type **iSCSI Access Team 1**, and then click **OK**.
10. Close the NIC Teaming window.
11. Click to the Start screen, click the **Control Panel** tile, click **Network and Internet**, and then click **Network and Sharing Center**.
12. In the Network and Sharing Center, click **Change adapter settings**.
13. In the Network Connections window, right-click **iSCSI Access Team 1**, and then click **Properties**.
14. In the **iSCSI Access Team 1 Properties** dialog box, double-click **Internet Protocol Version 4 (TCP/IPv4)**.
15. Click **Use the following IP address**, and then enter the following parameters:
 - IP address: **172.16.1.20**
 - Subnet mask: **255.255.0.0**
 - Default gateway: **172.16.0.1**
16. Click **Use the following DNS server addresses**, and then in the **Preferred DNS server** field, type **172.16.0.10**.
17. Click **OK**, and then click **OK** to close the **iSCSI Access Team 1 Properties** dialog box.
18. Close the Network Connections window.
19. Close the Network and Sharing Center window.
20. In Server Manager, click **Local Server**.
21. In the Properties for TOR-SVR1 pane, to the right of **NIC Teaming**, click **Enabled**.
22. In the NIC Teaming window, in the Adapters and Interfaces pane, click **Ethernet 3**. Press and hold the **Ctrl** key, and then click **Ethernet 4**.
23. In the NIC Teaming window, in the Adapters and Interfaces pane, click **Tasks**. In the drop-down list box, click **Add to New Team**.
24. In the **Team name** dialog box, type **iSCSI Access Team 2**, and then click **OK**.
25. Close the NIC Teaming window.

26. On the Start screen, click the **Control Panel** tile, click **Network and Internet**, and then click **Network and Sharing Center**.
27. In the Network and Sharing Center, click **Change adapter settings**.
28. Right-click **iSCSI Access Team 2**, and then click **Properties**.
29. In the **iSCSI Access Team 2 Properties** dialog box, double-click **Internet Protocol Version 4(TCP/IPv4)**.
30. Click **Use the following IP address**, and then enter the following parameters:
 - IP address: **131.107.1.10**
 - Subnet mask: **255.255.0.0**
 - Default gateway: Leave blank
31. Click **Use the following DNS server addresses**, and then next to **Preferred DNS server**, type **172.16.0.10**.
32. Click **OK**, and then click **OK** to close the **iSCSI Access Team 2 Properties** dialog box.
33. Close the Network Connections window.
34. Close the Network and Sharing Center window.

► **Task 2: Configure iSCSI initiators and MPIO**

1. Sign in to TOR-SVR1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**. If Server Manager is not open, click the **Server Manager** icon on the task bar.
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 a server from the server pool** is selected, and then click **Next**.
6. On the **Select server roles** page, click **Next**.
7. On the **Select features** page, click **Multipath I/O**, and then click **Next**.
8. On the **Confirm installation selections** page, click **Install**.
9. When installation is complete, click **Close**.
10. In Server Manager, click **Tools**, and then in the Tools drop-down list box, click **iSCSI Initiator**.
11. In the **Microsoft iSCSI** dialog box, click **Yes**.
12. In the **iSCSI Initiator Properties** dialog box, on the **Targets** tab, in the **Target** field, type **172.16.1.25**, and then click **Quick Connect**. In the **Discovered targets** section, you should now see a discovered target with a status of **Connected**.
13. In the **Quick Connect** box, click **Done**.
14. Click **OK** to close the **iSCSI Initiator Properties** dialog box.
15. In Server Manager, click **Tools**, and then in the Tools drop-down list box, click **MPIO**.
16. In **MPIO Properties** dialog box, click the **Discover Multi-Paths** tab.
17. On the **Discover Multi-Paths** tab, select the **Add support for iSCSI devices** check box, and then click **Add**. When you are prompted to reboot the computer, click **Yes**.

18. After the computer restarts, sign in to TOR-SVR1 with the username **Adatum\Administrator** and the password **Pa\$\$w0rd**.
19. In Server Manager, click **Tools**, and then in the Tools drop-down list box, click **MPIO**.
20. In the **MPIO Properties** dialog box, on the **MPIO Devices** tab, notice that **Device Hardware ID MSFT2005iSCSIBusType_0x9** is added to the list.
21. Click **OK** to close the **MPIO Properties** dialog box.

► **Task 3: Configure Storage Spaces by using iSCSI targets**

1. On TOR-SVR1, in Server Manager, in the left pane, click **File and Storage Services**, and then in the Servers pane, click **Storage Pools**.
2. In the STORAGE POOLS pane, click **TASKS**, and then in the **TASKS** drop-down list box, click **New Storage Pool**.
3. In the New Storage Pool Wizard, on the **Before You Begin** page, click **Next**.
4. On the **Specify a storage pool name and subsystem** page, in the **Name** text box, type **iSCSIPool**, and then click **Next**.
5. On the **Select physical disks for the storage pool** page, click the following physical disks, and then click **Next**:
 - **PhysicalDisk1**
 - **PhysicalDisk2**
 - **PhysicalDisk3**
 - **PhysicalDisk4**
6. On the **Confirm selections** page, click **Create**.
7. On the **View results** page, wait until the task completes, and then click **Close**.
8. Select **iSCSIPool**, in the VIRTUAL DISKS pane, click **TASKS**, and then click **New Virtual Disk**.
9. In the New Virtual Disk Wizard, on the **Before You Begin** page, click **Next**.
10. On the **Select the storage pool** page, click **iSCSIPool**, and then click **Next**.
11. On the **Specify the virtual disk name** page, in the **Name** text box, type **DFSDisk**, and then click **Next**.
12. On the **Select the storage layout** page, in the **Layout** list, click **Parity**, and then click **Next**.
13. On the **Specify the provisioning type** page, click **Fixed**, and then click **Next**.
14. On the **Specify the size of the virtual disk** page, click **Maximum size**, and then click **Next**.
15. On the **Confirm selections** page, click **Create**.
16. On the **View results** page, wait until the task completes. Ensure that the **Create a volume when this wizard closes** check box is selected, and then click **Close**.
17. In the New Volume Wizard, on the **Before You Begin** page, click **Next**.
18. On the **Select the server and disk** page, in the Disk pane, click the **DFSDisk** virtual disk, and then click **Next**.
19. On the **Specify the size of the volume** page, click **Next** to confirm the default selection.
20. On the **Assign to a drive letter or folder** page, in the **Drive letter** drop-down list box, ensure that **H** is selected, and then click **Next**.

21. On the **Select file system settings** page, in the **File system** drop-down list box, click **NTFS**, in the **Volume label** box, type **DFS Volume**, and then click **Next**.
22. On the **Confirm selections** page, click **Create**.
23. On the **Completion** page, wait until the creation completes, and then click **Close**.

► **Task 4: Validate the high availability of the deployment against the loss of a single network adapter**

1. On TOR-SVR1, on the taskbar, click the **Windows PowerShell** icon.
2. In the Windows PowerShell® window, at the command prompt, type the following command, and then press Enter:

```
Copy C:\windows\system32\notepad.exe H:\
```

3. Close the Windows PowerShell window.
4. On the taskbar, click the **File Explorer** icon.
5. In File Explorer, click **DFSVolume (H:)**.
6. Verify that notepad.exe displays in the file list.
7. Close File Explorer.
8. On the host machine, in Hyper-V® Manager, in the Virtual Machines pane, right-click **20414C-TOR-SVR1**, and then click **Settings**.
9. In Settings for 20414C-TOR-SVR1, in the Hardware pane, click the first occurrence of **Network Adapter**. In the **Virtual Switch** drop-down list box, click **Not connected**, and then click **OK**.
10. On TOR-SVR1, open File Explorer, and then click **DFSVolume (H:)**.
11. Verify that notepad.exe displays in the file list, even though TOR-SVR1 is not connected to the network.
12. On the host machine, in Hyper-V Manager, in the Virtual Machines pane, right-click **20414C-TOR-SVR1**, and then click **Settings**.
13. In Settings for 20414C-TOR-SVR1, in the Hardware pane, click the first occurrence of **Network Adapter**. In the **Virtual Switch** drop-down list box, click **External Network**, and then click **OK**.

Results: After completing this exercise, you should have implemented a high availability solution for file storage.

Exercise 4: Implementing a High Availability Solution by Using NLB

► Task 1: Configure an NLB cluster

1. Sign in to LON-DC1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. Click to the Start screen, and then click the **Internet Explorer** icon.
3. In Windows® Internet Explorer®, in the Address bar, type the address **http://LON-SVR1.adatum.com**, and then press Enter. If a pop-up window displays, click **Yes**. Verify that the name of the server displays at the top of the page.
4. In the **First name** text box, type your name, and then click **OK**. Verify that a message displays with your name on it.
5. In Internet Explorer, click **Back**. Verify that the message **Hello <your name>** displays. If the message does not display, press F5 to refresh the page. This is because the web application maintains session state.
6. In Internet Explorer, open a new tab. In the address bar, type **http://LON-SVR2.adatum.com**, and then press Enter. Verify that the name of the server displays at the top of the page.
7. In Server Manager, click **Tools**, and then in the **Tools** drop-down list box, click **DNS**.
8. In DNS Manager, expand **LON-DC1**, and then expand **Forward Lookup Zones**.
9. Select and then right-click **Adatum.com**, and then click **New Host (A or AAAA)**.
10. In the **New Host** dialog box, in the **Name** text box, type **www**. In the **IP Address text** box, type **172.16.0.111**, click **Add Host**, click **OK**, and then click **Done**.
11. Switch to LON-SVR1.
12. In Server Manager, click the **Tools** menu, and then click **Windows PowerShell ISE**.
13. In the Windows PowerShell Integrated Scripting Environment (ISE) window, enter the following command, and then press Enter:

```
Invoke-Command -Computersname LON-SVR1,LON-SVR2 -command {Install-WindowsFeature NLB,RSAT-NLB}
```


14. In the Windows PowerShell ISE window, type the following command, and then press Enter:

```
New-NlbCluster -InterfaceName "Ethernet" -OperationMode Multicast -ClusterPrimaryIP 172.16.0.111 -ClusterName LON-NLB
```

15. In the Windows PowerShell ISE window, type the following command, and then press Enter:

```
Add-NlbClusterNode -InterfaceName "Ethernet" -NewNodeName "LON-SVR2" -NewNodeInterface "Ethernet"
```

16. In Server Manager, click **Tools**, in the **Tools** drop-down list box, click **Network Load Balancing Manager**, and then click **OK**.


 **Note:** NLB is not fully configured yet. Therefore, a warning message will display, which you can disregard.

17. In the Network Load Balancing Manager, right-click **LON-NLB**, and then click **Cluster Properties**.
18. In the **Cluster Properties** dialog box, on the **Port Rules** tab, click the **All** port rule, and then click **Remove**.


19. On the **Port Rules** tab, click **Add**.
20. In the **Add/Edit Port Rule** dialog box, enter the following information, and then click **OK**:
 - Port range: **80 to 80**
 - Protocols: **TCP**
 - Filtering mode: **Multiple Host**
 - Affinity: **Single**
21. Click **OK** to close the **Cluster Properties** dialog box.

► **Task 2: Configure the Web servers to use highly availability storage**

1. Sign in to LON-SVR1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. In Server Manager, click **Tools**, and then in the **Tools** drop-down list box, click **DFS Management**.
3. In the DFS Management console, expand the **Namespaces** node, and then click the **\\Adatum.com\Website** namespace.
4. On the middle pane, click the **Namespace Servers** tab. Notice that there are three servers in use:
 - LON-SVR1
 - LON-SVR2
 - TOR-SVR1
5. Expand the **Replication** node, and then click **adatum.com\website\wwwroot**.
6. On the middle pane, click the **Connections** tab. Notice that each server replicates to the other two servers, forming a mesh topology.
7. Close the DFS Management console.
8. In Server Manager, click **Tools**, and then in the **Tools** drop-down list box, click **Internet Information Services (IIS) Manager**.

 **Note:** If a dialog box with the text **Do you want to get started with the Microsoft Web Platform to stay connected with latest Web Platform Components?** displays, select the **Do not show this message** checkbox, and then click **Yes**.

9. In the Connections pane, expand **LON-SVR1**, expand **Sites**, and then click **Default Web Site**.
10. In the Actions pane, click **Basic Settings**.
11. In the **Edit Site** dialog box, in the **Physical path** box, type **\\adatum.com\website\wwwroot**.
12. Click **Connect as**.
13. In the **Connect as** dialog box, click **Specific user**, and then click **Set**.
14. In the **Set Credentials** dialog box, type **Adatum\administrator** as the **User name**, and **Pa\$\$w0rd** in the **Password** and **Confirm password** fields. Click **OK** three times, and then close Internet Information Services (IIS) Manager.
15. Switch to LON-SVR2, and repeat steps eight through 14.

 **Note:** Do not use the administrator account to connect to the shared folder in a production environment. You should use a normal user account with limited permissions.

► Task 3: Validate the deployment

1. Switch to LON-CL1.
2. Click the **Desktop** tile, and then on the taskbar, click the **Internet Explorer** icon.
3. In the Internet Explorer Address bar, type the address **http://www.adatum.com**, and then press Enter. Note the name of the server that displays at the top of the page.
4. In the **First name** text box, type your name, and then click **OK**. Verify that a message displays with your name in it.
5. Click the **Refresh** icon 20 times. Notice that the server name does not change, due to affinity.
6. Switch to LON-SVR1.
7. In the Server Manager console, click the **Tools** menu, and then click **Windows PowerShell ISE**.
8. In the Windows PowerShell ISE window, type the following command, and then press Enter:

```
Stop-NTbClusterNode servername
```

Where *servername* is the name of the server to which you connected in the previous steps.

9. Switch to LON-CL1.
10. Click the **Refresh** icon. Verify that you connect to a different server and lose the session state.
11. Switch to LON-SVR1.
12. In the Windows PowerShell ISE window, type the following command, and then press Enter:

```
Start-NTbClusterNode servername
```

Where *servername* is the same server name as you used in the previous command.

13. Switch to LON-CL1.
14. Click the **Refresh** icon. The server name should not change.

► Task 4: 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 Hyper-V Manager, on the **Virtual Machines** list, right-click **20414C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps two and three for 20414C-LON-SVR1, 20414C-LON-SVR2, 20414C-LON-CL1, 20414C-TOR-SS1, and 20414C-TOR-SVR1.

Results: After completing this exercise, you should have implemented a high availability solution by using NLB.

Module 8: Planning and Implementing a High Availability Infrastructure Using Failover Clustering

Lab: Planning and Implementing a Highly Available Infrastructure by Using Failover Clustering

Exercise 1: Designing Highly Available Server Roles

► Task 1: Analyze server roles that need to be highly available

- Analyze the list of server roles that the exercise scenario provides. Based on what you have learned so far, think about highly available solutions for these server roles.

► Task 2: Propose a highly available design

Propose a design solution for making roles from exercise scenario highly available, by answering following questions:

1. How will you make Active Directory® Domain Services (AD DS) domain controllers highly available, and which highly available technology will you use?

Answer: For AD DS, you should not implement any specific highly available technology. It is enough to implement multiple AD DS domain controllers. AD DS does not support failover clustering or Network Load Balancing (NLB).

2. What technology or technologies will you use to make web servers highly available?

Answer: Usually, for web servers we use NLB. However, if web servers connect to a database in the background, then we can also implement failover clustering to make databases highly available.

3. How can you make a Dynamic Host Configuration Protocol (DHCP) server highly available, and are there any alternative solutions?

Answer: DHCP service supports failover clustering, so you can use it to make a DHCP server highly available. However, in Windows Server® 2012, you can also use DHCP failover as a technology to provide clusterless implementation of highly available DHCP servers.

4. What is the recommended solution for making Microsoft® SQL Server® highly available?

Answer: The only recommended highly available solution for SQL Server is failover clustering.

5. How can you make your file server highly available?

Answer: File servers support clustering. However, depending on the purpose of your file server, you also can consider technologies like Distributed File System (DFS).

6. How can you make Exchange Server highly available? Does the same approach apply for all Exchange Server roles?

Answer: Exchange Server Mailbox server supports failover clustering but only through implementation of database availability groups (DAGs). For a Client Access server role, you should use Client Access server arrays and NLB.

► **Task 3: Examine the suggested proposals in the Lab Answer Key**

Read the answers in the Lab Answer Key section, and then compare them with your answers. Discuss with the class.

Results: After completing this exercise, you will have completed the design of high availability for various server roles.

Exercise 2: Deploying a Failover Cluster

► Task 1: Connect to iSCSI targets from both host machines

1. On LON-HOST1, open **Server Manager**, click **Tools**, and then click **iSCSI Initiator**. At the **Microsoft iSCSI** prompt, click **Yes**.
2. Click the **Discovery** tab, click **Discover Portal...**, in the **IP address or DNS name** box, type **172.16.0.10**, and then click **OK**.
3. Click the **Targets** tab, click **Refresh**, in the **Targets** list, select **iqn.1991-05.com.microsoft:lon-dc1-target1-target**, and then click **Connect**.
4. Select **Add this connection to the list of Favorite Targets**, and then click **OK**.
5. Click **OK** to close iSCSI Initiator Properties.
6. On LON-HOST2, open **Server Manager**, click **Tools**, and then click **iSCSI Initiator**.
7. In the **Microsoft iSCSI** dialog box, click **Yes**.
8. Click the **Discovery** tab.
9. Click **Discover Portal...**
10. In the **IP address or DNS name** box, type **172.16.0.10**, and then click **OK**.
11. Click the **Targets** tab.
12. Click **Refresh**.
13. In the **Discovered Targets** list, select **iqn.1991-05.com.microsoft:lon-dc1-target1-target**, and then click **Connect**.
14. Select **Add this connection to the list of Favorite Targets**, and then click **OK**. Click **OK** to close iSCSI Initiator Properties.
15. On LON-HOST2, in the Server Manager window, click **Tools**, and then click **Computer Management**.
16. Expand **Storage**, and then click **Disk Management**.
17. Right-click **Disk 2**, and then click **Online**.
18. Right-click **Disk 2**, and then click **Initialize Disk**. 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**. If a window appears with a prompt to format disk, click **Cancel**.
25. On LON-HOST1, in Server Manager, click **Tools**, and then click **Computer Management**.
26. Expand **Storage**, and then click **Disk Management**.
27. Right-click **Disk Management**, and then click **Refresh**.
28. Right-click **Disk 2**, and then click **Online**.

► **Task 2: Install and configure failover clustering on both host machines**

1. On LON-HOST1, on the taskbar, click the **Server Manager** icon to open **Server Manager**.
2. From the **Dashboard**, click **Add roles and features**.
3. On the **Before You Begin** page, click **Next**.
4. On the **Select installation type** page, click **Next**.
5. On the **Select destination server** page, make sure that **Select a 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**, click **Add Features**, and then click **Next**.
8. On the **Confirm installation selections** page, click **Install**.
9. When installation is complete, click **Close**.
10. Repeat steps one through nine on LON-HOST2.
11. On LON-HOST1, in the Server Manager console, click **Tools**, and then click **Failover Cluster Manager**.
12. In Failover Cluster Manager, in the center pane, under **Management**, click **Create Cluster**.
13. In the Create Cluster Wizard, on the **Before You Begin** page, read the information. Click **Next**.
14. In the **Enter server name** box, type **LON-HOST1**, and then click **Add**. Type **LON-HOST2**, and then click **Add**.
15. Verify the entries, and then click **Next**.
16. On the **Validation Warning** page, click **No, I do not require support from Microsoft for this cluster**, and then click **Next**.
17. In the **Access Point for Administering the Cluster** page, in the **Cluster Name** box, type **VMCluster**.
18. Under **Address**, type **172.16.0.126**, and then click **Next**.
19. In the **Confirmation** dialog box, verify the information, clear the check mark next to **Add all eligible storage to the cluster**, and then click **Next**.
20. In the **Create Cluster Wizard Summary** page, click **Finish**.

► **Task 3: Configure the cluster**

1. On LON-HOST1, in the Failover Cluster Manager console, expand **VMCluster.Adatum.com**, expand **Storage**, and then right-click **Disks**.
2. Click **Add Disk**.
3. In the **Add Disks to a Cluster** dialog box, verify that **Cluster Disk 1** is selected, and then click **OK**.
4. Verify that the disk appears available for cluster storage in **Failover Cluster Manager**.
5. Right-click **VMCluster.Adatum.com**, select **More Actions**, and then click **Configure Cluster Quorum Settings**. Click **Next**.
6. On the **Select Quorum Configuration Option** page, click **Use default quorum configuration**, and then click **Next**.
7. On the **Confirmation** page, click **Next**.
8. On the **Summary** page, click **Finish**.

► **Task 4: Validate the cluster**

1. On LON-HOST2, open **Server Manager**, click **Tools**, and then click **Failover Cluster Manager**.
2. Click **VMCluster.Adatum.com** in the left pane, and then in the **Actions** pane, click **Validate Cluster**.
3. In the Validate a Configuration Wizard, on the **Before You Begin** page, click **Next**.
4. On the **Testing Options** page, make sure that **Run all tests (recommended)** is selected, and then click **Next**.
5. On the **Review Storage Status** page, select **Cluster Disk 1**, and then click **Next**.
6. On the **Confirmation** page, click **Next**.
7. Wait until cluster testing finishes, and then click **View Report**. Review the report. Some warnings are expected, but you should not have any errors.
8. Close Internet Explorer®, and then click **Finish**.

Results: After completing this exercise, you will have deployed a failover cluster.

Exercise 3: Implementing a Scale-Out File Server

► Task 1: Install the file server role and failover clustering on LON-SVR1 and LON-SVR2

1. On LON-SVR1, in Server Manager, click **Dashboard**, and then click **Add roles and features**.
2. On the **Before You Begin** page, click **Next**.
3. On the **Select installation type** page, click **Next**.
4. On the **Select destination server** page, click **Next**.
5. On the **Select server roles** page, expand **File and Storage Services (Installed)**, expand **File and iSCSI Services (Installed)**, and then verify that **File Server** is selected. Click **Next**.
6. On the **Select features** page, select **Failover Clustering**, click **Add Features**, and then click **Next**.
7. On the **Confirm installation selections** page, click **Install**.
8. When the **Installation Succeeded** message appears, click **Close**.
9. Repeat steps one through eight on LON-SVR2.

