

OFFICIAL MICROSOFT LEARNING PRODUCT

20743A

**Upgrading Your Skills to MCSA:
Windows Server 2016**

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.

© 2016 Microsoft Corporation. All rights reserved.

Microsoft and the trademarks listed at <http://www.microsoft.com/trademarks> are trademarks of the Microsoft group of companies. All other trademarks are property of their respective owners.

Product Number: 20743A

Part Number: X21-18464

Released: 09/2016

MICROSOFT LICENSE TERMS MICROSOFT INSTRUCTOR-LED COURSEWARE

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

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

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

1. DEFINITIONS.

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

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

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

2.1 Below are five separate sets of use rights. Only one set of rights apply to you.

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

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

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

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

b. If you are a Microsoft Learning Competency Member:

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

c. If you are a MPN Member:

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

d. If you are an End User:

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

e. If you are a Trainer.

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

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

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

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

2.4 Third Party Notices. The Licensed Content may include third party code that Microsoft, not the third party, licenses to you under this agreement. Notices, if any, for the third party code are included for your information only.

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

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

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

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

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

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

This limitation applies to

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

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

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

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

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

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

Cette limitation concerne:

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

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

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

Revised July 2013

Welcome!

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

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

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

Sincerely,

Microsoft Learning
www.microsoft.com/learning

Microsoft | Learning

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

Acknowledgments

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

Andrew J. Warren – Content Developer

Andrew Warren has more than 25 years of experience in the information technology (IT) industry, many of which he has spent teaching and writing. He has been involved as a Subject Matter Expert for many of the Windows Server 2012 courses, and the technical lead on many Windows 8 courses. He also has been involved in developing TechNet sessions on Microsoft Exchange Server. Based in the United Kingdom, Andrew runs his own IT training and education consultancy.

Claus Jacob Wordenskjold – Content Developer

Claus Jacob Wordenskjold is an independent consultant and trainer based in Denmark. He founded his company, Chinchilla Data, in 1995, and he has more than 25 years of experience in IT. Claus has been a Microsoft Certified Trainer (MCT) since 2002, and he has delivered training throughout Europe. He specializes in Windows Client and Windows Server courses, but he conducts occasional training in Microsoft SharePoint. Claus holds certifications in every Windows operating system since Windows 2000, and he provides consulting services on Windows Server, Active Directory Domain Services (AD DS), and Group Policy. Claus has been a speaker at local Danish Microsoft events, and he has authored several Windows-related courses.

Dave Franklyn – Content Developer

David M. Franklyn, MCSE, MCITP, Microsoft MVP Windows and Devices for IT, is also an Eastern USA Regional Lead MCT. Dave has been a Microsoft MVP since 2011, and a Senior Information Technology Trainer and Consultant at Auburn University in Montgomery Alabama, since 1998. He is the owner of DaveMCT, Inc. LLC and is a training partner with Dunn Training. Working with computers since 1976, Dave started out in the mainframe world and moved early into the networking arena. Before joining Auburn University, Dave spent 22 years in the United States Air Force as an electronic communications and computer systems specialist, retiring in 1998. Dave is president of the Montgomery Windows IT Professional Group, and a guest speaker at many events involving Microsoft products.

Joshua Shackelford – Content Developer

Joshua Shackelford has more than 12 years of experience in the IT industry as an engineer, architect, consultant, and administrator. Joshua has extensive experience consulting on the System Center suite, Hyper-V, and Active Directory. His clients include large energy corporations, financial institutions, and retail organizations. Joshua has been involved as a SME on Windows Server, Hyper-V, and Failover Clustering.

Slavko Kukrika – Technical Reviewer

Slavko Kukrika has been an MCT for more than 18 years. He holds many technical certifications, and he is honored to be a Microsoft Most Valuable Professional (MVP). Slavko specializes in the Windows operating system, Active Directory, and virtualization. He is a Windows Insider, and he has worked with Windows 10 since it became available publicly. Slavko regularly presents at technical conferences, and he is the author of several Microsoft Official Courses. In his private life, Slavko is the proud father of two sons, and he tries to extend each day to at least 25 hours or more!

Contents

Module 1: Installing and configuring Windows Server 2016

Module Overview	1-1
Lesson 1: Introducing Windows Server 2016	1-2
Lesson 2: Installing Windows Server 2016	1-7
Lesson 3: Configuring Windows Server 2016	1-15
Lab: Installing and configuring Nano Server	1-21
Lesson 4: Preparing for upgrades and migrations	1-26
Lesson 5: Migrating server roles and workloads	1-34
Lesson 6: Windows Server activation models	1-37
Module Review and Takeaways	1-40

Module 2: Overview of storage in Windows Server 2016

Module Overview	2-1
Lesson 1: Overview of storage in Windows Server 2016	2-2
Lesson 2: Implementing Data Deduplication	2-11
Lesson 3: Configuring iSCSI storage	2-30
Lab A: Implementing and managing storage	2-38
Lesson 4: Configuring the Storage Spaces feature in Windows Server 2016	2-45
Lab B: Implementing and managing advanced storage solutions	2-55
Module Review and Takeaways	2-63

Module 3: Implementing Directory Services

Module Overview	3-1
Lesson 1: Deploying Active Directory domain controllers	3-2
Lesson 2: Implementing service accounts	3-14
Lab: Implementing and managing AD DS	3-18
Lesson 3: Azure AD	3-22
Module Review and Takeaways	3-31

Module 4: Implementing AD FS

Module Overview	4-1
Lesson 1: Overview of AD FS	4-2
Lesson 2: Deploying AD FS	4-12
Lesson 3: Implementing AD FS for a single organization	4-18
Lab A: Implementing AD FS	4-24