► Task 2: Connect to the iSCSI target from both file server cluster nodes

1. On LON-SVR1, open **Server Manager**, click **Tools**, and then click **iSCSI Initiator**. At the **Microsoft iSCSI** prompt, click **Yes**.
2. Click the **Discovery** tab.
3. Click **Discover Portal**, in the **IP address or DNS name** box, type **172.16.0.10**, and then click **OK**.
4. Click the **Targets** tab.
5. In the **Targets** list, select **iqn.1991-05.com.microsoft:lcn-dc1-target1-target**, and then click **Connect**.
6. Select **Add this connection to the list of Favorite Targets**, and then click **OK**.
7. Click **OK** to close iSCSI Initiator Properties.
8. On LON-SVR2, open **Server Manager**, click **Tools**, and then click **iSCSI Initiator**.
9. In the **Microsoft iSCSI** dialog box, click **Yes**.
10. Click the **Discovery** tab.
11. Click **Discover Portal**, in the **IP address or DNS name** box, type **172.16.0.10**, and then click **OK**.
12. Click the **Targets** tab.
13. In the **Discovered Targets** list, select **iqn.1991-05.com.microsoft:lcn-dc1-target1-target**, and then click **Connect**.
14. Select **Add this connection to the list of Favorite Targets**, and then click **OK**. Click **OK** to close iSCSI Initiator Properties.
15. On LON-SVR2, in the Server Manager window, click **Tools**, and then click **Computer Management**.
16. Expand **Storage**, and then click **Disk Management**.
17. Right-click **Disk 2**, and then click **Online**. (Note: Make sure that you do not click Disk 1, as it is the disk in use in another cluster that you created previously).
18. Right-click **Disk 2**, and then click **Initialize Disk**. 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 **ClusterDisk1**. Select the **Perform a quick format** check box, and then click **Next**.
24. Click **Finish**. If the Microsoft Windows® prompt appears, click **Cancel**.
25. Right-click **Disk 3**, and then click **Online**.
26. Right-click **Disk 3**, and then click **Initialize Disk**. In the **Initialize Disk** dialog box, click **OK**.
27. Right-click the unallocated space next to **Disk 3**, and then click **New Simple Volume**.
28. On the **Welcome** page, click **Next**.
29. On the **Specify Volume Size** page, click **Next**.
30. On the **Assign Drive Letter or Path** page, click **Next**.
31. On the **Format Partition** page, in the **Volume label** box, type **Quorum**. Select the **Perform a quick format** check box, and then click **Next**.
32. Click **Finish**. If the Microsoft Windows prompt appears, click **Cancel**.
33. Close Computer Management.
34. On LON-SVR1, in Server Manager, click **Tools**, and then click **Computer Management**.
35. Expand **Storage**, and then click **Disk Management**.
36. Right-click **Disk Management**, and then click **Refresh**.
37. Right-click **Disk 3**, and then click **Online**. (Note: Make sure you do not click Disk 2)
38. Right-click **Disk 4**, and then click **Online**. (Note: LON-SVR1 already has two disks, so these additional disks are marked as Disk 3 and Disk 4, unlike LON-SVR2 where the same disks are marked as Disk 2 and Disk 3).
39. Close Computer Management.

► Task 3: Configure the Scale-Out File Server

1. On LON-SVR1, in the Server Manager console, click **Tools**, and then click **Failover Cluster Manager**.
2. In Failover Cluster Manager, in the center pane, under **Management**, click **Create Cluster**.
3. In the Create Cluster Wizard, on the **Before You Begin** page, read the information, and then click **Next**.
4. In the **Enter server name** box, type **LON-SVR1**, and then click **Add**. Type **LON-SVR2**, and then click **Add**.
5. Verify the entries, and then click **Next**.
6. On the **Validation Warning** page, click **No, I do not require support from Microsoft for this cluster**, and then click **Next**.
7. In the **Access Point for Administering the Cluster** page, in the **Cluster Name** box, type **FSCluster**.
8. Under **Address**, type **172.16.0.127**, and then click **Next**.
9. In the **Confirmation** dialog box, verify the information, clear the check mark next to **Add all eligible storage to the cluster**, and then click **Next**.
10. In the **Create Cluster Wizard Summary** page, click **Finish**.
11. In Failover Cluster Manager, expand **FSCluster.Adatum.com**, expand **Storage**, and then right-click **Disks**.
12. Click **Add Disk**.
13. In the **Add Disks to a Cluster** dialog box, verify that **Cluster Disk 1** and **Cluster Disk 2** are selected, and then click **OK**.

14. Verify that the disks appear available for cluster storage in **Failover Cluster Manager**.
15. Right-click **Cluster Disk 1**, and then select **Add to Cluster Shared Volumes**.
16. Right-click **FSCluster.Adatum.com**, select **More Actions**, and then click **Configure Cluster Quorum Settings**. Click **Next**.
17. On the **Select Quorum Configuration Option** page, click **Use default quorum configuration**, and then click **Next**.
18. On the **Confirmation** page, click **Next**.
19. On the **Summary** page, click **Finish**.
20. Right-click **Roles**, and then select **Configure Role**.
21. On the **Before You Begin** page, click **Next**.
22. On the **Select Role** page, select **File Server**, and then click **Next**.
23. On the **File Server Type** page, click **Scale-Out File Server for application data**, and then click **Next**.
24. On the **Client Access Point** page, in the **Name** box, type **AdatumFS**, and then click **Next**.
25. On the **Confirmation** page, click **Next**.
26. On the **Summary** page, click **Finish**.

► **Task 4: Create a continuously available file share**

1. On LON-SVR1, in Failover Cluster Manager, click **Roles**, and then in the central pane, right-click **AdatumFS**.
2. Select **Add File Share**. (If you receive a message that the Client Access Point is not ready, perform these steps on LON-SVR2.)
3. In the New Share Wizard, on the **Select the profile for this share** page, select **SMB Share-Applications**, and then click **Next**.
4. On the **Select the server and path for this share** page, click **Select by volume**, and then click **Next**.
5. On the **Specify share name** page, in the **Share name** text box, type **VMachines**, and then click **Next**.
6. On the **Configure share settings** page, click **Next**.
7. On the **Specify permissions to control access** page, click **Customize permissions**.
8. In the Advanced Security Settings for VMachines window, click **Add**.
9. In the Permission Entry for VMachines window, click **Select a principal**.
10. In the Select User, Computer, Service Account, or Group window, click **Object Types**. In the Object Types window, select **Computers**, and then click **OK**.
11. In the text box, type **LON-HOST1**. Click **Check Names**, and then click **OK**.
12. In the **Basic Permissions** section, select **Full Control**, and then click **OK**.
13. Repeat steps eight through 12 to assign **Full control permissions** to LON-HOST2.
14. Click **OK** to close Advanced Security Settings for VMachines.
15. On the **Specify permissions to control access** page, make sure that both **ADATUM\LON-HOST1\$** and **ADATUM\LON-HOST2\$** are present in the list with Full Control Access, and then click **Next**.
16. On the **Confirm selections** page, click **Create**, and then click **Close**.

Results: After completing this exercise, you will have configured a highly available file server.

Exercise 4: Configuring CAU

► Task 1: Configure CAU

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 the **Add features that are required for Failover Clustering?** dialog box, click **Add Features**. Click **Next**.
7. On the **Confirm installation selections** page, click **Install**.
8. When installation is complete, click **Close**.
9. Switch to LON-SVR1. Open **Server Manager**, click **Tools**, and then click **Windows Firewall with Advanced Security**.
10. In the Windows Firewall with Advanced Security window, click **Inbound Rules**.
11. In the rules list, find the rule **Inbound Rule for Remote Shutdown (RPC-EP-In)**. Right-click the rule, and then select **Enable Rule**.
12. In the rules list, find the rule **Inbound Rule for Remote Shutdown (TCP-In)**. Right-click the rule, and then select **Enable Rule**.
13. Close the Windows Firewall with Advanced Security window.
14. Switch to LON-SVR2, and repeat steps nine through 13.
15. On LON-DC1, in the **Server Manager** dashboard, click **Tools**, and then click **Cluster-Aware Updating**.
16. In the Cluster-Aware Updating window, in the **Connect to a failover cluster** drop-down list, select **FSCLUSTER**. Click **Connect**.
17. In the Cluster Actions pane, click **Preview updates for this cluster**.
18. In the FSCLUSTER-Preview Updates window, click **Generate Update Preview List**. After several minutes, updates will appear in the list. Review updates and then click **Close**.

► Task 2: Update the failover cluster

1. On LON-DC1, in the Cluster-Aware Updating console, click **Apply updates to this cluster**.
2. On the **Getting Started** page, click **Next**.
3. On the **Advanced Options** page, review the options for updating, and then click **Next**.
4. On the **Additional Update Options** page, click **Next**.
5. On the **Confirmation** page, click **Update**, and then click **Close**.
6. In the Cluster nodes pane, you can review the progress of updating.



Note: Remember that one node of the cluster is in a waiting state and the other node is restarting after it updates.

7. Wait until the process is finished.



Note: This may require a restart of both the nodes.

The process is finished when both nodes show **Succeeded** in the **Last Run status** column.

8. Sign in to LON-SVR1 with the username **Adatum\Administrator** and the password **Pa\$\$w0rd**.
9. On LON-SVR1, in Server Manager, click **Tools**, and then click **Cluster-Aware Updating**.
10. In the **Cluster-Aware Updating** dialog box, in the **Connect to a failover cluster** drop-down list box, select **FSCLUSTER**. Click **Connect**.
11. Click **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 completing this exercise, you will have configured Cluster-Aware Updating (CAU).

Exercise 5: Implementing Highly Available Virtual Machines

► Task 1: Move a .vhd file to the highly available storage

1. On the LON-HOST1, open **File Explorer**.
2. Browse to the **E:\Program Files\Microsoft Learning\20414\Drives\20414C-LON-CORE\Virtual Hard Disks** folder. Note that your drive letter may differ based upon your host machine configuration.
3. Right-click **20414C-LON-CORE.vhd**, and then click **Copy**.
4. In the File Explorer title bar, type **\\AdatumFS\VMachines**, and then press Enter.
5. Right-click the empty space in the folder, select **New**, click **Folder**, type **LON-CORE**, and then press Enter.
6. Double-click the **LON-CORE** folder, right-click inside the folder, and then click **Paste**.
7. Wait for a few minutes until the .vhd file copies to the highly available file share.
8. Close File Explorer.

► Task 2: Configure the Hyper-V nodes to use the scale-out file server cluster

1. On LON-HOST1, open **Hyper-V Manager**.
2. Right-click **LON-HOST1**, and then select **Hyper-V Settings**.
3. In the Hyper-V Settings for LON-HOST1 window, select **Virtual Hard Disks** in the left pane.
4. In the right pane, type **\\adatumfs\VMachines** in the text box.
5. Click **Virtual Machines** in the left pane.
6. In the right pane, type **\\adatumfs\VMachines** in the text box. Click **OK** to close the Hyper-V Settings window.
7. Repeat steps one through six on LON-HOST2.

► Task 3: Configure the virtual machine as highly available

1. On LON-HOST1, open the Failover Cluster Manager console.
2. Expand **VMCluster.Adatum.com**, and then right-click **Roles**.
3. Select **Virtual Machines**, and then click **New Virtual Machine**.
4. In the New Virtual Machine window, select **LON-HOST1**, and then click **OK**.
5. In the New Virtual Machine Wizard, on the **Before You Begin** page, click **Next**.
6. On the **Specify Name and Location** page, type **LON-CORE** for the name, and then click **Next**.
7. On the **Specify Generation** page, make sure that **Generation 1** is selected, and then click **Next**.
8. On the **Assign Memory** page, type **768** in the **Startup memory** box, and then click **Next**.
9. On the **Configure Networking** page, select **External Network** from the drop-down list box, and then click **Next**.
10. On the **Connect Virtual Hard Disk** page, click **Use an existing virtual hard disk**, and then click **Browse**. In the Open window, double-click the **LON-CORE** folder, click **20414C-LON-CORE.vhd**, and then click **Open**.
11. Click **Next**.

12. On the **Completing the New Virtual Machine Wizard** page, click **Finish**.
13. On the **Summary Page**, make sure that the status of the process is **Success**, and then click **Finish**.
14. In the Hyper-V Manager, right-click the LON-CORE virtual machine, and then select **Settings**.
15. In the Settings for LON-CORE on LON-HOST1 window, expand **Processor**, click **Compatibility**, and then select the option **Migrate to a physical computer with a different processor version**.
16. Click **OK**.
17. In Failover Cluster Manager, click **Roles**, right-click **LON-CORE**, and then select **Start**. Make sure that the virtual machine starts.
18. After the virtual machine starts, right-click it, and then select **Connect**.
19. In the Virtual Machine Connection window, click **Ctrl + Alt+ Delete** and then sign in as **Adatum\Administrator** with the password **Pa\$\$w0rd**.

► **Task 4: Perform a Live Migration for the virtual machine**

1. On the LON-HOST2, open **Windows PowerShell**.
2. In the Windows PowerShell® window, type **ping lon-core -t**, and then press Enter.
3. Verify that you receive replies from the LON-CORE IP address. Leave the Windows PowerShell window open.
4. Open the Failover Cluster Manager console, expand **VMCluster.Adatum.com**, and then click **Roles**.
5. Right-click the **20414C-LON-CORE** virtual machine, select **Move**, select **Live Migration**, and then click **Select Node**.
6. Click **LON-HOST2**, and then click **OK**.
7. Switch to the Windows PowerShell window, and then monitor the Internet Control Message Protocol (ICMP) packets. You should have minimal (one or two) or no lost packets.
8. Keep the Windows PowerShell window open.

► **Task 5: Validate high availability in the event of storage failure**

1. On LON-SVR1, open **Failover Cluster Manager**.
2. Expand **FSCluster.Adatum.com**, and then click **Roles**.
3. In the central pane, make sure that AdatumFS is running. Read the value of the **Owner** node column (it will be either LON-SVR1 or LON-SVR2).
4. Open **Hyper-V Manager** on LON-HOST1, right-click the virtual machine that owns the AdatumFS (from step three), and then click **Settings**.
5. Click **Network Adapter** (connected to External Network) in the left pane. In the right pane, in the **Virtual Switch** drop-down list box, select **Not Connected**, and then click **OK**.
6. Immediately switch back to Failover Cluster Manager on the current owner of AdatumFS, and then ensure that the owner node changes for **AdatumFS**. (You might have to refresh the console to see the change).
7. Click on **Nodes**, and then verify that one node is down.
8. Switch to the Windows PowerShell window on LON-HOST2, and then make sure that there are no lost packets.

9. Open **Hyper-V Manager** on LON-HOST1, right-click the virtual machine that you disconnected in step six, and then click **Settings**.
10. Click **Network Adapter** in the left pane. In the right pane, in the **Virtual switch** drop-down list box, select **External Network**, and then click **OK**.
11. Switch back to Failover Cluster Manager on LON-SVR1, and then ensure that both nodes are up and running (it might take up to a minute for the console to refresh the status. If you still cannot view the FSCluster.Adatum.com node, then reopen **Failover Cluster Manager**).
12. On LON-HOST1, in Failover Cluster Manager, right-click **LON-CORE**, and then click **Shut Down**.
13. On LON-HOST1 and LON-HOST2, close Failover Cluster Manager.

Results: After completing this lab, students will have virtual machines implemented in a highly available infrastructure.

Exercise 6: Implementing Operations Manager and VMM Integration

► Task 1: Import management packs

1. Switch to **LON-OM1**.
2. Open the Operations console from the taskbar.
3. Click the **Administration** workspace.
4. Expand **Administration**, expand **Device Management**, and then click **Management Packs**. Wait for the list of Management Packs to appear.
5. Click **Import Management Packs**.
6. In the Import Management Packs window, click **Add**, and then click **Add from disk**.
7. In the Online Catalog Connection window, click **No**.
8. In **Select Management Packs to import**, in the **File Name** text box, type `\\lon-vmm1\C$`, and then click **Open**.
9. Browse to **Program Files** then open **Microsoft System Center 2012 R2**, open **Virtual Machine Manager** and then open **Management Packs**.
10. Select all files, and then click **Open**.
11. In the Select Management Packs window, click **Install**.
12. In the Operations Manager window, click **Yes**. Wait until the management packs are imported, and then click **Close**.

► Task 2: Enable VMM integration with Operations Manager

1. On LON-VMM1, open the Virtual Machine Manager Console by clicking its icon on the task bar and then clicking **Connect**. (Note: If the VMM console does not start, open **Services** from Administrative Tools and check if all VMM services set to start automatically are running)
2. Click the **Settings** workspace.
3. Click **System Center Settings**.
4. In the central pane, double-click **Operations Manager Server**.
5. In the Add Operations Manager window, on the **Introduction** page, click **Next**.
6. On the **Connection to Operations Manager** page, in the **Server name** text box, type `lon-om1.adatum.com`. Click **Use a Run As account**, and then click **Browse**.
7. Select **Administrator**, and then click **OK**.
8. Click **Next**.
9. On the **Connection to VMM** page, type `Adatum\Administrator` for User name and `Pa$$w0rd` for password, and then click **Next**.
10. On the **Summary Page**, click **Finish**.
11. Wait until the job is completed, and then close the **Jobs** window. If the job does not complete successfully, restart LON-OM1 and LON-VMM1 and repeat steps one through 10.

► Task 3: Validate PRO integration

1. On LON-OM1, in the Operations Manager console, click the **Monitoring** workspace.
2. In the **Monitoring** workspace, expand **Monitoring**, expand **Microsoft System Center Virtual Machine Manager PRO**, and then click **PRO Object State**.
3. In the central pane, ensure that **LON-VMM1.adatum.com** appears.
4. Ensure that its state is **Healthy**. If its state is **Not monitored**, wait for a couple of minutes, and then refresh the console.
5. Switch to the LON-VMM1 machine.
6. In the VMM Manager console, click **Settings**.
7. Click **System Center Settings**, and then in the right pane, double-click **Operations Manager Server**.
8. In the Operations Manager Settings window, click **Test PRO**.
9. In the Virtual Machine Manager window, click **OK**, and then in the Operations Manager Settings window, click **OK**.
10. Click the **Jobs Workspace**, look for a job called **PRO Diagnostic**, wait for the status **Completed**, and then close the **PRO** dialog box. Click **History**, look for a job called **PRO Diagnostic**, select it, and then click **Details**.
11. Make sure that all tasks in this job completed successfully.

► Task 4: 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 **20414C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps two and three for 20414C-LON-SVR1, 20414C-LON-SVR2, 20414C-LON-VMM1, 20414C-LON-WSUS and 20414C-LON-OM1.
5. On LON-HOST1, from Server Manager, remove the Failover Clustering feature by using the Server Manager console. Repeat the same procedure on LON-HOST2.
6. Restart LON-HOST1 and LON-HOST2. When the server starts, select **20414C-LON-HOST1** in the start menu.

Results: After completing this exercise, students will have Performance and Resource Optimization (PRO) implemented.

Module 9: Planning and Implementing a Business Continuity Strategy

Lab: Implementing a Virtual Machine Backup Strategy with DPM


Exercise 1: Configuring DPM

► Task 1: Configure a storage pool in Microsoft System Center 2012 Data Protection Manager

1. Switch to LON-DM1, and then minimize Server Manager.
2. On the desktop, double-click **Microsoft System Center 2012 R2 Data Protection Manager**.
3. In System Center 2012 R2 DPM Administrator Console, click the **Management** workspace, and then, on the navigation bar, click the **Disks** link.
4. On the ribbon, click **Add**. The **Add Disks to Storage Pool** dialog box displays.
5. In the Add Disks to Storage Pool window, select **Disk1**, click **Add**, and then click **OK**.
6. You should see **Disk1 (Virtual HD ATA Device)** added in the **DPM Storage Pool**, marked with a green check box.

► Task 2: Deploy DPM protection agents

1. Switch to LON-SVR1.
2. On the taskbar, click the **Server Manager** icon.
3. In Server Manager, on the upper right side of the window, click **Tools**, and then click **Windows Firewall with Advanced Security**.
4. In the Windows Firewall with Advanced Security window, in the navigation pane, right-click **Windows Firewall with Advanced Security on Local Computer**, and then click **Properties**.
5. In the Windows Firewall with Advanced Security on Local Computer window, click the **Inbound connections** drop-down list box, select **Allow**, and then click **OK**.
6. Close the Windows Firewall with Advanced Security window.
7. Switch to LON-HOST1.
8. Repeat steps 1 to 6 on LON-HOST1.

 **Note:** In a production environment, you might customize firewall settings on corporate servers to allow traffic from IP addresses and ports needed for communication with DPM.

9. Switch to LON-DM1.
10. In the DPM Administrator Console, click the **Management** workspace, and then, on the navigation bar, click the **Agents** link.
11. On the ribbon, click **Install**. The Protection Agent Installation Wizard starts.
12. On the **Select Agent Deployment Method** page, click **Install agents**, and then click **Next**.

13. On the **Select Computers** page, select **LON-SVR1** and **LON-HOST1** from the list, click **Add**, and then click **Next**.
14. On the **Enter Credentials** page, in the **User name** box, type **Administrator**, in the **Password** box, type **Pa\$\$w0rd**, and then click **Next**.
15. On the **Choose Restart Method** page, click **Yes**, and then click **Next**.
16. On the **Summary** page, click **Install**. Wait until the **Results** column shows a **Success** status for each of the servers on which you installed the agent.
17. Click **Close** to close the Protection Agent Installation Wizard.

► Task 3: Configure protection groups

Create a data protection group

1. On LON-DM1, in the DPM Administrator Console, click the **Protection** workspace.
2. On the ribbon, click **New**. The Create New Protection Group Wizard starts.
3. In the Welcome to the New Protection Group Wizard, click **Next**.
4. On the **Select Protection Group Type** page, select **Servers**, and then click **Next**.
5. On the **Select Group Members** page, under **Available members**, expand **LON-SVR1**, expand **All Volumes**, expand drive **C:**, select the **Financial Data** folder check box, click **OK**, and then click **Next**.
6. On the **Select Data Protection Method** page, in the **Protection group name** box, type **Protection Group Financial Data Folder**, ensure that **I want short-term protection using Disk** is selected, and then click **Next**.
7. On the **Specify Short-Term Goals** page, accept the default settings, and then click **Next**.
8. On the **Review Disk Allocation** page, click **Next**.
9. On the **Choose Replica Creation Method** page, click **Next**.
10. On the **Consistency check options** page, click **Next**.
11. On the **Summary** page, click **Create Group**.
12. On the **Status** page, verify that both processes show **Success** in the **Results** column, and then click **Close**.

Create a virtual machine protection group

1. On LON-DM1, in the DPM Administrator Console, click the **Protection** workspace.
2. On the ribbon, click **New**. The Create New Protection Group Wizard starts.
3. In the Welcome to the New Protection Group Wizard, click **Next**.
4. On the **Select Protection Group Type** page, select **Servers**, and then click **Next**.
5. On the **Select Group Members** page, under **Available members**, expand **LON-HOST1**, expand **HyperV**, select the **Offline\20414C-LON-TEST** check box, and then click **Next**.
6. On the **Select Data Protection Method** page, in the **Protection group name** box, type **VM Protection Group**, ensure that **I want short-term protection using Disk** is selected, and then click **Next**.

7. On the **Specify Short-Term Goals** page, accept the default settings, and then click **Next**.
8. On the **Review Disk Allocation** page, click **Next**.
9. On the **Choose Replica Creation Method** page, click **Next**.
10. On the **Consistency check options** page, click **Next**.
11. On the **Summary** page, click **Create Group**.
12. On the **Status** page, verify that both processes show **Success** in the **Results** column, and then click **Close**.

Results: After completing these tasks, you will have created a storage pool in the Microsoft® System Center 2012 R2 Data Protection Manager containing Disk1. Next, you will have deployed Data Protection Manager (DPM) protection agents on LON-SVR1 and LON-HOST1. At the end of this exercise, you will have configured two protection groups. You will use the first protection group to protect data located in the Financial Data folder within the virtual machine LON-SVR1. You will use the second protection group to protect the virtual machine LON-TEST located on the physical host LON-HOST1.

Exercise 2: Implementing Backup and Restore for Virtual Machine Data

► Task 1: Configure DPM to back up virtual machine data

1. On LON-DM1, in the DPM Administrator Console, click the **Protection** workspace.
2. In the details pane, ensure that the status of the **Protection Group Financial Data Folder** is marked with a green check box.
3. In the details pane, right-click **C:\Financial Data**, and then click **Create recovery point**. The **Create recovery point** dialog box opens.
4. In the **Create recovery point** dialog box, click **OK**.
5. On the **Create Recovery Point** page, verify that both processes show **Success** in the **Results** column, and then click **Close**.

► Task 2: Delete data

1. Switch to LON-SVR1.
2. On the taskbar, click **File Explorer**, and then, in the navigation pane, click **Local Disk (C:)**.
3. In File Explorer, in the details pane, right-click the **Financial Data** folder, and then click **Delete**.
4. Close File Explorer.

► Task 3: Restore the deleted data

1. Switch to LON-DM1.
2. In the DPM Administrator Console, click the **Recovery** workspace.
3. In the navigation pane, expand **LON-SVR1**, expand **All Protected Volumes**, and then click **C:**.
4. In the results pane, select and right-click **Financial Data**, and then click **Recover**. The Recovery Wizard starts.
5. On the **Review Recovery Selection** page, click **Next**.
6. On the **Select Recovery Type** page, click **Recover to the original location**, and then click **Next**.
7. On the **Specify Recovery Options** page, click **Next**.
8. On the **Summary** page, click **Recover**.
9. On the **Recovery Status** page, verify that the **Recovery status** is **Successful**, and then click **Close**.
10. Click the **Protection** workspace, and then click **All Protection Groups**.
11. In the details pane, right-click **C:\Financial Data**, and then click **Perform consistency check**. Click **Yes** in the message box. In a few moments, the **Protection Status** should show a green check mark.
12. Switch to LON-SVR1.
13. On LON-SVR1, on the taskbar, click **File Explorer**, in the navigation pane, click **Local Disk (C:)**, and then ensure that the **Financial Data** folder has been restored.
14. Close File Explorer.

Results: After completing the exercise, you will have configured DPM to back up virtual machine data from the Protection Group Financial Data Folder created in the first exercise. After completing the backup, you will simulate data loss by deleting the Financial Data folder on LON-SVR1. Then you will restore the deleted data by using DPM.


Exercise 3: Implementing Virtual Machine Backup and Recovery by using DPM

► Task 1: Back up a virtual machine by using DPM

1. On LON-DM1, in the DPM Administrator Console, click the **Protection** workspace, and then, in the details pane, click **Protection Group: VM Protection Group**.
2. In the details pane, ensure that the status of the **VM Protection Group** is marked with a green check mark. It may take as long as 10 minutes for the status to show **OK**.
3. In the details pane, right-click **\Offline\20414C-LON-TEST**, and then click **Create recovery point**. The **Create recovery point** dialog box opens.
4. In the **Create recovery point** dialog box, click **OK**.
5. On the **Create Recovery Point** page, verify that the task shows **Success** in the **Results** column, and then click **Close**.

► Task 2: Change a configuration in the virtual machine

- a. Switch to LON-HOST1.
- b. In Hyper-V Manager, right-click **20414C-LON-TEST**, and then click **Settings**.
- c. In **Settings for 20414C-LON-TEST**, in the left pane, under **Hardware**, click **Memory**, and then, in the right pane, next to **Startup RAM**, type **256**, and then click **OK**.

 **Note:** This change will disrupt LON-TEST from normal operation because of the small amount of memory allocated. In the next task, you will restore the original memory setting for LON-TEST from backup.

► Task 3: Restore the virtual machine

1. Switch to LON-DM1.
2. In the DPM Administrator Console, click the **Recovery** workspace.
3. In the navigation pane, expand **Recoverable Data\Adatum.com\LON-HOST1**, and then click **All Protected Hyper-V Data**.
4. In the results pane, under **Recoverable Item**, right-click **Offline\20414C-LON-TEST**, and then click **Recover**. The Recovery Wizard starts.
5. On the **Review Recovery Selection** page, click **Next**.
6. On the **Select Recovery Type** page, click **Recover to original instance**, and then click **Next**.
7. On the **Specify Recovery Options** page, click **Next**.
8. On the **Summary** page, click **Recover**.
9. On the **Recovery Status** page, verify that the **Recovery status** is **Successful**, and then click **Close**.
10. Switch to LON-HOST1.
11. In Hyper-V Manager, right-click **20414C-LON-TEST**, and then click **Settings**.
12. In **Settings for 20414C-LON-TEST**, in the left pane, under **Hardware**, click **Memory**, and then, in the right pane, next to **Startup RAM**, verify that a value of **1024** has been restored, and then click **OK**.

► **Task 4: To prepare for the next module**

When you are finished with 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 **20414C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20414C-LON-SVR1** and **20414C-LON-DM1**.

Results: After completing this exercise, you will have configured DPM to back up the virtual machine LON-TEST located on the LON-HOST1 physical host from the VM Protection Group created in the first exercise. After completing the backup, you will simulate corruption of the virtual machine by changing the configuration of LON-TEST in the Microsoft Hyper-V® console on LON-HOST1. At the end, you will restore the corrupted virtual machine configuration by using DPM.

Module 10: Planning and Implementing a Public Key Infrastructure

Lab: Planning and Implementing an Active Directory Certificate Services Infrastructure

Exercise 1: Planning an Active Directory Certificate Services Deployment

► Task 1: Read the supporting documentation

- Read the documentation that the student handbook provides.

► Task 2: Propose a solution and plan a course of action

Based on the lab scenario, propose a solution for a public key infrastructure (PKI) design. Use the following questions as guidance for your PKI design development:

1. How will you address the requirement that the root certification authority (CA) at A. Datum must remain completely isolated from the network at all times?

Answer: You can address the requirement for isolation of the root CA by deploying an offline root CA.

2. How will you make the CA role highly available?

Answer: You can achieve high availability of the CA role by deploying a CA cluster or by deploying multiple CAs for each region.

3. How will you achieve certificate-based authentication for managers and applications?

Answer: You can achieve certificate-based authentication by deploying smart cards to all managers and to all users who use applications that require certificate-based authentication. Also, you can force managers to sign in only with smart cards, by configuring options in their user accounts in Active Directory® Domain Services (AD DS).

4. How will you configure a web server certificate to address requirements?

Answer: To configure a web server certificate, you will duplicate an existing web server certificate template, and then configure the required options on the new template. The new template will be used instead of the old one by configuring superseding.

5. How will you address the requirement that only laptops receive certificates for deployment of DirectAccess?

Answer: You should place all laptop computers in a separate organizational unit (OU). In addition, you should duplicate the computer certificate template, and make a new one called Windows® 8 DirectAccess with the options required in the lab scenario. You will configure a discretionary access control list (DACL) on that certificate template, so that only laptop computers can enroll for that certificate. Then you will configure and link the new Group Policy Object (GPO) for the laptop's OU. In that GPO, you will configure autoenrollment for computer certificates.

6. How will you address the requirement for revocation checking between sites?

Answer: To minimize the bandwidth used for revocation checking between sites, you can deploy Online Responder functionality in each site.

7. How will you address the requirement for Internet-based revocation checking?

Answer: To enable revocation checking from the Internet, you should configure a new certificate revocation list distribution point (CDP) and authority information access (AIA) locations in the properties of the issuing CA. Then, you should publish CRL and AIA files to that location and make that location accessible from the Internet.

► **Task 3: Examine the suggested proposals in the Lab Answer Key**

Compare your solution with the solution that the Lab Answer Key provides. Discuss alternative solutions with the class and instructor.

Results: After completing this exercise, students will have a plan for Active Directory Certificate Services (AD CS) deployment.

Exercise 2: Deploying a CA Infrastructure

► Task 1: Install the Active Directory Certificate Services role on a stand-alone server

1. Sign in 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 the **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 **Configure Active Directory Certificate Services on the destination server**.
12. In the Active Directory Certificate Services (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 provider 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 certificate revocation for the subordinate CA

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 box, select **CRL Distribution Point (CDP)**, and then click **Add**.

5. In the **Location** text box, type **http://lon-svr1.adatum.com/CertData/**, and, 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 the URL, type **.crl**, and then click **OK**.
8. Select the following options: **Include in the CDP extension of issued certificate**, 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/**, and then, in the **Variable** drop-down box, click **<ServerDNSName>**, and then click **Insert**.
11. In the **Location** text box, type an underscore (**_**), and, 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 the 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 the 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 File Explorer window, and then 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 File Explorer address bar, type **\\lon-svr1\C\$**, and then press Enter.
30. Right-click the empty space, and then click **Paste**.
31. Close File Explorer.

► **Task 3: Publish the Root CA certificate in the domain**

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 the **RootCA.cer** file, and then click **Open**.
8. Click **Next** two times, and then click **Finish**.
9. When the Certificate Import Wizard prompt appears, click **OK**.
10. Close the Group Policy Management Editor.
11. Close the Group Policy Management console.

► **Task 4: Configure an AD CS role on the subordinate enterprise CA**

1. Sign in to LON-SVR1 as **Adatum\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**.
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, and then 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 **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, click **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 5: Configure the certificates to enable the subordinate trust**

1. On LON-SVR1, open a File 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**. When the Certificate Import Wizard window appears, click **OK**.
7. In the File Explorer window, select the **AdatumRootCA.crl** and **LON-CA1_AdatumRootCA.crt** files, right-click the files, and then click **Copy**.
8. Double-click **inetpub**.
9. Double-click **wwwroot**.
10. Create a new folder, and then name it **CertData**.
11. Paste the two files that you copied into that folder.
12. Switch to **Local Disk (C:)**.
13. Right-click the file **LON-SVR1.Adatum.com_Adatum-LON-SVR1-CA.req**, and then click **Copy**.
14. In the File Explorer address bar, type **\\LON-CA1\C\$**, and then press Enter.
15. In the File Explorer window, right-click an empty space, and then click **Paste**. Make sure that the request file copies to LON-CA1.
16. Switch to the LON-CA1 server.
17. In the Certification 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 the **LON-SVR1.Adatum.com_Adatum-LON-SVR1-CA.req** file, and then click **Open**.
19. In the Certification Authority console, expand **AdatumRootCA**, and then click the **Pending Requests** container. Right-click **Pending Requests**, and then 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 to 20 seconds, and then, on the toolbar, click the green icon to start the CA service.
35. Ensure that the CA starts successfully.

Results: After completing this exercise, students will have deployed a CA infrastructure.

Exercise 3: Implementing Certificate Templates

► Task 1: Configuring OUs and groups for laptop computers

1. On the LON-DC1, from **Server Manager**, click **Tools**, and then open **Active Directory Users and Computers**.
2. In the Active Directory Users and Computers console, right-click the **Adatum.com** domain object, click **New**, and then click **Organizational Unit**.
3. In the New Object – Organizational Unit window, in the **Name** text box, type **Laptops**, and then click **OK**.
4. In the Active Directory Users and Computers console, right-click the **Laptops** OU, click **New**, and then click **Group**.
5. In the New Object – Group window, in the **Group name** text box, type **Laptops**, and then click **OK**.
6. In the Active Directory Users and Computers console, click the **Computers** container.
7. Right-click **LON-CL1**, click **Move**, and, in the Move window, select the **Laptops** OU, and then click **OK**.
8. Click the **Laptops** OU in the console, right-click the **LON-CL1** computer object, and then select **Properties**.
9. Click the **Member Of** tab, click **Add**, in the Select Groups window, type **Laptops**, and then click **OK**.
10. Click **OK**.
11. Close the Active Directory Users and Computers console.

► Task 2: Configure the web server certificate template

1. On LON-SVR1, in the Certification Authority console, expand **Adatum-IssuingCA**, right-click **Certificate Templates**, and then click **Manage**.
2. In the Certificate Templates console, locate and right-click the **Web Server** template, and then click **Duplicate Template**.
3. Click the **General** tab.
4. In the **Template display name** field, type **Adatum Web Server Certificate**, and then 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 3: Configure the DirectAccess certificate template

1. On LON-SVR1, in the Certificate Templates console, locate and right-click the **Computer** template, and then click **Duplicate Template**.
2. Click the **General** tab.
3. In the **Template display name** field, type **DirectAccess Clients**, and then set the **Validity period** to **6 months**.
4. Click the **Security** tab.
5. Click **Add**.
6. In the Select Users, Computers, Service Accounts, or Groups window, type **Laptops**, and then click **Check Names**.
7. Click **OK**.

8. Select the **Laptops** user group, and then, in the lower pane, select the **Allow** option for **Read, Enroll** and **Autoenroll**.

9. Click **OK**.

► **Task 4: Configure the smart card certificate**

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 then, in the **Template display name** text box, type **Adatum Smart Card User**.
3. On the **Subject Name** tab, clear 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. Click the **Security** tab.
10. Click **Add**.
11. In the Select Users, Computers, Service Accounts, or Groups window, type **Managers**, and then click **Check Names**.
12. Click **OK**.
13. Select the **Managers** user group, and then, in the lower pane, select the **Allow** option for **Read, Enroll** and **Autoenroll**.
14. Click **OK**.
15. Close the **Certificate Templates** console.

► **Task 5: Issue the certificate templates**

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**, **DirectAccess Clients**, and **Adatum Web Server**, and then click **OK**.

Results: After completing this exercise, students will have configured certificate templates to meet security requirements.

Exercise 4: Implementing Certificate Revocation and Distribution

► Task 1: Configure certificate revocation checking

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 (2 of 6 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 indicates that the installation was successful, 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 OCSP Management console, right-click **Revocation Configuration**, and then click **Add Revocation Configuration**.

26. In the Add Revocation Configuration wizard, click **Next**.
27. On the **Name the Revocation Configuration** page, in the **Name** box, type **AdatumCA Online Responder**, and then click **Next**.
28. On the **Select CA Certificate Location** page, click **Next**.
29. On the **Choose CA Certificate** page, click **Browse**, click the **Adatum-IssuingCA** certificate, click **OK**, and then click **Next**.
30. On the **Select Signing Certificate** page, verify that **Automatically select a signing certificate** 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.

► **Task 2: Configure certificate revocation checking for DirectAccess clients**

1. On LON-SVR1, in the Server Manager console, click **Tools**, and then click **Certification Authority**.
2. In the certsrv – [Certification Authority (Local)] console, right-click **Adatum-IssuingCA**, and then click **Properties**.
3. In the Adatum-IssuingCA Properties window, click the **Extensions** tab.
4. In the **Extensions** tab, in the **Select extension:** drop-down list box, select **CRL Distribution Point (CDP)**, and then click **Add**.
5. In the **Location** text box, type **http://www.adatum.com/CertData/**, and, 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 the URL, type **.crl**, and then click **OK**.
8. Select the **Include in the CDP extension of issued certificates** and **Include in CRLs. Clients use this to find Delta CRL locations** options, and then 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://www.adatum.com/CertData/**, and, in the **Variable** drop-down box, click **<ServerDNSName>**, and then click **Insert**.
11. In the **Location** text box, type an underscore (**_**), and, 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 the 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 the Certification Authority service.

► **Task 3: Configure autoenrollment for user and computer certificates**

1. On LON-DC1, in **Server Manager**, open **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 the **Laptops** OU, and then select the **Create a GPO in this domain, and Link it here** option.
3. In the New GPO window, in the **Name** text box, type **DirectAccessCert**, and then click **OK**.
4. Click the **Laptops** OU, and, in the right pane, right-click **DirectAccessCert**, and then click **Edit**.
5. In the **Group Policy Management Editor**, under the **Computer Configuration** node, expand **Policies**, expand **Windows Settings**, expand **Security Settings**, and then click **Public Key Policies**.
6. In the right pane, double-click **Certificate Services Client – Auto-Enrollment**.
7. In the Certificate Services Client – Auto-Enrollment Properties window, in the **Configuration Model** drop-down list box, select **Enabled**.
8. Select the **Renew expired certificates, update pending certificates, and remove revoked certificates** check box, select the **Update certificates that use certificate templates** check box, and then click **OK**.
9. Close the Group Policy Management Editor.
10. In the Group Policy Management console, expand **Forest:Adatum.com**, expand **Domains**, expand **Adatum.com**, right-click the **Managers** OU, and then select the **Create a GPO in this domain, and Link it here** option.
11. In the New GPO window, in the **Name** text box, type **SmartCardCert**, and then click **OK**.
12. Click the **Managers** OU, and, in the right pane, right-click **SmartCardCert**, and then select **Edit**.
13. In the **Group Policy Management Editor**, under the **User Configuration** node, expand **Policies**, expand **Windows Settings**, expand **Security Settings**, and then click **Public Key Policies**.
14. In the right pane, double-click **Certificate Services Client – Auto-Enrollment**.
15. In the Certificate Services Client – Auto-Enrollment Properties window, in the **Configuration Model** drop-down list box, select **Enabled**.
16. Select the **Renew expired certificates, update pending certificates, and remove revoked certificates** check box, select the **Update certificates that use certificate templates** check box, and then click **OK**.
17. Close the **Group Policy Management Editor** and close the **Group Policy Management** console.

► **Task 4: Validate certificate enrollment methods for smart card and DirectAccess certificates**

1. Sign in to LON-CL1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. On **LON-CL1**, open the Command Prompt window by typing **cmd.exe** on the Start screen.
3. Type **gpupdate /force**, and then press Enter.
4. Type **mmc.exe**, and then press Enter.
5. In the Console1 window, click **File**, and then select **Add/Remove Snap-in**. In the Add or Remove Snap-ins console, select **Certificates**, and then click **Add**.
6. In the **Certificates** snap-in, select **Computer account**, click **Next**, and then click **Finish**. Click **OK**.
7. In the Console1 window, expand **Certificates**, expand **Personal**, and then click **Certificates**.

8. Ensure that the certificate is issued to LON-CL1.Adatum.com. In the right pane, scroll to the right, and verify that the template used for the certificate is **DirectAccess Clients**.
9. Sign out from LON-CL1, and then sign in as **Adatum\Aidan** with the password **Pa\$\$w0rd**.
10. On **LON-CL1**, open the Command Prompt window by typing **cmd.exe** on the Start screen.
11. Type **gpupdate /force**, and then press Enter. Wait until Group Policy refreshes.
12. Type **mmc.exe**, and then press Enter.
13. In the Console1 window, click **File**, and then select **Add/Remove Snap-in**. In the Add or Remove Snap-ins console, select **Certificates**, and then click **Add**.
14. Click **OK**.
15. In the Console1 window, expand **Certificates**, expand **Personal**, and then click **Certificates**.
16. Verify that the user Aidan Delaney is issued a certificate based on the Adatum Smart Card User template.

► **Task 5: Prepare for the next module**

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

1. On the host computer, start Microsoft Hyper-V® Manager.
2. On the **Virtual Machines** list, right-click **20414C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps two and three for 20414C-LON-SVR1, 20414C-LON-CA1, and 20414C-LON-CL1.

Results: After completing this exercise, students will have configured certificate revocation and distribution.

Module 11: Planning and Implementing an Identity Federation Infrastructure

Lab: Planning and Implementing AD FS Infrastructure

Exercise 1: Designing the AD FS Deployment


► **Task 1:** Choose the appropriate AD FS deployment scenario

Based on the requirements identified, determine whether A. Datum Corporation requires a web single sign-on (SSO) or a federated web SSO deployment to enter into a partnership with Trey Research.

A. Datum will provide SSO functionality initially for internal users. Then it will partner with Trey Research and other companies and customers. Therefore, you will deploy a federated web SSO.

► **Task 2: Identify the number of AD FS servers that the deployment requires, and the role and location of each of the federation servers that you will deploy**

1. Identify how many AD FS servers you will require for the deployment of AD FS in A. Datum and Trey Research.
 - Because this will be a federated web SSO deployment, you will require two servers for this deployment.
2. Determine the role and location of each of the federation servers that you will deploy.
 - Datum will host a resource federation server, while an account federation server will be located in Trey Research.

 **Note:** Trey Research is not hosting any applications, so their AD FS server is an account federation server. Note, however, that A. Datum's AD FS server is a resource federation server from the perspective of federated web SSO. However, you can consider it an account federation server also, when providing SSO functionality to A. Datum users.

Results: In this exercise, you should have identified the appropriate Active Directory® Federation Services (AD FS) deployment scenario to use to meet the defined requirements. You also should have identified the number of AD FS servers required for the deployment, their locations, and the role of each AD FS server in the deployment.

Exercise 2: Configuring Prerequisite Components for AD FS

► Task 1: Configure DNS forwarders and add DNS records

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 the **IP address** column, and then type **172.16.10.10**. Press Enter, and then click **OK**.
6. In the left pane, expand **Forward Lookup Zones**, and then click **Adatum.com**.
7. Right-click **Adatum.com**, and then click **New Host (A or AAAA)** from the context menu.
8. In the New Host window, type **adfs** in the **Name** text box and type **172.16.0.10** in the **IP address** text box.
9. Click **Add Host**. Click **OK**, and then click **Done**.
10. Close the DNS Manager.
11. On TREY-DC1, in Server Manager, click **Tools**, and then click **DNS**.
12. Expand **TREY-DC1**, and then click **Conditional Forwarders**.
13. Right-click **Conditional Forwarders**, and then click **New Conditional Forwarder**.
14. In the **DNS Domain** box, type **Adatum.com**.
15. Click the **IP address** column, and then type **172.16.0.10**. Press Enter, and then click **OK**.
16. In the left pane, expand **Forward Lookup Zones** and then click **TreyResearch.net**.
17. Right-click **TreyResearch.net**, and then click **New Host (A or AAAA)** from the context menu.
18. In the New Host window, type **adfs** in the **Name** text box and type **172.16.10.10** in the **IP address** text box.
19. Click **Add Host**. Click **OK**, and then click **Done**.
20. Close the DNS Manager.

► Task 2: Exchange root certificates to enable certificate trusts

1. On LON-DC1, click to the **Start** page.
2. On the **Start** page, type **\\TREY-DC1.treyresearch.net\certenroll**, and then press Enter.
3. In the CertEnroll window, right-click the **TREY-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. On the Start screen, type **MMC.exe**, 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 the **Select Group Policy Object** window, 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**, expand **Policies**, expand **Windows Settings**, expand **Security Settings**, expand **Public Key Policies**, and then click **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 **TREY-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 when the Certificate Import Wizard prompt displays, click **OK**.
 18. Close the Group Policy Management Editor without saving changes.
 19. On TREY-DC1, click to the **Start** page.
 20. On the **Start** page, type `\\LON-DC1.adatum.com\certenroll`, and then press Enter.
 21. In the CertEnroll window, right-click the **LON-DC1.Adatum.com_AdatumCA.crt** file, and then click **Copy**.
 22. In the left pane, click **Documents**, and then paste the file into the **Documents** folder.
 23. Right-click **Start**, click **Run**, 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 then click **Next**.
 27. Verify that **Local computer** is selected, click **Finish**, and then click **OK**.
 28. Expand **Certificates**, and then click **Trusted Root Certification Authorities**.
 29. Right-click **Trusted Root Certification Authorities**, point to **All Tasks**, and then click **Import**.
 30. On the **Welcome to the Certificate Import Wizard** page, click **Next**.
 31. On the **File to Import** page, click **Browse**.
 32. In the Open window, click **LON-DC1.Adatum.com_AdatumCA.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 when the Certificate Import Wizard prompt displays, click **OK**.
 35. Close Console1 without saving changes.
- **Task 3: Request and install a certificate for the web server**
1. On LON-SVR3, in Server Manager, click **Tools**, and then click **Internet Information Services (IIS) Manager**.
 2. In the console tree, click **LON-SVR3 (Adatum\Administrator)**.
 3. In the 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-SVR3.adatum.com**
 - Organization: **A. Datum**
 - Organizational unit: **IT**
 - City/locality: **London**
 - State/province: **England**
 - Country/region: **GB**
6. On the **Online Certification Authority** page, in the **Specify Online Certification Authority** text box, click **Select** to search for a certification authority (CA) server in the domain.
7. Select **AdatumCA**, and then click **OK**.
8. In the **Friendly Name** text box, type **LON-SVR3.adatum.com**, and then click **Finish**.

► **Task 4: Bind the certificate to the claims-aware application on the web server, and then verify application access**

1. On LON-SVR3, 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-SVR3.adatum.com**, and then click **OK**.
4. Click **Close**, and then close IIS.
5. On LON-DC1, open Internet Explorer®.
6. In the address bar, type **https://lon-svr3.adatum.com/adatumtestapp**, and then press Enter.
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.

► **Task 5: Prepare a certificate template and install a certificate for the Trey Research AD FS server**

Prepare the certificate template

1. On Trey-DC1, in Server Manager, click **Tools**, and then click **Certification Authority**.
2. Expand **TreyResearchCA**. Click **Certificate Templates**.
3. Right-click **Certificate Templates**, and then click **Manage**.
4. In the Certificate Templates Console window, right-click the **Web Server** template, and then click **Duplicate Template** from the context menu.
5. In the Properties of New Template window, click the **General** tab.
6. In the **Template display name** text box, delete the existing text, and then type **Trey Research Web Server**.
7. Click the **Security** tab. Click **Authenticated Users**, and then select the **Allow** checkbox for the **Enroll** permission. Click **OK**.
8. Close the Certificate Templates Console window.

9. In the certsrv – [Certification Authority (Local)\TreyResearchCA\Certificate Templates] window, right-click **Certificate Templates**, click **New**, and then click **Certificate Template to Issue**.
10. In the **Enable Certificate Templates** window, click **Trey Research Web Server**, and then click **OK**.

Request the certificate

1. On Trey-DC1, right-click **Start**, click **Run**, type **MMC**, and then click **OK**.
2. In **Console1**, click **File**, click **Add/Remove Snap-in**, click **Certificates**, and then click **Add**.
3. In the Certificates snap-in window, click **Computer account**, and then click **Next**.
4. In the Select computer window, ensure that **Local computer: (the computer this console is running on)** is selected, and then click **Finish**. Click **OK** to close the Add or Remove Snap-ins window.
5. In the Console1 window, expand **Certificates**, expand **Personal**, right-click **Certificates**, click **All Tasks**, and then click **Request New Certificate**.
6. On the **Before You Begin** page, click **Next**.
7. On the **Select Certificate Enrollment Policy** page, click the **Active Directory Enrollment Policy**, and then click **Next**.
8. In the Request Certificates window, click to select the **Trey Research Web Server** policy, and then click the down arrow icon next to **Details**. Click **Properties**.
9. In the Certificate Properties window, click the **Private Key** tab, expand **Key options**, and then click **Make private key exportable**.
10. Click the **Subject** tab. In the **Subject name** configuration area, click the **Type** dropdown menu, and then click **Common name**.
11. In the **Value** text box for the common name, type **adfs.treyresearch.net**, click **Add**, and then click **OK**.
12. In the Request Certificates window, click **Enroll** to enroll for the certificate.
13. In the Certificate Installation Results window, click **Finish**.

Results: In this exercise, you should have configured Domain Name System (DNS) forwarding to enable name resolution between A. Datum and Trey Research. In addition, you should have exchanged root certificates between the two organizations, and installed and configured a web certificate on the application server.

Exercise 3: Deploying AD FS for Internal Users

► Task 1: Install and configure the AD FS server role

1. On LON-DC1, click the Windows PowerShell® shortcut on the taskbar. At the Windows PowerShell prompt, run the **Add-KdsRootKey –EffectiveTime (Get-Date).AddHours(-10)** command.
2. Close the Windows PowerShell window.
3. In Server Manager, click **Manage**, and then click **Add Roles and Features**.
4. On the **Before you begin** page, click **Next**.
5. On the **Select installation type** page, click **Next**.
6. On the **Select destination server** page, click **Next**.
7. On the **Select server roles** page, select **Active Directory Federation Services**, and then click **Next**.
8. On the **Select features** page, click **Next**.
9. On the **AD FS** page, click **Next**.
10. On the **Confirm installation selections** page, click **Install**, and then wait for the installation to finish. Click **Close** when the installation completes.
11. In Server Manager, click the yellow notification icon, and then click the **Configure the federation service on this server** link.
12. On the **Welcome** page, ensure that **Create the first federation server in a federation server farm** is selected, and then click **Next**.
13. On the **Connect to Active Directory Domain Services** page, click **Next** to use the **ADATUM\Administrator** account.
14. On the **Specify Service Properties** page, select the **SSL certificate** that is named **adfs.adatum.com**. Type **adfs.adatum.com** as the **Federation Service Display Name**, and then click **Next**.
15. On the **Specify Service Account** page, click **Use an existing domain user account or group Managed Service Account**.
16. Click **Select** to select the account name. In the Select User or Service Account window, type **ADFS-SVC** in the text box, and then click **OK**.
17. On the **Specify Service Account** page, type **Pa\$\$w0rd** in the **Account Password** field, and then click **Next**.
18. On the **Specify Configuration Database** page, ensure that **Create a database on this server using Windows Internal Database** is selected, and then click **Next**.
19. On the **Review Options** page, verify that the correct configuration settings are listed, and then click **Next**.
20. On the **Pre-requisite Checks** page, click **Configure**.
21. Wait for the configuration to finish and then click **Close**.
22. On LON-DC1, in Server Manager, click **Tools**, and then click **Windows PowerShell**.
23. At the Windows PowerShell prompt, type **set-ADFSProperties –AutoCertificateRollover \$False**, and then press Enter. You must perform this step so that you can modify the certificates that AD FS uses.
24. Close the Windows PowerShell window.
25. In Server Manager, click **Tools**, and then click **AD FS Management**.

26. In the AD FS console, in the left pane, expand **Service**, and then click **Certificates**.
27. Right-click **Certificates**, and then click **Add Token-Signing Certificate**.
28. In the **Select a token-signing certificate** dialog box, click the certificate with the name **adfs.adatum.com**.
29. Click **OK** to close the **Windows Security** dialog box.
30. When the **AD FS Management** warning dialog box displays, click **OK**.



Note: Verify that the certificate has a subject of **CN=adfs.adatum.com**. If no name displays under the **Subject** when you add the certificate, delete the certificate, and then add the next certificate in the list.

31. Under **Token-signing**, right-click the newly added certificate, and then click **Set as Primary**. Review the warning message, and then click **Yes**.
32. 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 claim 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, on 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 Lightweight Directory Access Protocol (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 WIF application to trust the AD FS deployment

1. On LON-SVR3, click to the **Start** screen, and then type **Windows Identity Foundation Federation Utility**. In the search results, click **Windows Identity Foundation Federation Utility**.
2. On the **Welcome to the Federation Utility Wizard** page, in **Application configuration location**, type **C:\inetpub\wwwroot\AdatumTestApp\web.config** for the location of the web.config file of the WIF sample application.
3. In the **Application URI** text box, type **https://lon-svr3.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://adfs.adatum.com/federationmetadata/2007-06/federationmetadata.xml>** for the STS WS-Federation metadata document location, and then click **Next** to continue.
5. If the **STS signing certificate chain validation error** page is displayed, ensure that **Disable certificate chain validation** is selected, and then click **Next**.
6. On the **Security Token Encryption** page, ensure that **No encryption** is selected, and then click **Next**.
7. On the **Offered claims** page, review the claims that the federation server will offer, and then click **Next**.
8. On the **Summary** page, review the changes that the Federation Utility Wizard will make to the sample application, scroll through the items to understand what each item is doing, and then click **Finish**.
9. Click **OK** in the **Success** dialog box.

► **Task 4: Configure a relying party trust and claims rules**

1. On LON-DC1, in the AD FS Management console, click **AD FS**.
2. In the right pane, click **Add Relying Party Trust**.
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-svr3.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 **Configure Multi-Factor Authentication Now?** page, ensure that **I do not want to configure multi-factor authentication settings for this relying party trust at this time** is selected, and then click **Next**.
8. On the **Choose Issuance Authorization Rules** page, ensure that **Permit all users to access this relying party** is selected, and then click **Next**.
9. On the **Ready to Add Trust** page, review the relying party trust settings, and then click **Next**.
10. On the **Finish** page, click **Close**. The Edit Claim Rules for ADatum Test App window opens.
11. In the **Edit Claim Rules for Adatum Test App properties** dialog box, on the **Issuance Transform Rules** tab, click **Add Rule**.
12. 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 to the user by using integrated Windows® authentication.
13. On the **Configure Rule** page, in the **Claim rule name** text box, type **Pass Through Windows Account Name Rule**. In the **Incoming claim type** drop-down list, select **Windows account name**, and then click **Finish**.
14. Click **Add Rule**.
15. On the **Select Rule Template** page, under **Claim rule template**, select **Pass Through or Filter an Incoming Claim**, and then click **Next**.

16. On the **Configure Rule** page, in the **Claim rule name** text box, type **Pass Through Email Address Rule**, in the **Incoming Claim Type** drop-down list, select **E-mail Address**, and then click **Finish**.
17. Click **Add Rule**.
18. On the **Select Rule Template** page, under **Claim rule template**, select **Pass Through or Filter an Incoming Claim**, and then click **Next**.
19. On the **Configure Rule** page, in the **Claim rule name** text box, type **Pass Through UPN Rule**, in the **Incoming Claim Type** drop-down list, select **UPN**, and then click **Finish**.
20. Click **Add Rule**.
21. On the **Select Rule Template** page, under **Claim rule template**, select **Pass Through or Filter an Incoming Claim**, and then click **Next**.
22. On the **Configure Rule** page, in the **Claim rule name** text box, type **Pass Through Name Rule**, in the **Incoming Claim Type** drop-down list, select **Name**, and then click **Finish**.
23. Click **Apply**, and then click **OK**.

► **Task 5: Verify that approved internal users can access the application**

1. Sign in to LON-CL1 as **Adatum\Brad** with the password **Pa\$\$w0rd**.
2. On LON-CL1, click to the desktop, and then open Internet Explorer.
3. Connect to **https://lon-svr3.adatum.com/AdatumTestApp/**.



Note: Ensure that you type the trailing forward slash (/).

4. If you are prompted for credentials, type **Adatum\Brad** with the 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.
5. Close Internet Explorer.

Results: In this exercise, you should have installed and configured the AD FS server role on LON-DC1. You should also have configured the claim provider and relying party trusts, and then configured the necessary claims rules. Finally, you should have configured the Windows Identity Foundation (WIF) application to trust the AD FS deployment, and then verified that approved internal users could access the application.

Exercise 4: Deploying AD FS for a Partner Organization

► Task 1: Install and configure the AD FS server role on Trey-DC1

1. On Trey-DC1, click the Windows PowerShell shortcut on the taskbar. At the Windows PowerShell prompt, run the **Add-KdsRootKey –EffectiveTime (Get-Date).AddHours(-10)** command.
2. Close the Windows PowerShell window.
3. In Server Manager, click **Manage**, and then click **Add Roles and Features**.
4. On the **Before you begin** page, click **Next**.
5. On the **Select installation type** page, click **Next**.
6. On the **Select destination server** page, click **Next**.
7. On the **Select server roles** page, select **Active Directory Federation Services**, and then click **Next**.
8. On the **Select features** page, click **Next**.
9. On the **AD FS** page, click **Next**.
10. On the **Confirm installation selections** page, click **Install**, and then wait for the installation to finish. Click **Close** when the installation completes.
11. In Server Manager, click the yellow notification icon, and then click the **Configure the federation service on this server** link.
12. On the **Welcome** page, ensure that **Create the first federation server in a federation server farm** is selected, and then click **Next**.
13. On the **Connect to Active Directory Domain Services** page, click **Next** to use the **Treyresearch\Administrator** account.
14. On the **Specify Service Properties** page, select the **SSL certificate** that is named **adfs.treyresearch.net**. Type **adfs.treyresearch.net** as the **Federation Service Display Name**, and then click **Next**.
15. On the **Specify Service Account** page, click **Create a Group Managed Service Account**.
16. In the **Account Name** box, type **ADFS**, and then click **Next**.
17. On the **Specify Configuration Database** page, ensure that **Create a database on this server using Windows Internal Database** is selected, and then click **Next**.
18. On the **Review Options** page, verify that the correct configuration settings are listed, and then click **Next**.
19. On the **Pre-requisite Checks** page, click **Configure**.
20. Wait for the configuration to finish (note that a service principal name registration error may occur), and then click **Close**.
21. On the Trey-DC1 virtual machine, in Server Manager, click **Tools**, and then click **Windows PowerShell**.
22. At the Windows PowerShell prompt, type **set-ADFSProperties –AutoCertificateRollover \$False**, and then press Enter. You must perform this step so that you can modify the certificates that AD FS uses.
23. Close the Windows PowerShell window.
24. In Server Manager, click **Tools**, and then click **AD FS Management**.
25. In the AD FS console, in the left pane, expand **Service**, and then click **Certificates**.

26. Right-click **Certificates**, and then click **Add Token-Signing Certificate**.
27. Choose the certificate named **adfs.treyresearch.net**. Click **OK** to close the **Windows Security** dialog box.
28. When the **AD FS Management** warning dialog box displays, click **OK**.



Note: Verify that the certificate has a subject of **CN=adfs.treyresearch.net**. If no name displays under the **Subject** when you add the certificate, delete the certificate, and then add the next certificate in the list.

29. Under **Token-signing**, right-click the newly added certificate, and then click **Set as Primary**. Review the warning message, and then click **Yes**.
30. Select the certificate that has just been superseded, right-click the certificate, and then click **Delete**. Click **Yes** to confirm the deletion.

► **Task 2: 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, ensure that **Import data about the claims provider published online or on a local network** is selected, type **https://adfs.treyresearch.net** as the host name, 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 adfs.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**, click **Finish**, and then 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 Windows PowerShell prompt, type the following command, and then press Enter:

```
Set-ADFSClaimsProviderTrust -TargetName "adfs.treyresearch.net" -  
SigningCertificateRevocationCheck None
```


17. Close the Windows PowerShell window.

► **Task 3: Configure a relying party trust on TREY-DC1 for the A. Datum claims-aware application**

1. On TREY-DC1, in Server Manager, click **Tools**, and then click **AD FS Management**.
2. In the AD FS console, in the left pane, expand **Trust Relationships**. Right-click **Relying Party Trusts**, and then click **Add Relying Party Trust**.
3. On the **Welcome** page, click **Start**.
4. On the **Select Data Source** page, ensure that **Import data about the relying party published online or on a local network** is selected, type **https://adfs.adatum.com** for the host name, 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 **Configure Multi-factor Authentication Now?** page, ensure that **I do not want to configure multi-factor authentication settings for this relying party trust at this time** is selected, and then click **Next**.
7. On the **Choose Issuance Authorization Rules** page, ensure that **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** to save the configuration.
9. On the **Finish** page, click **Close**. The Edit Claim Rules for Adatum TestApp window opens.
10. In the Edit Claim Rules for Adatum TestApp window, on the **Issuance Transform Rules** tab, click **Add Rule**.
11. In the **Claim rule template** list, select **Pass Through or Filter an Incoming Claim**, and then click **Next**.
12. In the **Claim rule name** box, type **Pass Through Windows Account Name Rule**, and then in the **Incoming claim type** drop-down list, select **Windows account name**.
13. Select **Pass through all claim values**, and then click **Finish**.
14. Click **OK**, and then close the AD FS console.

► **Task 4: Verify access to the A. Datum test application for Trey Research users**

1. On TREY-DC1, open Internet Explorer, and then connect to **https://lon-svr3.adatum.com/adatumtestapp/**.

 **Note:** The sign-in process has changed, and now you must select an authority that can authorize and validate the access request. The **Home Realm Discovery** page (the **Sign In** page) appears, and then you must select an authority.

2. On the **Sign In** page, select **adfs.treyresearch.net**.
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 then connect to **https://lon-svr3.adatum.com/adatumtestapp/** again.

- When prompted for credentials, type **TreyResearch\April** with the password **Pa\$\$w0rd**, and then press Enter. You should be able to access the application.
- Close Internet Explorer.



Note: You will not be prompted for a home realm again. Once users have selected a home realm and a realm authority has authenticated them, the relying party federation server issues a LSRealm cookie. The default lifetime for the cookie is 30 days. Therefore, to sign in multiple times, you should delete that cookie after each sign-in attempt to return to a clean state.

► Task 5: Configure claim rules for the claim provider trust and the relying party trust to allow access only for a specific group

- On TREY-DC1, open the AD FS console, expand **Trust Relationships**, and then click **Relying Party Trusts**.
- Select **Adatum TestApp**, and in the Actions pane, click **Edit Claim Rules**.
- On the Edit Claim Rules for Adatum TestApp window, on the **Issuance Transform Rules** tab, click **Add Rule**.
- On the **Select Rule Template** page, under **Claim rule template**, select **Send Group Membership as a Claim**, and then click **Next**.
- On the **Configure Rule** page, in the **Claim rule name** field, type **Permit Production Group Rule**.
- Next to **Users group**, click **Browse**, type **Production**, and then click **OK**.
- Under **Outgoing claim type**, click **Group**.
- Under **Outgoing claim value**, type **Production**, click **Finish**, and then click **OK**.
- On LON-DC1, if required, open the AD FS Management console.
- In the AD FS console, expand **Trust Relationships**, and then click **Claims Provider Trusts**.
- Select **adfs.treyresearch.net**, and then in the Actions pane, click **Edit Claim Rules**.
- On the **Acceptance Transform Rules** tab, click **Add Rule**.
- On the **Select Rule Template** page, under **Claim rule template**, select **Pass Through or Filter an Incoming Claim**, and then click **Next**.
- On the **Configure Rule** page, in the **Claim rule name** text box, type **Send Production Group Rule**.
- In the **Incoming claim type** drop-down list, click **Group**, and then click **Finish**. Click **Yes**, and then click **OK**.
- In the AD FS console, under **Trust Relationships**, click **Relying Party Trusts**.
- Select the **Adatum Test App**, and in the Actions pane, click **Edit Claim Rules**.
- On the **Issuance Transform Rules** tab, click **Add Rule**.
- Under **Claim Rule Template**, click **Pass Through or Filter an Incoming Claim**, and then click **Next**.
- Under **Claim rule name**, type **Send TreyResearch Group Name Rule**.
- In the **Incoming claim type** drop-down list, click **Group**, and then click **Finish**.
- 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**, ensure that **Permit or Deny Users Based on an Incoming Claim** is selected, and then click **Next**.
25. On the **Configure Rule** page, in the **Claim rule name** text box, type **Permit TreyResearch Production Group Rule**, and then in the **Incoming claim type** drop-down list, select **Group**.
26. In **Incoming claim value**, type **Production**, ensure that **Permit access to users with this incoming claim** is selected, and then click **Finish**.
27. On the **Issuance Authorization Rules** tab, click **Add Rule**.
28. On the **Select Rule Template** page, under **Claim rule template**, ensure that **Permit or Deny Users Based on an Incoming Claim** is selected, and then click **Next**.
29. On the **Configure Rule** page, in the **Claim rule name** field, type **Temp**, and then in the **Incoming Claim Type** drop-down list, select **UPN**.
30. In the **Incoming Claim Value** field, type **@adatum.com**, ensure that **Permit access to users with this incoming claim** is selected, and then click **Finish**.
31. Click the **Temp** rule, and then click **Edit Rule**.
32. In the **Edit Rule –Temp** dialog box, click **View Rule Language**.
33. Press Ctrl+C to copy the rule language to the clipboard, and then click **OK**.
34. Click **Cancel**.
35. Click the **Temp** rule, click **Remove Rule**, and then click **Yes**.
36. On the **Issuance Authorization Rules** tab, click **Add Rule**.
37. On the **Select Rule Template** page, under **Claim rule template**, select **Send Claims Using a Custom Rule**, and then click **Next**.
38. On the **Configure Rule** page, in the **Claim rule name** field, type **ADatum User Access Rule**.
39. 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 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 the claim must match. 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.

40. Click **OK** to close the property page, and then save the changes to the relying party trust.

► **Task 6: Verify restrictions and accessibility to the claims-aware application**

1. On TREY-DC1, open Internet Explorer, and then connect to **https://lon-svr3.adatum.com/adatumtestapp/**.
2. When prompted for credentials, type **TreyResearch\April** with the 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. Reopen Internet Explorer, click the **Settings** icon in the top right corner, and then click **Internet Options**.
5. Under **Browsing History**, click **Delete**, click **Delete** again, and then click **OK**.
6. Connect to **https://lon-svr3.adatum.com/adatumtestapp/**.
7. On the **Sign In** page, click **ads.treyresearch.net**.
8. When prompted for credentials, type **TreyResearch\Morgan** with the 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 should have configured a claims provider trust for Trey Research on Adatum.com and a relying party trust for Adatum on Trey Research. You should also have verified access to the A. Datum claims-aware application and configured the application to restrict access from Trey Research to specific groups.

Exercise 5: Deploy the Web Application Proxy

► Task 1: Configure certificates for the Web Application Proxy server

1. On LON-DC1, on the **Start** screen, type **mmc**, and then press **Enter**.
2. In the Microsoft Management Console, click **File**, and then click **Add/Remove Snap-in**.
3. In the Add or Remove Snap-ins window, in the **Available snap-ins** column, double-click **Certificates**.
4. In the Certificates snap-in window, click **Computer account**, and then click **Next**.
5. In the **Select Computer** window, click **Local Computer (the computer this console is running on)**, and then click **Finish**.
6. In the **Add or remove Snap-ins** window, click **OK**.
7. In the Microsoft Management Console, expand **Certificates (Local Computer)**, expand **Personal**, and then click **Certificates**.
8. Right-click the **adfs.adatum.com** certificate, point to **All Tasks**, and then click **Export**.
9. In the Certificate Export Wizard, click **Next**.
10. On the **Export Private Key** page, click **Yes, export the private key**, and then click **Next**.
11. On the **Export File Format** page, click **Next**.
12. On the **Security** page, select the **Password** check box.
13. In the **Password** and **Confirm password** boxes, type **Pa\$\$w0rd**, and then click **Next**.
14. On the **File to Export** page, in the **File name** box, type **C:\adfs.pfx**, and then click **Next**.
15. On the **Completing the Certificate Export Wizard** page, click **Finish**, and then click **OK** to close the success message.
16. Close the Microsoft Management Console, and then do not save the changes.
17. On LON-SVR2, on the **Start** screen, type **mmc**, and then press **Enter**.
18. In the Microsoft Management Console, click **File**, and then click **Add/Remove Snap-in**.
19. In the Add or Remove Snap-ins window, in the **Available snap-ins** column, double-click **Certificates**.
20. In the Certificates snap-in window, click **Computer account**, and then click **Next**.
21. In the Select Computer window, click **Local Computer (the computer this console is running on)**, and then click **Finish**.
22. In the Add or remove Snap-ins window, click **OK**.
23. In the Microsoft Management Console, expand **Certificates (Local Computer)**, and then click **Personal**.
24. Right-click **Personal**, point to **All Tasks**, and then click **Import**.
25. In the Certificate Import Wizard, click **Next**.
26. On the **File to Import** page, in the **File name** box, type **\\LON-DC1\c\$\adfs.pfx**, and then click **Next**.
27. On the **Private key protection** page, in the **Password** box, type **Pa\$\$w0rd**.
28. Select the **Mark this key as exportable** check box, and then click **Next**.
29. On the **Certificate Store** page, click **Place all certificates in the following store**.
30. In the **Certificate store** dialog box, select **Personal**, and then click **Next**.

31. On the **Completing the Certificate Import Wizard** page, click **Finish**, and then click **OK** to clear the success message.
32. Close the Microsoft Management Console, and then do not save the changes.

► **Task 2: Install the Web Application Proxy role**

1. On LON-SVR2, in Server Manager, click **Manage**, and then click **Add Roles and Features**.
2. In the Add Roles and Features Wizard, on the **Before you begin** page, click **Next**.
3. On the **Select installation type** page, click **Role-based or feature-based installation**, and then click **Next**.
4. On the **Select destination server** page, click **LON-SVR2.Adatum.com**, and then click **Next**.
5. On the **Select server roles** page, select the **Remote Access** check box, and then click **Next**.
6. On the **Select features** page, click **Next**.
7. On the **Remote Access** page, click **Next**.
8. On the **Select role services** page, select **Web Application Proxy**.
9. In the Add Roles and Features Wizard, click **Add Features**.
10. On the **Select role services** page, click **Next**.
11. On the **Confirm installation selections** page, click **Install**.
12. On the **Installation progress** page, click **Close**.

► **Task 3: Configure access to an internal website**

1. On LON-SVR2, click **Tools**, and then click **Remote Access Management**.
2. In the Remote Access Management console, click **Web Application Proxy**.
3. Click **Run the Web Application Proxy Configuration Wizard**.
4. On the **Welcome** page, click **Next**.
5. On the **Federation Server** page, type **adfs.adatum.com** in the **Federation service name** text box.
6. On the **Federation Server** page, type **administrator** in the **User name** text box, type **Pa\$\$w0rd** in the **Password** text box, and then click **Next**.
7. On the **AD FS Proxy Certificate** page, click the **Select a certificate to be used by the AD FS proxy** dropdown menu, and then click **adfs.adatum.com**. Click **Next**.
8. On the **Confirmation** page, click **Configure**.
9. On the **Results** page, click **Close**.
10. In the Remote Access Management Console window, in the Tasks pane, click **Publish**.
11. On the **Welcome** page, click **Next**.
12. On the **Preauthentication** page, click **Pass-through**, and then click **Next**.
13. On the **Publishing Settings** page, type **AdatumTestApp** for the **Name**, type **https://lon-svr3.adatum.com/AdatumTestApp/** for the **External URL**, select **adfs.adatum.com** for the external certificate, and then click **Next**.
14. On the **Confirmation** page, click **Publish**.
15. On the **Results** page, click **Close**.

► **Task 4: Verify access to the internal web site from the client computer**

1. Sign in to LON-CL1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. On the Start screen, type **Notepad**.
3. Right-click **Notepad**, and then click **Run as administrator**.
4. In Notepad, click **File**, and then click **Open**.
5. In the **File name** box, type **C:\Windows\System32\Drivers\etc\hosts**, and then click **Open**.
6. At the bottom of the file, add the following line, click **File**, and then click **Save**:
 - o **172.16.0.13 lon-svr3.adatum.com**
7. Close Notepad.
8. Open Internet Explorer.
9. In Internet Explorer, in the address bar, type **https://lon-svr3.adatum.com/adatumtestapp/**, and then press Enter. On the **There is a problem with this website's security certificate** page, click **Continue to this website (not recommended)**.
10. On the **adfs.adatum.com** page, under **Sign in with one of these accounts**, click **adfs.treyresearch.net**.
11. In the **Windows Security** dialog box, sign in as **TreyResearch\Morgan** with the password **Pa\$\$w0rd**.
12. After the application loads, close Internet Explorer.

► **Task 5: To prepare for the next module**

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

1. On the host computer, start Hyper-V® Manager.
2. In the **Virtual Machines** list, right-click **20414C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20414C-LON-SVR2**, **20414C-LON-SVR3**, **20414C-TREY-DC1**, and **20414C-LON-CL1**.

Results: In this exercise, you should have installed the Web Application Proxy role, configured certificates for the proxy server, configured access to an internal web site, and verified access to the internal web site.

Module 12: Planning and Implementing Data Access for Users and Devices

Lab A: Implementing DAC and Access-Denied Assistance

Exercise 1: Planning and Implementing DAC

► Task 1: Read the supporting documentation

- Read the documentation that the student handbook provides.

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

- Answer the questions in the proposals section of the File Security Strategy document.

Proposals

- How will you design DAC to fulfill the requirements for access control described in the scenario?

Answers may vary, but possible answers may include:

- Only employees who belong to Research Department should access and modify folders that belong to the Research Department.
- Only managers should be able to access files classified as highly confidential.
- Managers should access confidential files only from workstations that belong to the ManagersWKS security group.

To meet these requirements, you must implement claims, resource properties, and file classifications, and then use them together in DAC.

To implement this solution, you should:

1. Create the appropriate claims for users and devices. The user claim uses "department" as its source attribute, while the device claim uses "description" as its source attribute.
2. Configure a Resource Property for the Research Department.
3. Configure Central Access Rules and Central Access Policies to protect resources.
4. Configure a file classification for confidential documents.
5. Apply the Central Access Policy to folders where Research and Managers files are located.

As a solution for users who receive error messages, you should implement access-denied assistance.

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

- Compare your proposals with the ones in the previous lab task.

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

- Be prepared to discuss your proposals with the class.

Results: After completing this exercise, students should have successfully planned and implemented DAC.

Exercise 2: Preparing DAC Deployment

► Task 1: Prepare AD DS for DAC, and review the default claim types

1. On LON-DC1, if necessary, sign in as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. In Windows 8® Server Manager, click **Tools**, and then click **Active Directory Users and Computers**.
3. In the Active Directory Users and Computers console, click and then right-click **Adatum.com**, click **New**, and then click **Organizational Unit**.
4. In the **New Object – Organizational Unit** dialog box, in the **Name** field, type **Test**, and then click **OK**.
5. In the Active Directory Users and Computers console, expand **Adatum.com**, and then click the **Computers** container.
6. Press the Ctrl key. Click the **LON-SVR1** and **LON-CL1** computers, right-click, and then select **Move**.
7. In the Move window, click **Test**, and then click **OK**.
8. Close the Active Directory Users and Computers console.
9. In Server Manager, click **Tools**, and then click **Group Policy Management**.
10. Expand **Forest: Adatum.com**, expand **Domains**, expand **Adatum.com**, and then click the **Group Policy Objects** container.
11. In the Results pane, right-click **Default Domain Controllers Policy**, and then click **Edit**.
12. In the Group Policy Management Editor, under **Computer Configuration**, expand **Policies**, expand **Administrative Templates**, expand **System**, and then click **KDC**.
13. In the right pane, double-click **KDC support for claims, compound authentication and Kerberos armoring**.
14. In the KDC support for claims, compound authentication and Kerberos armoring window, click **Enabled**. In the Options section, on the drop-down list box, click **Supported**, and then click **OK**.
15. Close the Group Policy Management Editor and Group Policy Management Console.
16. On the taskbar, click the **Windows PowerShell** icon.
17. In the Windows PowerShell window, type **gpupdate /force** and then press Enter.
18. After Group Policy updates, close Windows PowerShell®.
19. In Server Manager, click **Tools**, and then click **Active Directory Users and Computers**.
20. In the Active Directory Users and Computers console, right-click **Users**, click **New**, and then click **Group**.
21. For the Group name, type **ManagersWKS** and then click **OK**.
22. Click the **Test** container.
23. Right-click **LON-CL1**, and then click **Properties**.
24. In the **LON-CL1 Properties** dialog box, click the **Member Of** tab, and then click **Add**.
25. On the **Select Groups** page, type **ManagersWKS**. Click **Check Names**, click **OK**, and then click **OK** again.
26. Click the **Managers** organizational unit (OU).
27. Right-click **Aidan Delaney**, and then click **Properties**.

28. In the **Aidan Delaney Properties** dialog box, 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 then click **Properties**.
31. In the **Allie Bellew Properties** dialog box, click the **Organization** tab. Ensure that the **Department** field is populated with the value **Research**, and then click **Cancel**.
32. On LON-DC1, in Server Manager, click **Tools**, and then click **Active Directory Administrative Center**.
33. In the Active Directory Administrative Center console, in the navigation pane, click **Dynamic Access Control**.
34. In the central pane, double-click **Claim Types**.
35. Verify that there are no default claims defined, except AuthenticationSilo.
36. In the navigation pane, click **Dynamic Access Control**, and then double-click **Resource Properties**.
37. Review the default resource properties.



Note: Note that all properties are disabled by default.

38. In the navigation pane, click **Dynamic Access Control**, and then double-click **Resource Property Lists**.
39. In the central pane, right-click **Global Resource Property List**, and then click **Properties**.
40. 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 and devices**

1. In the Active Directory Administrative Center, in the navigation pane, click **Dynamic Access Control**.
2. Double-click **Claim Types**.
3. In the Tasks pane, click **New**, and then click **Claim Type**.
4. In the **Create Claim Type** dialog box, in the Source Attribute section, click **Department** in the list. In the **Display name** text box, type **Company Department**.
5. Select the **User** and the **Computer** check boxes, and then click **OK**.
6. In the Active Directory Administrative Center, in the Tasks pane, click **New**, and then click **Claim Type**.
7. On the **Create Claim Type** page, in the Source Attribute section, click **description**.
8. Clear the **User** check box, select the **Computer** check box, and then click **OK**.


► **Task 3: Configure resource properties for files**

1. In the Active Directory Administrative Center, click **Dynamic Access Control**.
2. In the central pane, double-click **Resource Properties**.
3. In the **Resource Properties** list, locate and right-click **Department**, and then click **Enable**.
4. In the **Resource Properties** list, locate and right-click **Confidentiality**, and then click **Enable**.
5. Ensure that the **Department** and **Confidentiality** properties are enabled in the list.

6. Double-click **Department**.
7. Scroll down to the Suggested Values section, and then click **Add**.
8. In the **Add a suggested value** window, type **Research** in the **Value** and **Display name** text boxes, and then click **OK** twice.
9. Click **Dynamic Access Control**, and then double-click **Resource Property Lists**.
10. In the central pane, double-click **Global Resource Property List**.
11. Ensure that both **Department** and **Confidentiality** display in the **Resource Properties** list, and then click **Cancel**.
12. Close the Active Directory Administrative Center.

► **Task 4: Classify files and folders**

1. Switch to LON-SVR1.
2. In Server Manager, click **Add roles and features**.
3. In the Add Roles and Features Wizard, click **Next** three times.
4. On the **Select server roles** page, expand **File and Storage Services (Installed)**, expand **File and iSCSI Services**, and then select the **File Server Resource Manager** check box.
5. When prompted, click **Add Features**.
6. Click **Next** twice, and then click **Install**. After installation completes, click **Close**.
7. On the desktop, on the taskbar, click the **File Explorer** icon.
8. In File Explorer, in the address bar, type **C:**, and then press Enter.
9. In File Explorer, right-click an area of free space, point to **New**, and then click **Folder**.
10. Type **Docs** and then press Enter.
11. Double-click the **Docs** folder.
12. Right-click an area of free space, point to **New**, and then click **Text Document**.
13. Type **Doc1** and then press Enter.
14. Double-click **Doc1**.
15. In Notepad, type **This is a secret document**.
16. Close the file, and, when prompted, click **Save**.
17. In File Explorer, in the Docs folder, right-click an area of free space, point to **New**, and then click **Text Document**.
18. Type **Doc2**, and then press Enter.
19. Double-click **Doc2**.
20. In Notepad, type **This is a secret document**.
21. Close the file, and, when prompted, click **Save**.
22. In File Explorer, in the Docs folder, right-click an area of free space, point to **New**, and then click **Text Document**.
23. In Notepad, type **Doc3**, and then press Enter.
24. Double-click **Doc3**.

25. Type **This is a document**.
 26. Close the file, and, when prompted, click **Save**.
 27. In the File Explorer address bar, type **C:**, and then press Enter.
 28. Right-click **Docs**, point to **Share with**, and then click **Specific people**.
 29. In the **File Sharing** dialog box, in the text box, type **Authenticated Users**, and then click **Add**.
 30. In the **Name** list, click **Authenticated Users**, and then, in **Permission Level**, click **Read/Write**.
 31. Click **Share**, and then click **Done**.
 32. In Server Manager, click **Tools**, and then click **File Server Resource Manager**.
 33. In the File Server Resource Manager console, expand **Classification Management**.
 34. Right-click **Classification Properties**, and then click **Refresh**.
 35. Verify that the **Confidentiality** and **Department** properties display in the list.
 36. Click **Classification Rules**.
 37. In the Actions pane, click **Create Classification Rule**.
 38. In the **Create Classification Rule** window, in the **Rule name** text box, type **Set Confidentiality**.
 39. Click the **Scope** tab, and then click **Add**.
 40. In the **Browse For Folder** dialog box, expand **Local Disk (C:)**, click the **Docs** folder, and then click **OK**.
 41. Click the **Classification** tab.
 42. Ensure that the following settings are set, and then click **Configure**:
 - Classification method: **Content Classifier**
 - Property: **Confidentiality**
 - Value: **High**
 43. In the **Classification Parameters** dialog box, click the **Regular expression** drop-down list box, and then click **String**.
 44. In the **Expression** field, next to the word String, type **secret**, and then click **OK**.
 45. Click the **Evaluation Type** tab. Select the **Re-evaluate existing property values** check box, click **Overwrite the existing value**, and then click **OK**.
 46. In the File Server Resource Manager, in the Actions pane, click **Run Classification with all rules now**.
 47. Click **Wait for classification to complete**, and then click **OK**.
 48. After the classification completes, a report displays. Verify that two files were classified.
-  **Note:** You can see that the two files were classified in the Report Totals section of the report.
49. Close the report.
 50. Switch to File Explorer, expand drive **C**, and then click the **Docs** folder.
 51. Right-click **Doc1**, and then click **Properties**.

52. In the **Doc1.txt Properties** dialog box, click the **Classification** tab, and then verify that **Confidentiality** is set to **High**.
53. Repeat steps 51 and 52 on files **Doc2** and **Doc3**.



Note: Doc2.txt should have the same confidentiality as Doc1.txt, while Doc3.txt should have no value. This is because only Doc1 and Doc2 have the word *secret* in their contents.


54. In the File Explorer address bar, type **C:** and then press Enter.
55. In File Explorer, right-click an area of free space, point to **New**, and then click **Folder**.
56. Type **Research**, and then press Enter.
57. In File Explorer, double-click **Research**, right-click an area of free space, point to **New**, and then click **Text Document**.
58. Type **Research1**, and then press Enter.
59. Double-click **Research1**.
60. In Notepad, type **This is a research document** and then close the file. When prompted, click **Save**.
61. In the File Explorer address bar, type **C:** and then press Enter.
62. Right-click **Research**, point to **Share with**, and then click **Specific people**.
63. In the **File Sharing** dialog box, in the text box, type **Authenticated Users** and then click **Add**.
64. In the **Name** list, click **Authenticated Users**, and then, in **Permission Level**, click **Read/Write**.
65. Click **Share**, and then click **Done**.
66. In File Explorer, right-click the **Research** folder, and then click **Properties**.
67. In the **Research Properties** dialog box, click the **Classification** tab. Click **Department**, in the Value section, click **Research**, click **Apply**, and then click **OK**.

Results: After completing this exercise, students will have prepared for DAC deployment.


Exercise 3: Implementing DAC

► Task 1: Configure Central Access Policy rules

1. Switch to LON-DC1.
2. In Server Manager, click **Tools**, and then click **Active Directory Administrative Center**.
3. In the Active Directory Administrative Center console, in the Navigation pane, click **Dynamic Access Control**.
4. Double-click **Central Access Rules**.
5. In the Tasks pane, click **New**, and then click **Central Access Rule**.
6. In the **Create Central Access Rule** dialog box, in the **Name** box, type **Department Match**. In the **Target Resources** section, click **Edit**.
7. In the **Central Access Rule** dialog box, click **Add a condition**.
8. Set a condition as follows: **Resource-Department-Equals-Value-Research**. Click **OK**.
9. In the **Permissions** section, click **Use following permissions as current permissions**.
10. In the **Permissions** section, click **Edit**.
11. Click **Administrators (ADATUM\Administrators)**, and then click **Remove**.
12. In the **Advanced Security Settings for Permissions** dialog box, click **Add**.
13. In the **Permission Entry for Permissions** dialog box, click **Select a principal**.
14. In the **Select User, Computer, Service Account or Group** dialog box, type **Authenticated Users**, click **Check Names**, and then click **OK**.
15. In the Basic permissions section, click **Modify, Read and Execute, Read**, and **Write**, and then click **Add a condition**.
16. Click the **Group** drop-down list box, and then click **Company Department**.
17. On the **Value** drop-down list box, click **Resource**.
18. In the last drop-down list box, click **Department**. Click **OK** three times.

 **Note:** As a result, you should have the following expression: **User-Company Department-Equals-Resource-Department**.

19. In the Tasks pane, click **New**, and then click **Central Access Rule**.
20. For the name of the rule, type **Access Confidential Docs**.
21. In the **Target Resources** section, click **Edit**.
22. In the **Central Access Rule** window, click **Add a condition**.
23. In the last drop-down list box, click **High**, and then click **OK**.

 **Note:** As a result, you should have the following expression: **Resource-Confidentiality-Equals-Value-High**.

24. In the Permissions section, click **Use following permissions as current permissions**, and then click **Edit**.
 25. Click **Administrators (ADATUM\Administrators)**, and then click **Remove**.
 26. In the **Advanced Security Settings for Permissions** dialog box, click **Add**.
 27. In the **Permission Entry for Permissions** dialog box, click **Select a principal**.
 28. In the **Select User, Computer, Service Account or Group** dialog box, type **Authenticated Users**, click **Check Names**, and then click **OK**.
 29. In the Basic permissions section, click **Modify, Read and Execute, Read**, and **Write**, and then click **Add a condition**.
 30. Click **Add items**.
 31. In the **Select User, Computer, Service Account or Group** dialog box, type **Managers**, and then click **OK**.
 32. In the **Multiple Names Found** dialog box, click **Managers**, and then click **OK**.
 33. Click **Add a condition**, and then click the lower **Add items**.
 34. In the **Select User, Computer, Service Account or Group** dialog box, type **Managers**, and then click **OK**.
 35. In the **Multiple Names Found** dialog box, click **Managers**, and then click **OK**.
 36. Set the second condition to: **Device-Group-Member of each-Value-**.
 37. Click lower **Add items**.
 38. In the **Select Computer or Group** dialog box, type **Managers**, and then click **OK**.
 39. In the **Multiple Names Found** dialog box, click **ManagersWKS**, and then click **OK**, and then click **OK** three times.
- **Task 2: Create and publish the Central Access Policy**
1. On LON-DC1, in Active Directory Administrative Center, click **Dynamic Access Control**, and then double-click **Central Access Policies**.
 2. In the Tasks pane, click **New**, and then click **Central Access Policy**.
 3. In the **Name** text box, type **Protect confidential docs**, and then click **Add**.
 4. Click the **Access Confidential Docs** rule, and then click the **Move (>>)** icon. Click **OK** twice.
 5. In the Tasks pane, click **New**, and then click **Central Access Policy**.
 6. In the **Name** box, type **Department Match**, and then click **Add**.
 7. Click the **Department Match** rule, click the **More (>>)** icon, and then click **OK** twice.
 8. Switch to Server Manager, click **Tools**, and then click **Group Policy Management**.
 9. In the Group Policy Management, under Domains, expand **Adatum.com**, right-click **Test**, and then click **Create a GPO in this domain, and link it here**.
 10. Type **DAC Policy**, and then click **OK**.
 11. Expand Test, right-click **DAC Policy**, and then click **Edit**.

12. In the Group Policy Management Editor, 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**.
13. Click both **Department Match** and **Protect confidential docs**, click **Add**, and then click **OK**.
14. Close the Group Policy Management Editor.
15. Close the Group Policy Management Console.

► **Task 3: Apply a Central Access Policy**

1. Switch to LON-SVR1.
2. On the taskbar, click the **Windows PowerShell** icon.
3. In the Windows PowerShell window, type **gpupdate /force**, and then press Enter.
4. Close the Windows PowerShell window.
5. Switch to File Explorer.
6. In the File Explorer window, in the address bar, type **C:**, and then press Enter.
7. Right-click the **Docs** folder, and then click **Properties**.
8. In the **Docs Properties** dialog box, click the **Security** tab, and then click **Advanced**.
9. In the Advanced Security Settings for Docs window, click the **Central Policy** tab, and then click **Change**.
10. In the drop-down list box, click **Protect confidential docs**, and then click **OK** twice.
11. Right-click the Research folder, and then click **Properties**.
12. In the **Research Properties** dialog box, click the **Security** tab, and then click **Advanced**.
13. In the Advanced Security Settings for Research window, click the **Central Policy** tab, and then click **Change**.
14. In drop-down list box, click **Department Match**, and then click **OK** twice.

Results: After completing this exercise, students should have implemented DAC.

Exercise 4: Validating and Remediating DAC

► Task 1: Configure access-denied remediation settings

1. Switch to LON-DC1.
2. In Server Manager, click **Tools**, and then click **Group Policy Management**.
3. Expand **Forest: Adatum.com**, expand **Domains**, expand **Adatum.com**, and then click **Group Policy Objects**.
4. Right-click **DAC Policy**, and then click **Edit**.
5. Under Computer Configuration, expand **Policies**, expand **Administrative Templates**, expand **System**, and then click **Access-Denied Assistance**.
6. In the right pane, double-click **Customize Message for Access Denied errors**.
7. In the Customize Message for Access Denied errors window, click **Enabled**.
8. In the **Display the following message to users who are denied access** text box, type **You are denied access because of permission policy. Please request access**.
9. Select the **Enable users to request assistance** check box.
10. Review the other options without making any changes, and then click **OK**.
11. In the right pane of the Group Policy Management Editor, double-click **Enable access-denied assistance on client for all file types**, click **Enabled**, and then click **OK**.
12. Close the Group Policy Management Editor, and then close the Group Policy Management Console.
13. Switch to LON-SVR1.
14. On LON-SVR1, on the taskbar, click the **Windows PowerShell** icon.
15. In Windows PowerShell, type **gpupdate /force**, and then press Enter.


► Task 2: Verify DAC functionality

1. Restart LON-CL1.
2. Sign in to LON-CL1 as **Adatum\April** with the password **Pa\$\$w0rd**.
3. Click the **Desktop**, and then, on the taskbar, click the **File Explorer** icon.
4. In the File Explorer address bar, type **\\LON-SVR1\Docs**, and then press Enter.
5. Try to open **Doc3**. You should be able to open that document.
6. In the File Explorer address bar, type **\\LON-SVR1\Research**, and then press Enter.




Note: You should be unable to access this folder.

7. Click **Request Assistance**. Review the options for sending a message, and then click **Close**.
8. Sign out of LON-CL1.
9. Sign back in to LON-CL1 as **Adatum\Allie** with the password **Pa\$\$w0rd**.
10. Click the **Desktop** tile, and then, on the taskbar, click the **File Explorer** icon.
11. In File Explorer, in the address bar, type **\\LON-SVR1\Research**, and then press Enter.

 **Note:** You should be able to access this folder and open documents inside, because Allie is a member of the Research Department.

12. Sign out of LON-CL1.
13. Sign back in to LON-CL1 as **Adatum\Aidan** with the password **Pa\$\$w0rd**.
14. Click the **Desktop**, and then, on the taskbar, click the **File Explorer** icon.
15. In File Explorer, in the address bar, type **\\LON-SVR1\Docs**, and then press Enter.

 **Note:** You should be able to access this folder and open documents inside, because Aidan is a member of the Managers Department, and he is accessing the documents from a computer that is a member of ManagersWKS group.

► Task 3: View effective permissions

1. On LON-SVR1, switch to **File Explorer**.
2. In File Explorer, in the address bar, type **C:**, and then press Enter.
3. In File Explorer, right-click the **Research** folder, and then click **Properties**.
4. In the **Research Properties** dialog box, click the **Security** tab, click **Advanced**, and then click **Effective Access**.
5. Click **select a user**.
6. In the Select User, Computer, Service Account, or Group window, type **April**, click **Check Names**, and then click **OK**.
7. Click **View effective access**.
8. Review the results. April should not have access to this folder.
9. Click **Include a user claim**.
10. On the drop-down list box, click **Company Department**, and then, in the **Value** text box, type **Research**.
11. Click **View Effective access**. April should now have access.
12. Close all windows.

► Task 4: Prepare for the next lab

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

1. On the host computer, start Microsoft® Hyper-V® Manager.
2. On the **Virtual Machines** list, right-click **20414C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machines** dialog box, click **Revert**.
4. Repeat steps two and three for 20414C-LON-SVR1 and 20414C-LON-CL1.

Results: After completing this exercise, students should have validated functionality of DAC.

Lab B: Implementing Work Folders

Exercise 1: Preparing and Implementing an Infrastructure for Work Folders

► Task 1: Installing Work Folders functionality and configuring a Secure Sockets Layer certificate

1. Sign in to LON-SVR1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. On LON-SVR1, in Server Manager, click **Add roles and features**.
3. On the **Before You Begin** page, click **Next**.
4. On the **Select installation type** page, ensure that **Role - based or feature - based installation** is selected, and then click **Next**.
5. On the **Select destination server** page, click **Next**.
6. On the **Select server roles** page, expand **File and Storage Services**, expand **File and iSCSI Services**, and then select **Work Folders**.
7. In the **Add features that are required for Work Folders** dialog box, note the features, and then click **Add Features**.
8. On the **Select server roles** page, click **Next**.
9. On the **Select features** page, click **Next**.
10. On the **Confirm installation selection** pages, click **Install**.
11. When the installation finishes, click **Close**.
12. In Server Manager on LON-SVR1, click **Tools**, and then click **Internet Information Services (IIS) Manager**.
13. In the **Internet Information Services (IIS) Manager** console, click **LON-SVR1**, click **No** if prompted, and then double-click **Server Certificates** in the middle pane.
14. In the Actions pane, click **Create Domain Certificate**.
15. In the Create Certificate window, fill in the text fields as follows:
 - Common name: **lon-svr1.adatum.com**
 - Organization: **Adatum**
 - Organizational Unit: **IT**
 - City/locality : **Seattle**
 - State/province : **WA**
 - Country/region: **US**
16. Click **Next**.
17. On the **Online Certification Authority** page, click **Select**.
18. In the Select Certification Authority window, select **AdatumCA**, and then click **OK**.
19. In the **Friendly name** text box, type **lon-svr1.adatum.com**, and then click **Finish**. Wait until the certificate is issued.
20. In the **IIS** console, expand **Sites**, and then click **Default Web Site**.

21. In the Actions pane, click **Bindings**.
22. In the Site Bindings window, click **Add...**
23. In the Add Site Binding window, under **Type**, select **https**. In the **SSL certificate** drop-down list, select **lon-svr1.adatum.com**.
24. Click **OK**, and then click **Close**.
25. Close **Internet Information Services (IIS) Manager**.

► **Task 2: Provision a share for Work Folders**

1. On LON-SVR1, in Server Manager, in the navigation pane, click **File and Storage Services**.
2. Click **Shares**, and, in the SHARES area, click **Tasks**, and then select **New Share...**
3. In the New Share Wizard, on the **Select the profile for this share** page, ensure that **SMB Share – Quick** is selected, and then click **Next**.
4. On the **Select the server and path for this share** page, accept the defaults, and then click **Next**.
5. On the **Specify share name** page, in the **Share name** field, type **WF-Share**, and then click **Next**.
6. On the **Configure share settings** page, select **Enable access - based enumeration**, leave the other settings as their default settings, and then click **Next**.
7. On the **Specify permissions to control access** page, note the default settings, and then click **Next**.
8. On the **Confirm selections** page, click **Create**.
9. On the **View results** page, click **Close**.

► **Task 3: Configure and implement Work Folders**

1. On LON-SVR1, in Server Manager, expand **File and Storage Services**, and then click **Work Folders**.
2. In the **WORK FOLDERS** tile, click **Tasks**, and then click **New Sync Share...**
3. In the New Sync Share Wizard, on the **Before You Begin** page, click **Next**.
4. On the **Select the server and path** page, select **Select by file share**, ensure that the share you created in the previous task, WF-Share, is highlighted, and then click **Next**.
5. On the **Specify the structure for user folders** page, accept the default selection, User alias, and then click **Next**.
6. On the **Enter the sync share name** page, accept the default, and then click **Next**.
7. On the **Grant sync access to groups** page, note the default selection to disable inherited permissions and grant users exclusive access, and then click **Add**.
8. In the **Select User or Group** dialog box, in the **Enter the object names to select** field, type **WFSync**, click **Check Names**, and then click **OK**.
9. On the **Grant sync access to groups** page, click **Next**.
10. On the **Specify device policies** page, clear both check boxes, and then click **Next**. (Note: In a production environment, you should apply device policies).
11. On the **Confirm selections** page, click **Create**.
12. On the **View results** page, click **Close**.
13. Switch to LON-DC1.

14. Open **Server Manager**, click **Tools**, and then click **Group Policy Management**. (Note: If the console does not open, and a dialog box appears that tells you Group Policy Management is loading, close the dialog box and try to open the Group Policy Management console again).
15. Expand **Forest: Adatum.com-Domains-Adatum.com**, and then click **Group Policy Objects**. Right-click **Group Policy Objects**, and then click **New**.
16. In the New GPO window, type **Work Folders GPO** in the **Name** field, and then click **OK**.
17. Right-click **Work Folders GPO**, and then click **Edit**.
18. In the Group Policy Management Editor, expand **User Configuration/Policies/Administrative Templates/Windows Components**, and then click **Work Folders**.
19. Double-click **Specify Work Folders settings** in the details pane, and then, in the **Specify Work Folders settings** dialog box, click **Enabled**.
20. In the **Work Folders URL** text box, type **https://lon-svr1.adatum.com**, and then select **Force automatic setup**.
21. Click **OK** to close the **Specify Work Folders** settings dialog box, and then close the Group Policy Management Editor.
22. In the Group Policy Management Console, right-click the **Adatum.com** domain object, and then select **Link an Existing GPO...**
23. In the Select GPO window, select **Work Folders GPO**, and then click **OK**.
24. Close the Group Policy Management Console.

Results: After completing this exercise, students will have configured a Work Folders server infrastructure.

Exercise 2: Configuring AD FS and Web Application Proxy for Work Folders Publishing

► Task 1: Install AD FS

1. On LON-DC1, in Server Manager, click **Tools**, and then click **DNS**.
2. In the DNS Manager, expand **LON-DC1**, expand **Forward Lookup Zones**, and then click **Adatum.com**.
3. Right-click **Adatum.com**, and then click **New Host (A or AAAA)**.
4. In the New Host window, in the **Name** box, type **adfs**.
5. In the **IP address** box, type **172.16.0.10**, and then click **Add Host**.
6. In the DNS window, click **OK**.
7. Click **Done**, and then close the DNS Manager.
8. On LON-DC1, in Server Manager, click **Manager**, and then click **Add Roles and Features**.
9. In the Add Roles and Features Wizard, on the **Before You Begin** page, click **Next**.
10. On the **Select installation type** page, click **Role-based or feature-based installation**, and then click **Next**.
11. On the **Select destination server** page, click **Select a server from the server pool**, click **LON-DC1.Adatum.com**, and then click **Next**.
12. On the **Select server roles** page, select the **Active Directory Federation Services** check box, and then click **Next**.
13. On the **Select features** page, click **Next**.
14. On the **Active Directory Federation Services (AD FS)** page, click **Next**.
15. On the **Confirm installation selections** page, click **Install**.
16. When the installation is complete, click **Close**.
17. On LON-DC1, on the taskbar, click **Windows PowerShell**.
18. At the Windows PowerShell command-line interface command prompt, type **Add-KdsRootKey –EffectiveTime (Get-Date).AddHours(-10)**, and then press Enter. Wait until you see that you get Guid as response.
19. Close Windows PowerShell.

► Task 2: Configure AD FS

1. On LON-DC1, in Server Manager, click the **Notifications** icon, and then click **Configure the federation service on this server**.
2. In the Active Directory Federation Services Configuration Wizard, on the **Welcome** page, click **Create the first federation server in a federation server farm**, and then click **Next**.
3. On the **Connect to Active Directory Domain Services** page, click **Next** to use **Adatum\Administrator** to perform the configuration.
4. On the **Specify Service Properties** page, in the **SSL Certificate** box, select **ADFS.adatum.com**.
5. In the **Federation Service Display Name** box, type **A. Datum Corporation**, and then click **Next**.
6. On the **Specify Service Account** page, click **Create a Group Managed Service Account**.
7. In the **Account Name** box, type **ADFS**, and then click **Next**.

8. On the **Specify Configuration Database** page, click **Create a database on this server using Windows Internal Database**, and then click **Next**.
9. On the **Review Options** page, click **Next**.
10. On the **Pre-requisite Checks** page, click **Configure**.
11. On the **Results** page, click **Close**.

► Task 3: Install Web Application Proxy

1. On LON-SVR2, in Server Manager, click **Manage**, and then click **Add Roles and Features**.
2. In the Add Roles and Features Wizard, on the **Before You Begin** page, click **Next**.
3. On the **Select installation type** page, click **Role-based or feature-based installation**, and then click **Next**.
4. On the **Select destination server** page, click **LON-SVR2.Adatum.com**, and then click **Next**.
5. On the **Select server roles** page, select the **Remote Access** check box, and then click **Next**.
6. On the **Select features** page, click **Next**.
7. On the **Remote Access** page, click **Next**.
8. On the **Select role services** page, select **Web Application Proxy**.
9. In the Add Roles and Features Wizard, click **Add Features**.
10. On the **Select role services** page, click **Next**.
11. On the **Confirm installation selections** page, click **Install**.
12. On the **Installation progress** page, click **Close**.

► Task 4: Configure Web Application Proxy

1. On LON-DC1, on the Start screen, type **mmc**, and then press Enter.
2. In the Microsoft Management Console, click **File**, and then click **Add/Remove Snap-in**.
3. In the Add or Remove Snap-ins window, in the **Available snap-ins** column, double-click **Certificates**.
4. In the Certificates snap-in window, click **Computer account**, and then click **Next**.
5. In the Select Computer window, click **Local Computer (the computer this console is running on)**, and then click **Finish**.
6. In the Add or remove Snap-ins window, click **OK**.
7. In the Microsoft Management Console, expand **Certificates (Local Computer)**, expand **Personal**, and then click **Certificates**.
8. Right-click **ADFS.adatum.com**, point to **All Tasks**, and then click **Export**.
9. In the Certificate Export Wizard, click **Next**.
10. On the **Export Private Key** page, click **Yes, export the private key**, and then click **Next**.
11. On the **Export File Format** page, click **Next**.
12. On the **Security** page, select the **Password** check box.
13. In the **Password** and **Confirm password** boxes, type **Pa\$\$w0rd**, and then click **Next**.

14. On the **File to Export** page, in the **File name** box, type **C:\adfs.pfx**, and then click **Next**.
15. On the **Completing the Certificate Export Wizard** page, click **Finish**, and then click **OK** to close the success message.
16. Close the Microsoft Management Console, and then do not save the changes.
17. On LON-SVR2, on the Start screen, type **mmc**, and then press Enter.
18. In the Microsoft Management Console, click **File**, and then click **Add/Remove Snap-in**.
19. In the Add or Remove Snap-ins window, in the **Available snap-ins** column, double-click **Certificates**.
20. In the Certificates snap-in window, click **Computer account**, and then click **Next**.
21. In the Select Computer window, click **Local Computer (the computer this console is running on)**, and then click **Finish**.
22. In the Add or remove Snap-ins window, click **OK**.
23. In the Microsoft Management Console, expand **Certificates (Local Computer)**, and then click **Personal**.
24. Right-click **Personal**, point to **All Tasks**, and then click **Import**.
25. In the Certificate Import Wizard, click **Next**.
26. On the **File to Import** page, in the **File name** box, type **\\LON-DC1\c\$\adfs.pfx**, and then click **Next**.
27. On the **Private key protection** page, in the **Password** box, type **Pa\$\$w0rd**.
28. Select the **Mark this key as exportable** check box, and then click **Next**.
29. On the **Certificate Store** page, click **Place all certificates in the following store**.
30. In the **Certificate store** box, ensure that **Personal** is selected, and then click **Next**.
31. On the **Completing the Certificate Import Wizard** page, click **Finish**, and then click **OK** to clear the success message.
32. Close the Microsoft Management Console, and then do not save the changes.
33. On LON-SVR2, in Server Manager, click the **Notifications** icon, and then click **Open the Web Application Proxy Wizard**.
34. In the Web Application Proxy Wizard, on the **Welcome** page, click **Next**.
35. On the **Federation Server** page, enter the following, and then click **Next**:
 - Federation service name: **adfs.adatum.com**
 - User name: **Adatum\Administrator**
 - Password: **Pa\$\$w0rd**
36. On the **AD FS Proxy Certificate** page, in the **Select a certificate to be used by the AD FS proxy** box, select **ADFS.adatum.com**, and then click **Next**.
37. On the **Confirmation** page, click **Configure**.
38. On the **Results** page, click **Close**.
39. The Remote Access Management Console opens automatically. Leave it open for the next task.

► **Task 5: Publish Work Folders through Web Application Proxy**


1. On LON-SVR2, in the Remote Access Management Console, in the Tasks pane, click **Publish**.
2. In the Publish New Application Wizard, on the **Welcome** page, click **Next**.
3. On the **Preauthentication** page, click **Pass-through**, and then click **Next**.
4. On the **Publishing Settings** page, in the **Name** box, type **A. Datum Work Folders**.
5. In the **External URL** box, type **https://wf.adatum.com/**.
6. In the **External certificate** box, select **ADFS.adatum.com**.
7. In the **Backend server URL** box, type **https://lon-svr1.adatum.com/**, and then click **Next**.
8. On the **Confirmation** page, click **Publish**.
9. On the **Results** page, click **Close**.

Results: After completing this exercise, students will have configured Active Directory® Federation Services (AD FS) and Web Application Proxy services for Work Folders publishing.

Exercise 3: Validating Work Folders Functionality

► Task 1: Validating Work Folders from a domain-joined client device

1. Sign in to LON-CL1 as **Adatum\Aidan** with the password **Pa\$\$w0rd**.
2. On the Start screen, start typing **PowerShell**, and then click the **Windows PowerShell** icon in the Search pane.
3. At the Windows PowerShell command prompt, type **gpupdate /force**, and then press Enter. (Note: If you receive an error message when refreshing group policy, restart the LON-CL1 machine and try again).
4. Open **File Explorer** from the taskbar.
5. Verify that the **Work Folders** folder has been created.

 **Note:** The presence of the Work Folders folder indicates that the Work Folders configuration through Group Policy is successful.

6. In **File Explorer**, create a few, two or three, text files and save them in the **Work Folders** folder.

 **Note:** File Explorer displays the synchronization status of the files in the Work Folders folder.

7. Right-click **Windows** on the taskbar, and then click **Control Panel**. In the Control Panel window, in View, select **Large icons** by the drop-down menu.
8. In the Control Panel window, click **Work Folders**.
9. Ensure that Work Folders are configured and working. (Note: there should be no errors or warnings in this window).
10. Close the Work Folders window, and then close the Control Panel window.

► Task 2: Validating Work Folders from a non-domain joined device

1. Sign in to LON-CL2 as **Delaney** with password **Pa\$\$w0rd**. Note that you are signing in to a device that is non-domain-joined.
2. On LON-CL2, on the Start screen, type **Notepad**.
3. Right-click **Notepad**, and then click **Run as administrator**.
4. In the User Account Control window, when prompted, type **Pa\$\$w0rd** in the **Password** field, and click **Yes**.
5. In Notepad, click **File**, and then click **Open**.
6. In the **File name** box, type **C:\Windows\System32\Drivers\etc\hosts**, and then click **Open**.
7. At the bottom of the file, add the following line, click **File**, and then click **Save**:
 - **131.107.0.2 wf.adatum.com**
8. Close Notepad.
9. On LON-CL2, on the Start screen, type **powershell**, and then open **Windows PowerShell**.
10. In the Windows PowerShell window, type **ping lon-svr1.adatum.com** and press Enter. Verify that the ping cannot find a host. (Note: This verifies that you cannot access the server that hosts Work Folders by using its internal name).
11. Type **ping wf.adatum.com** and press Enter. Verify that you receive a reply from 131.107.0.2. (Note: This verifies that you can access the Web Application Proxy server external interface and the name used to publish Work Folders).

12. Close the Windows PowerShell window.
13. Right-click **Windows** on the taskbar, and then click **Control Panel**.
14. In the **Control Panel**, click **System and Security**, and then click **Work Folders**.
15. Click **Set up Work Folders**.
16. On the **Enter your work email address** page, click **Enter a Work Folders URL instead**, then type **https://wf.adatum.com**, and then click **Next**.
17. In the Windows Security window, when prompted for user name and password type **Aidan@adatum.com** for the User name and **Pa\$\$w0rd** for Password. Select the **Remember my credentials** check box, and then click **OK**.
18. On the **Introducing Work Folders** page, click **Next**.
19. On the **Security policies** page, select the **I accept these policies on my PC** check box, and then click **Set up Work Folders**.
20. On the **Work Folders has started syncing with this PC** page, click **Close**.
21. The Work Folders folder will open. Verify that you see files that you created in the previous task for the LON-CL1 virtual machine.

► **Task 3: Validate Work Folders data synchronization**

1. On LON-CL2, in the Work Folders folder, create a few, two or three, new text document files.
2. Switch to LON-CL1.
3. Open **File Explorer** and navigate to the **Work Folders** folder. Check if the new files are present. Most likely, they will not appear immediately.
4. Create one new text file in **Work Folders** on LON-CL1.
5. Wait three to four minutes.
6. Refresh the content in Work Folders on LON-CL1. Verify that all files from both LON-CL1 and LON-CL2 are present.
7. Switch to LON-CL2. Make some changes in one of the documents. Wait three to four minutes, and then refresh the content in the Work Folders folder.
8. Ensure that all files appear in the folder and that they are the same as on LON-CL1.

► **Task 4: Prepare for the next module**

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

1. On the host computer, start Hyper-V Manager.
2. On the Virtual Machines list, right-click **20414C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machines** dialog box, click **Revert**.
4. Repeat steps two and three for 20414C-LON-SVR1, 20414C-LON-SVR2, 20414C-LON-CL1, and 20414C-LON-CL2.

Results: After completing this exercise, students will have validated Work Folders functionality from both domain-joined devices and devices that are non-domain joined.

Module 13: Planning and Implementing an Information Rights Management Infrastructure

Lab: Planning and Implementing an AD RMS Infrastructure

Exercise 1: Planning the AD RMS Deployment

► Task 1: Read the supporting documentation

Supporting Documentation

----- Original Message -----

From: Charlotte Weiss [Charlotte@contoso.com]

Sent: 03 Feb 2013 08:45

To: Ed@contoso.com

Subject: AD RMS

Ed,

Because of the highly confidential nature of the work that the A. Datum research team at performs, the security team at A. Datum Corporation wants to implement additional security for some of the research team's documents. In particular, the security team wants to ensure that users inside and outside the organization cannot share confidential documents with any unauthorized users. You must plan and implement an AD RMS solution that will provide the level of protection requested by the security team.

Please create a plan to install AD RMS at both A. Datum and Trey Research that incorporates the rest of the action items.

Thank you,
Charlotte

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

Answer the questions in the A. Datum Information Rights Management (IRM) Plan: AD RMS document, shown below.

A. Datum IRM Plan: AD RMS	
Document Reference Number: GW00612	
Document Author	Charlotte Weiss
Date	6 th February
<ul style="list-style-type: none"> • Requirements Overview • Design an AD RMS deployment for the A. Datum and the Trey Research companies. • The goal of the AD RMS deployment is to protect information, no matter where it goes. Once you add AD RMS protection to a digital file, the protection stays with the file. By default, only the content owner is able to remove the protection from the file. The owner grants rights to other users to perform actions on the content, such as the ability to view, copy, or print the file: • AD RMS must automatically designate as confidential and then protect all documents in the ConfidentialResearch folder. 	

A. Datum IRM Plan: AD RMS

- Only members from the Research and Managers groups should be able to access the documents.
- Configure the integration of this AD RMS deployment with Dynamic Access Control.
- A. Datum employees must be able to share the AD RMS–protected content with users at Trey Research.
- The A. Datum security team wants to ensure that you do not secure any AD RMS–protected content in such a way so that no one can access the information.

Additional Information

We need to cut down on the total number of servers required. Consider using the internal database instead of a full Microsoft® SQL Server® deployment. We have authorization to deploy AD RMS on the existing domain controllers.

Proposals

Question: How many AD RMS clusters do you need to deploy to satisfy the security requirements for both companies?

Answer: There are two separate forests in this configuration, so you will need at least 2 AD RMS clusters, one for each forest.

Question: What database solution will you deploy?

Answer: Each AD RMS cluster will need a SQL Server database solution. You could use the internal database provided as a feature in Windows Server® 2012, but you will not be able to scale up with such a deployment.

Question: What service accounts, if any, do you need to create?

Answer: We will need only one account to act as the service account. We cannot use the administrator's account, because the configuration of the AD RMS console requires an account other than the installing account to activate the database.

Question: What Secure Sockets Layer (SSL) certificate requirements do we have? How do we satisfy them in both forests?

Answer: We can use self-signed certificates. We can place the certificates in each computer's Trusted Root Certification Authorities store.

► **Task 3: Examine the suggested proposals in the Lab Answer Key**

Examine the completed AD RMS plan, and be prepared to discuss your solutions.

Lab answer key solutions:

1. You must deploy one AD RMS cluster in the A. Datum forest and another in Trey Research.
2. You will deploy the Windows Internal Database, in order to reduce the number of servers deployed. However, this will prevent you from deploying additional servers in each cluster.
3. You need only one account to act as the service account. You cannot use the administrator's account, because the configuration of the AD RMS console requires an account other than the installing account to activate the database.
4. You can use self-signed certificates, and place them in each computer's Trusted Certificate Authorities store.

Results: At the end of this lab exercise, you will have planned the deployment of an Active Directory® Rights Management Services (AD RMS) infrastructure for the Adatum.com and TreyResearch.Net forests, based on business requirements and management specifications.

Exercise 2: Deploying an AD RMS Infrastructure for Internal Users

► Task 1: Configure the AD RMS prerequisites

Create and configure accounts

1. On LON-DC1, in Server Manager, click **Tools**, and then click **Active Directory Users and Computers**.
2. In the console tree, expand **Adatum.com**.
3. Right-click **Users**, point to **New**, and then click **User**.
4. In the **New Object – User** dialog box, type **ADRMSSRVC** in the **First name** and **User logon name** boxes, and then click **Next**.
5. In the **New Object – User** dialog box, type **Pa\$\$w0rd** in the **Password** and **Confirm password** boxes. Clear the **User must change password at next logon** check box, click **Next**, and then click **Finish**.
6. Right-click **ADRMSSRVC**, and then click **Add to a group**. Type **Domain Admins**, and then click **OK**.
7. Click **OK** to close the message box.

To add email addresses to the users

1. Click the **Research** organizational unit (OU). Locate and then double-click the account named **Hani Loza**.
2. In the **Properties**, on the **General** tab, in the **E-mail** field, type **hani@adatum.com**, and then click **OK**.
3. Do the same for the following accounts, which are in the following OUs/containers:

OU/Container	User name	Email
Development	Toni Poe	toni@adatum.com
Sales	Limor Henig	limor@adatum.com
Marketing	Stuart Glasson	stuart@adatum.com

To add new groups to Active Directory

4. In the Active Directory Users and Computers console, right-click **Users**, point to **New**, and then click **Group**.
5. In the **New Object – Group** dialog box, in the **Group Name** box, type **Employees**, in the **Group Scope** section, click the **Universal** button, and then click **OK**.

To add email addresses to group objects

1. In the Active Directory Users and Computers console, click the **Research** OU, and then double-click the **Research** security group to open the **Properties** dialog box.
2. On the **General** tab, in the **E-mail** box, type **Research@adatum.com**, and then click **OK**.
3. Perform steps 1 and 2 above for the **Managers** found in the Managers OU and **Employees** security groups located in the **Users** container. Provide the Email format **group@adatum.com**.

To add user accounts to groups

1. In the Active Directory Users and Computers console, click **Users**, and then double-click **Employees**.
2. Click **Members**, and then click **Add**.
3. Type **limor@adatum.com;stuart@adatum.com**, and then click **OK** twice.
4. Perform steps 2 and 3 above to add one member to each of the remaining groups as follows:
 - Limor Henig: Managers
 - Toni Poe: Managers
5. Close the Active Directory Users and Computers console.

Create shared folders on LON-DC1

1. On LON-DC1, from the task bar, open File Explorer, and then right-click **Local Disk (C:)**. Point to **New**, and then click **Folder**.
2. Type **ConfidentialResearch** for the folder name, and then press Enter.
3. Right-click **ConfidentialResearch**, point to **Share with**, and then click **Specific people**. The File Sharing Wizard opens.
4. Under **Choose people on your network to share with**, type **Research**, and then click **Add**. Type **Managers**, and then click **Add**.
5. In the list, click the arrow for **Permission Level** on the group **Research**, and then select **Read/Write**. Repeat this for **Managers**.
6. Click **Share**, and then click **Done**.

Create a share on LON-DC1 to store the AD RMS Templates

1. From File Explorer, right-click **Local Disk (C:)**.
2. Point to **New**, and then click **Folder**.
3. Type **Public** for the folder name, and then press Enter.
4. Right-click **Public**, point to **Share with**, and then click **Specific people**. The File Sharing Wizard opens.
5. Under **Choose people on your network to share with**, click the arrow, select **Everyone**, and then click **Add**.
6. In the list, click the arrow for **Permission Level** on the group **Everyone**, and then confirm that **Read** is selected.
7. Click **Share**, and then click **Done**.
8. Close File Explorer.

► Task 2: Deploy the first server in an AD RMS cluster

Deploy the AD RMS cluster in Adatum.com

1. On LON-DC1, in Server Manager, click **Add roles and features**.
2. Click **Next** three times to get to the **Select server roles** page.
3. On the **Select server roles** page, select **Active Directory Rights Management Services**. When prompted to add features, click **Add Features**, and then click **Next**.
4. On the **Select features** page, click **Next**.

5. In the **Active Directory Rights Management Services** pop-up window, click **Next**.
6. In the **Select role services**, page, verify that **Active Directory Rights Management Server** is selected, and then click **Next**.
7. Click **Install** to add the role. Allow the installation to complete, and then click **Close**.

To configure a new AD RMS root cluster

1. In Server Manager, click the **Notifications** icon.
2. For the task event labeled **Configuration required for Active Directory Rights Management Services at LON-DC1**, click **Perform additional configuration**. The AD RMS Configuration Wizard opens.
3. In the AD RMS Configuration Wizard, on the **Active Directory Rights Management Services** page, click **Next**.
4. On the **Create or Join an AD RMS Cluster** page, accept the default selection (**Create a new AD RMS root cluster**) and then click **Next**.
5. On the **Select Configuration Database Server** page, click **Use Windows Internal Database on this server**, and then click **Next**.
6. In the **Specify Service Account**, page, click **Specify**, in the **Windows Security** dialog box, type **ADRMSSRVC** and the currently set password (**Pa\$\$w0rd**), and then click **OK**.
7. Verify that the **Domain User Account** is set to **ADATUM\ADRMSSRVC**, and then click **Next**.
8. For **Cryptographic Mode**, accept the default (**Cryptographic Mode 2**), and then click **Next**.
9. For **Cluster Key Storage**, accept the default (**Use AD RMS centrally managed key storage**), and then click **Next**.
10. For **Cluster Key Password**, type and confirm a password (**Pa\$\$w0rd**), and then click **Next**.
11. For **Cluster Web Site**, accept the default (**Default Web Site**), and then click **Next**.
12. For **Cluster Address**, select **Use an unencrypted connection (http://)**, for **Fully Qualified Domain Name**, type **LON-DC1.adatum.com** (be sure to use the FQDN, not just LON-DC1), and then click **Next**.
13. For **Licenser Certificate**, accept the default name (**LON-DC1** does not need to be the FQDN), and then click **Next**.
14. For **SCP Registration**, accept the default (**Register the SCP now**), and then click **Next**.
15. In the **Confirmation** page, review your installation selections, and then click **Install**. Click **Close**.
16. Sign out of LON-DC1, and then sign in again as **Adatum\Administrator** with the password **Pa\$\$w0rd** to update the security token of the signed-in user account.
17. Your AD RMS root cluster is now installed and configured.

To open the Active Directory Rights Management Services console

1. On LON-DC1, in Server Manager, click **Tools**, and then click **Active Directory Rights Management Services**. Verify that the console opens with no errors.
2. Close the Active Directory Rights Management Services console.

► **Task 3: Deploy the AD RMS cluster in the Trey Research forest**

To add the AD RMS server role

1. Sign in to TREY-DC1 as **TreyResearch\Administrator** with the password **Pa\$\$w0rd**.
2. On TREY-DC1, in Server Manager, click **Add roles and features**.
3. Click **Next** three times to get to the **Select server roles** page.
4. On the **Select server roles** page, select **Active Directory Rights Management Services**, click **Add Features**, and then click **Next**.
5. On the **Select features** page, click **Next**.
6. On the **Active Directory Rights Management Services** page, click **Next**.
7. On the **Select role services** page, verify that **Active Directory Rights Management Server** is selected, and then click **Next**.
8. Click **Install** to add the role, and then click **Close**.
9. In Server Manager, click **Notifications**, and then click **Perform additional configuration**.
10. In the configuration window, on the **Active Directory Rights Management Services** page, click **Next**.
11. On the **Create or Join an AD RMS Cluster** page, accept the default selection (**Create a new AD RMS root cluster**), and then click **Next**.
12. On the **Select Configuration Database Server** page, click **Use Windows Internal Database on this server**, and then click **Next**.
13. In **Specify Service Account**, click **Specify**, in the **Windows Security** dialog box, type **ADRMSSVC** and the currently set password (**Pa\$\$w0rd**), and then click **OK**.
14. Verify that the **Domain User Account** is set to **TreyResearch\ADRMSSRVC**, and then click **Next**.
15. For **Cryptographic Mode**, accept the default (**Cryptographic Mode 2**), and then click **Next**.
16. For **Cluster Key Storage**, accept the default (**Use AD RMS centrally managed key storage**), and then click **Next**.
17. For **Cluster Key Password**, type and confirm a password (**Pa\$\$w0rd**), and then click **Next**.
18. For **Cluster Web Site**, accept the default (**Default Web Site**), and then click **Next**.
19. For **Cluster Address**, select **Use an unencrypted connection (http://)**, type **TREY-DC1.TreyResearch.net**, and then click **Next**.
20. For **Licenser Certificate**, accept the default name (**TREY-DC1** - it does not need to be the FQDN), and then click **Next**.
21. For **SCP Registration**, accept the default (**Register the SCP now**), and then click **Next**.
22. In the **Confirmation** page, review your installation selections, and then click **Install**. Click **Close**.
23. Sign out of TREY-DC1, and then sign in again as **Treyresearch\Administrator** with the password **Pa\$\$w0rd** to update the security token of the logged-in user account. Your AD RMS root cluster is now installed and configured.

To open the Active Directory Rights Management Services console

1. On TREY-DC1, in Server Manager, click **Tools**, and then click **Active Directory Rights Management Services**.
2. Verify that the Active Directory Rights Management Services console opens without errors.
3. Close the Active Directory Rights Management Services console.

► Task 4: Configure the AD RMS Templates**To create a new AD RMS rights policy template**

1. Switch to LON-DC1. In Server Manager, click **Tools**, and then click **Active Directory Rights Management Services**.
2. In the Active Directory Rights Management Services console, expand **lon-dc1.adatum.com**.
3. Select and then right-click **Rights Policy Templates**, and then click **Properties**.
4. Select the **Enable export** check box, in the **Specify templates file location (UNC)** box, type **\\LON-DC1\public**, and then click **OK**.
5. In the Actions pane, click **Create Distributed Rights Policy Template** to start the Create Distributed Rights Policy Template Wizard.
6. Click **Add**. In the **Language** box, choose **English (United States)**.
7. In the **Name** box, type **Adatum.com RC**.
8. In the **Description** box, type **Adatum.com Research Confidential**, click **Add**, and then click **Next**.
9. Click **Add**, in the **The e-mail address of a user or group** box, type **Managers@adatum.com**, and then click **OK**.
10. Select the **View** check box to grant the **Managers@adatum.com** group **Read** access to any document created by using this AD RMS rights policy template.
11. Click **Add**, in the **The e-mail address of a user or group** box, type **research@adatum.com**, and then click **OK**.
12. Select the **Full Control** check box to grant the **research@adatum.com** group **full control** access to any document created by using this AD RMS rights policy template.
13. Click **Finish**, and then close the Active Directory Rights Management Services console.

To enable the automated scheduled task

1. Sign in to LON-CL1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. On the Start screen, type **Schedule tasks**, and then click on **Schedule tasks** item.
3. Expand **Task Scheduler Library**, expand **Microsoft**, expand **Windows**, and then click **Active Directory Rights Management Services Client**.
4. In the top details pane, right-click **AD RMS Rights Policy Template Management (Automated)**, and then click **Enable**.
5. Close Task Scheduler.
6. Right-click the **Windows Start** icon, click **Run**, in the **Run** text box, type **regedit.exe**, and then press Enter.
7. Expand the following registry key:

HKEY_CURRENT_USER\Software\Microsoft\Office\15.0\Common

8. Right-click **Common** and select **New**, and then click **Key**. Type **DRM**, and then press Enter.
9. Right-click **DRM**, click **New**, and then click **Expandable String Value**.
10. In the **Value name** box, type **AdminTemplatePath**, and then press Enter.
11. Double-click the **AdminTemplatePath** registry value, in the **Value data** box, type **%LocalAppData%\Microsoft\DRM\Templates**, and then click **OK**.
12. Close Registry Editor.
13. In the taskbar, click **File Explorer**.
14. In the **This PC** console tree, right-click **This PC**, and then select **Properties**.
15. In the **System** console tree, select **Remote settings**.
16. In the **System Properties** window, click **Select Users**.
17. In the **Remote Desktop Users** window, click **Add**.
18. In the **Select Users or Groups** window, in the **Enter the object names to select** text box, type **Domain Users**, and then click **OK** three times.
19. Sign out of LON-CL1.

► Task 5: Configure AD RMS Exclusion Policies

A Datum has a contract with a management consulting group. A Datum has given Toni Poe, a management consultant from that group, office space and an A. Datum computer system to perform office work. She has an Active Directory Domain Services (AD DS) account, and A. Datum has placed her in the Management universal group. However, she should not have Read permissions on restricted Research Confidential files to which the Management group has view access. To that end, we will create an exclusion policy on the AD RMS cluster for her.

Create a RAC for Toni Poe

1. On LON-CL1, sign in as **Adatum\Toni** with the password **Pa\$\$w0rd**.
2. From the Start screen, click the **Desktop** tile.
3. From the desktop taskbar, click **Internet Explorer**.
4. In the URL bar, type **http://lon-dc1.adatum.com**.
5. Close any pop-up warnings.
6. Click the gear icon in the upper-right of the **Internet Explorer** window, and then click **Internet options**.
7. Click the **Security** tab, click **Local intranet**, and then click **Sites**.
8. Click **Advanced**, in the **Add this website to the zone** box, type **http://LON-DC1.adatum.com**, and then click **Add**.
9. Click **Close**, and then click **OK** twice.
10. Close Internet Explorer®.
11. Click to the Start screen.
12. On the Start screen, type **Word**, in the **Search** area, select **Word 2013**, and then click **Blank document**.
13. Click the **File** tab in the ribbon, on the **Info** page, click **Protect Document**, and then click **Restrict Access, Connect to Digital Rights Management Servers and get templates**. Click **Protect Document**, click **Restrict Access**, and then click **Restricted Access**.

14. In the **Permission** dialog box, select the **Restrict permission to this document** check box, and then click **OK**.
15. Save the document as **Text.docx** to the local Documents library, close all open windows, and then sign out of LON-CL1.

Exclude Rights Account Certificates

1. On LON-DC1, in Server Manager, click **Tools**, and then click **Active Directory Rights Management Services**.
2. In the Active Directory Rights Management Services console, expand **lon-dc1.adatum.com**.
3. Expand **Exclusion policies**, and then click **Users**.
4. In the Actions pane, click **Enable User Exclusion**.
5. In the Actions pane, click **Exclude RAC**.
6. On the **Add RAC to be excluded** page, ensure that **Use this option for excluding rights accounts certificates of internal users who have an Active Directory Domain Services account** is selected. In the **User Name** box, type **toni@adatum.com**. Click **Finish**. Toni's email address and public key should now be in the table on the **User Exclusion Information** page. If the public key is missing, this is because Toni has never consumed information rights. In order to complete this action, Toni must sign in to a client computer, which you just did, and then use a protected document. Close the Active Directory Rights Management Services console.

► Task 6: Validate the internal deployment

To verify the functionality of the AD RMS deployment, you will sign in as Hani Loza, and then restrict permissions on a Microsoft Word document. The permissions will enable the members of the Management group to read the document but not make changes to it or print or copy it, while members of the Research group will have full control. Then you will sign in as Limor Henig, verifying that only the appropriate permission to read the document has been granted.

Before you can consume rights-protected content, you must add the AD RMS cluster URL to the Local Intranet security zone.

Add the AD RMS cluster URL to the Local Intranet security zone for all users who will be consuming rights-protected content.

To add AD RMS cluster to Local Intranet security zone

1. Sign in to LON-CL1 as Hani (**Adatum\Hani**) with the password **Pa\$\$w0rd**.
2. Click the **Desktop** tile.
3. From the taskbar, click **Internet Explorer**.
4. Click **Tools** (the gear icon in upper right), and then click **Internet options**.
5. Click the **Security** tab, click **Local intranet**, and then click **Sites**.
6. Click **Advanced**.
7. In the **Add this website to the zone** box, type **http://LON-DC1.adatum.com**, and then click **Add**. Click **Close**.
8. Click **OK** twice and then close Internet Explorer.
9. Sign out of LON-CL1.
10. Repeat steps 1 to 9 for **Adatum\Limor**.

To restrict permissions on a Microsoft Word document

1. Sign in to LON-CL1 as **Hani** with a password of **Pa\$\$w0rd**.
2. On the Start screen, type **Word**, and then press Enter. In the **Search** area, select **Word 2013**. In the **First things first** window, click **Use recommended settings**, and click **Accept**. On the **User Account Control** pop-up, click **No**, and on the **Microsoft Word** pop-up, click **OK**. In the **Office** dialog box, click **Next** three times, and then click **All done**.
3. Select the Blank document template in the **Recent** page, and on the blank document page, type **Managers can read this document, but they cannot change, print, or copy it. Research group members have Full control**.
4. From the **File** tab, on the **Info** page, click **Protect Document**, click **Restrict Access**, and then click **Connect to Digital Rights Management Servers** and get templates.
5. On the **Info** page, click **Protect Document**, click **Restrict Access**, and then click **Restricted Access**.
6. In **Permissions**, select the **Restrict permission to this document** check box, and then in the **Read** box, type **Managers@adatum.com**. In the **Change** box, type **research@adatum.com**.
7. Click **OK** to close the **Permission** dialog box.
8. From the File menu, click **Save As**, click **Browse**, in the **Save As File name** text box, type **\\LON-DC1\ConfidentialResearch\ADRMS-TST.docx**, and then click **Save**.
9. Close Microsoft Word, and then sign out of LON-CL1.

To view a rights-protected document

1. Sign in to LON-CL1 as Limor Henig (**ADATUM\limor**) with the password **Pa\$\$w0rd**.
2. Click the **Desktop** tile.
3. Open File Explorer, and then browse to **\\LON-DC1\ConfidentialResearch**. Double-click **ADRMS-TST.docx** to open it in Microsoft Word 2013.
4. Repeat the same steps you performed earlier for Hani when opening Word for the first time.
5. When the document opens, note that the Restricted Access yellow bar shows: **Permission is currently restricted. Only specified users can access this content**.
6. Click the **File** tab. Notice that the **Print** option is not available.
7. Close Microsoft Word and sign out of LON-CL1.

Results: You should have a working AD RMS cluster in both the Adatum.com and TreyResearch.net forests. In addition, you should be able to protect Microsoft Office documents with IRM, and see the results of that protection.

Exercise 3: Implementing AD RMS Integration with Dynamic Access Control

► Task 1: Enable resource properties

To enable resource properties

1. On LON-DC1, and from Server Manager, click **Tools**, click **Active Directory Administrative Center**, and then in the console tree, switch to **Tree View**.
2. Expand **Dynamic Access Control**, and then select **Resource Properties**.
3. In the **Display Name** column, scroll down to the **Impact** property. Right-click **Impact**, and then click **Enable**.
4. In the **Display Name** column, scroll down to the **Personally Identifiable Information** property. Right-click **Personally Identifiable Information**, and then click **Enable**.
5. To publish the resource properties in the **Global Resource List**, in the left pane, click **Dynamic Access Control**, in the details pane, double-click **Resource Property Lists**, and then double-click **Global Resource Property List** (expand this window).
6. Under **Resource Properties**, click **Add**, scroll down to and click **Impact**, and then add **Impact** to the list by clicking the >> button. Do the same for **Personally Identifiable Information**. Click **OK** twice to finish. Note that these resource properties may be in the list already. If so, the OK button may be grayed out. In that case, simply verify they are in the list, and then click **Cancel**.

► Task 2: Create classification rules

This task explains how to create the **High Impact** classification rule. This rule will search the content of documents, and if it finds the string "Adatum Confidential", it will classify this document as having high business impact. This classification will override any previously assigned classification of low-business impact.

You will also create a **High PII** rule. This rule searches the content of documents, and if it finds a Social Security number, it classifies the document as having high personally identifiable information (PII).

Create the high-impact classification rule

1. Sign in to LON-SVR1 by using **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. On the taskbar, click **File Explorer**.
3. In File Explorer, in the console tree, expand **This PC**, and then click **Local Disk (C:)**.
4. Right-click **Research Documents**, and then select **Share with** and then click **Specific People**.
5. In the blank text box, click the drop-down list box arrow, select **Everyone**, and then click **Add**.
6. In the Everyone entry, click the drop-down list box arrow by **Read** and change to **Read/Write**.
7. Click the **Share** button, and then click **Done**. Close File Explorer.
8. Maximize Server Manager, if it is not maximized already. In Server Manager, click **Add roles and features**.
9. Click **Next** three times until you reach the **Select server roles** page. Expand **File And Storage Services (Installed)** and expand **File And iSCSI Services (Installed)**. Select the check box next to **File Server Resource Manager**. Click **Add Features**, click **Next** two times, and then click **Install**. When the installation is finished, click **Close**.

10. You need to refresh the Global Resource Properties from Active Directory. Open Windows PowerShell®, type the following cmdlet at the command prompt, and then press Enter:

```
Update-FSRMClassificationPropertyDefinition
```

11. Close Windows PowerShell.
12. In Server Manager, click **Tools**, and then click **File Server Resource Manager**.
13. In the left pane of File Server Resource Manager, expand **Classification Management**, and then select **Classification Rules**. In the **Actions** pane, click **Configure Classification Schedule**. On the **Automatic Classification** tab, select **Enable fixed schedule**, select **Sunday**, and then select the **Allow continuous classification for new files** check box. Click **OK**.
14. In the Classification Rules node, in the Actions pane, click **Create Classification Rule**. This opens the **Create Classification Rule** dialog box.
15. In the **Rule name** box, type **High Business Impact**.
16. In the **Description** box, type **Determines if the document has a high business impact based on the presence of the string "Adatum Confidential"**.
17. On the **Scope** tab, click **Set Folder Management Properties**, select **Folder Usage**, click **Add**, click **Browse**, browse to **C:\Research Documents**, and then click **OK**.
18. Under **Value**, select the **Group Files** check box, click **OK**, and then click **Close**.
19. On the **Scope** tab, select **Group Files**.
20. Click the **Classification** tab. Under **Choose a method to assign the property to files**, select **Content Classifier** from the drop-down list box.
21. Under **Choose a property to assign to files**, select **Impact** from the drop-down list box.
22. Under **Specify a value**, select **High** from the drop-down list box.
23. Under **Parameters**, click **Configure**. In the **Classification Parameters** dialog box, in the **Expression Type** list, select **String**. In the **Expression** box, type **Adatum Confidential**, and then click **OK**.
24. Click the **Evaluation Type** tab. Click **Re-evaluate existing property values**, click **Overwrite the existing value**, and then click **OK** to finish.

To create the high-PII classification rule

25. In the left pane of File Server Resource Manager, expand **Classification Management**, and then click **Classification Rules**.
26. In the Actions pane, click **Create Classification Rule**.
27. In the **Rule name** box, type **High PII**. In the **Description** box, type **Determines if the document has a high PII based on the presence of a Social Security Number**.
28. Click the **Scope** tab, and then select the **Group Files** check box. The C:\Research Documents folder should show as included in the scope.
29. Click the **Classification** tab. Under **Choose a method to assign the property to files**, select **Content Classifier** from the drop-down list box.
30. Under **Choose a property to assign to files**, select **Personally Identifiable Information** from the drop-down list box.
31. Under **Specify a value**, select **High** from the drop-down list box.

32. Under **Parameters**, click **Configure**. In the Classification Parameters window, in the **Expression Type** list, select **Regular Expression**. In the **Expression** box, type the following expression without including any line breaks, and then click **OK**:

```
^(?!000)([0-7]\d{2}|7([0-7]\d|7[012]))([ -]?)?(?!00)\d\d3(?!0000)\d{4}$
```



Note: This expression will allow invalid social security numbers. This allows us to use fictitious social security numbers in the lab.

33. Click the **Evaluation Type** tab. Select **Re-evaluate existing property values, overwrite the existing value**, and then click **OK** to finish.
34. You should now have two classification rules:
- High Business Impact
 - High PII

► Task 3: Automatically protect documents with AD RMS

Now that you have created rules to classify documents automatically based on content, you must create a file management task that uses AD RMS to protect certain documents automatically based on their classification. In this step, you will create a file management task that protects any documents with a high PII automatically. Only members of the Research group will have access to documents that contain high PII.

To add LON-DC1 to the Local Intranet security zone on LON-SVR1

1. On LON-SVR1, click to the Start screen, and then click **Internet Explorer**.
2. Click **Tools** (the gear icon in the upper right), and then click **Internet options**.
3. Click the **Security** tab, click **Local intranet**, and then click **Sites**.
4. Click **Advanced**.
5. In the **Add this website to the zone** box, type **http://LON-DC1.adatum.com**, and then click **Add**. Click **Close**.
6. Click **OK** twice, and then close Internet Explorer.

To protect documents with AD RMS

1. In File Server Resource Manager, in the left pane, select **File Management Tasks**. In the Actions pane, select **Create File Management Task**.
2. In the **Task name** field, type **High PII**. In the **Description** field, type **Automatic RMS protection for high PII documents**.
3. Click the **Scope** tab, and then select the **Group Files** check box. The **C:\Research Documents** location should show as being part of the scope.
4. Click the **Action** tab. Under **Type**, select **RMS Encryption**. Select the **Adatum.com RC** template.
5. Click the **Condition** tab, and then click **Add**. Under **Property**, select **Personally Identifiable Information**. Under **Operator**, select **Equal**. Under **Value**, select **High**. Click **OK**.
6. Click the **Schedule** tab. In the **Schedule** section, click **Weekly**, and then select **Sunday**. Running the task once a week will ensure that you catch any documents that may have been missed due to a service outage or other disruptive event.
7. In the **Continuous operation** section, select **Run continuously on new files**, and then click **OK**. You should now have a file management task named High PII.

► **Task 4: Verify the deployment**

1. On LON-SVR1, open File Explorer, and then navigate to **C:\Research Documents**.
2. Right-click the **Finance Memo** document, click **Properties**, click the **Classification** tab, and then notice that both properties currently have no value. Click **Cancel**.
3. Right-click the **Request for Approval to Hire** document, and then select **Properties**.
4. Click the **Classification** tab, and notice that the both properties currently have no value. Click **Cancel**.
5. Switch to LON-CL1, and sign in as **Adatum\Hani** with the password **Pa\$\$w0rd**.
6. From the desktop, click the File Explorer icon in the taskbar, and in the URL text area, type **\\LON-SVR1\Research Documents**, and then press Enter.
7. Open the **Finance Memo** document. Type **Adatum Confidential**, and then press **Enter** twice. Save the document, and then close Microsoft Word.
8. Open the **Request for Approval to Hire** document. Type **Social Security #:**, press Enter, on a new line, type **777-77-7777**, and then press **Enter** twice. This must be on a new line for Dynamic Access Control to notice the expression quickly. Save the document, and then close Microsoft Word.
9. Switch to LON-SVR1. In File Explorer, navigate to **C:\Research Documents**.
10. Right-click **Finance Memo**, and click **Properties**. Click the **Classification** tab. Notice that the **Impact** property is now set to **High**. Click **Cancel**.
11. Right-click the **Request for Approval to Hire** document, and then click **Properties**.
12. Click the **Classification** tab. Notice that the **Personally Identifiable Information** property is now set to **High**. Click **Cancel**.
13. If either of the properties in step 11 or step 12 is not set properly, open the File Server Resource Manager, and then select **File Management Tasks**. Select **High PII** in the details pane, and in the Actions pane, click **Run File Management Task Now**. In the message box, select **Run task in Background**, and then click **OK**. The Run File Management Task item will be grayed out. After a few moments, it will be plain text again, indicating the management task is finished. Repeat steps 9 through 12 above.
14. On all machines except LON-DC1, close all open windows, and then sign out.

Results: You should have applied Dynamic Access Control classification rules to IRM-protected content.

Exercise 4: Implementing AD RMS Integration for External Users

► Task 1: Export the trusted user domain policy

Export a trusted user domain

1. On LON-DC1, from Server Manager, click **Tools**, and then click **Active Directory Rights Management Services**.
2. In the console tree, expand **lon-dc1.adatum.com**, expand **Trust Policies**, and then click **Trusted User Domains**.
3. In the Actions pane, click **Export Trusted User Domain**. The **Export Trusted User Domain As** dialog box opens.
4. In the **File name** box, type **C:\ADRMS_LON-DC1_LicensorCert.bin**.
5. Click **Save** to save the file with the name and location that you specified.
6. Repeat steps 1 to 5 on TREY-DC1, but use the name **C:\ADRMS_TREY-DC1_LicensorCert.bin** for the .bin file.

► Task 2: Export the trusted publishing domain policy

Export a trusted publishing domain policy

1. On LON-DC1, from Server Manager, click **Tools**, and then click **Active Directory Rights Management Services**.
2. In the console tree, expand **lon-dc1.adatum.com**, expand **Trust Policies**, and then click **Trusted Publishing Domains**.
3. In the results pane, select **LON-DC1**, and then in the Actions pane, click **Export Trusted Publishing Domain**.
4. In the **Export Trusted Publishing domain** dialog box, click **Save As**, and then type **C:\AdatumTrustedPubDomain.xml**. Click **Save**.
5. In the **Password** and **Confirm password** boxes, type **Pa\$\$w0rd**.
6. Click **Finish** to create the trusted publishing domain file.
7. Repeat steps 1 to 6 on TREY-DC1 by using the file name **C:\TreyTrustedPubDomain.xml** with the password **Pa\$\$word**.

► Task 3: Import the trusted user domain policy from the partner domain

1. On LON-DC1, in Server Manager, click **Tools**, and in the drop-down list box, click **DNS**. Expand **LON-DC1**, select and right-click **Conditional Forwarders**, and then click **New Conditional Forwarder**.
2. Under **DNS Domain**, type **TreyResearch.net**. In the **IP addresses of the master servers** box, type **172.16.10.10**. Press Enter, and then click **OK**. Close DNS Manager.
3. On TREY-DC1, repeat steps 1 and 2. Use the DNS domain **Adatum.com** and set the **IP Address of the master servers** as **172.16.0.10**.
4. On LON-DC1, open the Active Directory Rights Management Services console, and then expand **lon-dc1.adatum.com**.
5. In the console tree, expand **Trust Policies**, and then click **Trusted User Domains**.
6. In the Actions pane, click **Import Trusted User Domain**.
7. In the **Trusted user domain file** box, type **\\TREY-DC1\C\$\ADRMS_TREY-DC1_LicensorCert.bin**.

8. In the **Display name** box, type **TreyResearch**. Click **Finish**.
9. Repeat steps 4 to 8 on TREY-DC1, replacing the file name above with **\\LON-DC1\C\$\ADRMS_LON-DC1_LicensorCert.bin** and **Display name** with **Adatum**.

► **Task 4: Import the trusted publishing domains policy from the partner domain**

Add a trusted publishing domain

1. On LON-DC1, open the Active Directory Rights Management Services console and expand **lon-dc1.adatum.com**.
2. In the console tree, expand **Trust Policies**, and then click **Trusted Publishing Domains**.
3. In the Actions pane, click **Import Trusted Publishing Domain**.
4. In the **Trusted Publishing Domain file** box, type **\\TREY-DC1\c\$\TreyTrustedPubDomain.xml**.
5. Type **Pa\$\$w0rd** in the **Password** box.
6. In the **Display name** box, type **TreyResearch Domain**. Click **Finish**.
7. Repeat steps 1 to 6 on TREY-DC1 by using the file name **\\LON-DC1\C\$\AdatumTrustedPubDomain.xml** and a **Display name** of **Adatum Domain**.

► **Task 5: Configure anonymous access to the AD RMS licensing server**

1. Switch to LON-DC1.
2. In Server Manager, click **Tools**, and then click **Internet Information Services (IIS) Manager**. If the Internet Information Services (IIS) Manager message box appears, click **Cancel**.
3. In Internet Information Services (IIS) Manager, expand **LON-DC1\Sites\Default Web Site\wmcs**.
4. Right-click **licensing**, and then click **Switch to Content View**.
5. Right-click **license.aspx**, and then click **Switch to Features View**.
6. Double-click **Authentication**, click **Anonymous Authentication**, and then in the Actions pane, click **Enable**.
7. Click **Windows Authentication**, and then click **Disable**.
8. Right-click **licensing**, and then click **Switch to Content View**.
9. Right-click **ServiceLocator.aspx**, and then click **Switch to Features View**.
10. Double-click **Authentication**, click **Anonymous Authentication**, and then in the Actions pane, click **Enable**.
11. Click **Windows Authentication**, and then click **Disable**.
12. Close Internet Information Services (IIS) Manager.

► **Task 6: Verify user access to the protected document**

1. Sign in to LON-CL1 as user **Adatum\Hani** with the password **Pa\$\$w0rd**.
2. From the Start screen, type **\\LON-DC1\ConfidentialResearch**, and then press Enter.
3. Double-click **ADRMS-TST.docx**. On the **Restricted Access** yellow bar, click **Change Permission**.
4. In the box next to the **Read** permission, place a semicolon after **Managers@adatum.com**, and then add **liberty@treysresearch.net**. Click **OK**.
5. Save the **ADRMS-TST.docx** file, close Microsoft Word, and then sign out of LON-CL1.

6. Switch to LON-DC1, and then open File Explorer.
7. Copy the **ADRMS-TST.docx** file from **C:\ConfidentialResearch** to **\\TREY-DC1\Public**.
8. Sign in to **TREY-CL1** as user **TreyResearch\Liberty** with the password **Pa\$\$w0rd**.
9. Click the **Desktop** tile.
10. From the taskbar, click **Internet Explorer**.
11. Click **Tools** (the gear icon in upper right), and then click **Internet options**.
12. Click the **Security** tab, click **Local intranet**, and then click **Sites**.
13. Click **Advanced**.
14. In the Add this website to the zone box, type **http://TREY-DC1.TreyResearch.net**, and then click **Add**. Click **Close**.
15. Click **OK** twice and then close Internet Explorer.
16. From the Start screen, type **\\TREY-DC1\Public**, and then press **Enter**.
17. Double-click the **ADRMS-TST.docx** file.
18. An Active Directory Rights Management Services pop-up window will appear that says, **"To create and consume content with restricted access..."** Click **OK**.
19. Repeat the **First things first** and **Office first time** steps you used previously to open the document. In the document, in the yellow **Restricted Access** bar, click **View Permission**.



Note: The permission for Liberty@treyresearch.net is what Hani Loza assigned previously.

20. Click **OK**.
21. Close all windows, and sign out of all virtual machines.

► Task 7: 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 **20414C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20414C-LON-SVR1**, **20414C-LON-CL1**, **20414C-TREY-DC1**, and **20414C-TREY-CL1**.

Results: You should have both a working trusted user and trusted publishing domain policy between the Adatum.com and TreyResearch.net forests. In addition, you should be able to protect Microsoft Office documents with IRM for external users across the domains.

MCT USE ONLY. STUDENT USE PROHIBITED

Notes

MCT USE ONLY. STUDENT USE PROHIBITED

Notes

MCT USE ONLY. STUDENT USE PROHIBITED

Notes

MCT USE ONLY. STUDENT USE PROHIBITED

Notes