Lesson 4: Implementing Web Application Proxy	4-29
Lab B: Implementing Web Application Proxy	4-35
Lesson 5: Implementing SSO with Microsoft Online Services	4-39
Module Review and Takeaways	4-42
Module 5: Implementing network services	
Module Overview	5-1
Lesson 1: Overview of networking enhancements	5-2
Lesson 2: Implementing IPAM	5-20
Lesson 3: Managing IP address spaces with IPAM	5-30
Lab: Implementing network services	5-37
Module Review and Takeaways	5-45
Module 6: Implementing Hyper-V	
Module Overview	6-1
Lesson 1: Configure the Hyper-V role in Windows Server 2016	6-2
Lesson 2: Configuring Hyper-V storage	6-11
Lesson 3: Configuring Hyper-V networking	6-21
Lesson 4: Configuring Hyper-V virtual machines	6-28
Lab: Implementing server virtualization with Hyper-V	6-38
Module Review and Takeaways	6-45
Module 7: Configuring advanced networking features	
Module Overview	7-1
Lesson 1: Overview of high-performance networking features	7-2
Lesson 2: Configuring advanced Hyper-V networking features	7-12
Lab: Configuring advanced Hyper-V networking features	7-23
Module Review and Takeaways	7-27
Module 8: Implementing Software Defined Networking	
Module Overview	8-1
Lesson 1: Overview of Software Defined Networking	8-2
Lesson 2: Implementing network virtualization	8-11
Lesson 3: Implementing Network Controller	8-16
Lab: Deploying Network Controller	8-29
Module Review and Takeaways	8-34

Module 9: Implementing remote access

Module Overview	9-1
Lesson 1: Remote access overview	9-2
Lesson 2: Implementing DirectAccess	9-9
Lesson 3: Implementing VPN	9-24
Lab: Implementing DirectAccess	9-36
Module Review and Takeaways	9-41

Module 10: Deploying and managing Windows and Hyper-V containers

Module Overview	10-1
Lesson 1: Overview of containers in Windows Server 2016	10-2
Lesson 2: Deploying Windows Server and Hyper-V containers	10-8
Lesson 3: Installing, configuring, and managing containers by using Docker	10-16
Module Review and Takeaways	10-33

Module 11: Implementing failover clustering

Module Overview	11-1
Lesson 1: Overview of failover clustering	11-2
Lesson 2: Implementing a failover cluster	11-18
Lesson 3: Configuring highly available applications and services on a failover cluster	11-24
Lesson 4: Maintaining a failover cluster	11-30
Lesson 5: Implementing a stretch cluster	11-36
Lab: Implementing failover clustering	11-43
Module Review and Takeaways	11-49

Module 12: Implementing failover clustering with Windows Server 2016**Hyper-V**

Module Overview	12-1
Lesson 1: Overview of the integration of Hyper-V Server 2016 with failover clustering	12-2
Lesson 2: Implementing Hyper-V virtual machines on failover clusters	12-7
Lesson 3: Implementing Windows Server 2016 Hyper-V virtual machine migration	12-20
Lesson 4: Implementing Hyper-V Replica	12-24
Lab: Implementing failover clustering with Windows Server 2016 Hyper-V	12-29
Module Review and Takeaways	12-36

Lab Answer Keys

Module 1 Lab: Installing and configuring Nano Server	L1-1
Module 2 Lab A: Implementing and managing storage	L2-7
Module 2 Lab B: Implementing and managing advanced storage solutions	L2-16
Module 3 Lab: Implementing and managing AD DS	L3-23
Module 4 Lab A: Implementing AD FS	L4-27
Module 4 Lab B: Implementing Web Application Proxy	L4-32
Module 5 Lab: Implementing network services	L5-37
Module 6 Lab: Implementing server virtualization with Hyper-V	L6-47
Module 7 Lab: Configuring advanced Hyper-V networking features	L7-55
Module 8 Lab: Deploying Network Controller	L8-59
Module 9 Lab: Implementing DirectAccess	L9-63
Module 11 Lab: Implementing failover clustering	L11-67
Module 12 Lab: Implementing failover clustering with Windows Server 2016 Hyper-V	L12-77

About This Course

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

Course Description



Note: This first release ('A') MOC version of course 20743A has been developed on Windows Server 2016 Technical Preview 5. Microsoft Learning will release a 'B' version of this course after the RTM version of the software is available.

This five-day, instructor-led course explains how to implement and configure new Windows Server 2016 features and functionality. This course is for information technology (IT) professionals who want to upgrade their technical skills from Windows Server 2008 or Windows Server 2012 to Windows Server 2016. This course presumes a high level of knowledge about previous Windows Server technologies and skills equivalent to the Microsoft Certified Solutions Associate (MCSA): Windows Server 2008 or Windows Server 2012 credential.

This course is not a product-upgrade course, detailing considerations for migrating and upgrading students' specific environment to Windows Server 2016. Rather, this course provides updates to students' existing Windows Server knowledge and skills, as they pertain to Windows Server 2016.

Audience

This course is for IT professionals who are experienced Windows Server 2012 or Windows Server 2008 system administrators, with real-world experience working in a Windows Server 2008 R2 or Windows Server 2008 enterprise environment. Additionally, students should have obtained the MCSA credential for Windows Server 2008 or Windows Server 2012, or they should have equivalent knowledge.

IT professionals who plan to take the Microsoft Certified Solutions Expert (MCSE) exams might be interested in this course, as preparation for the MCSA exams, which are a prerequisite for the MCSE specialties.

Student Prerequisites

This course requires that you meet the following prerequisites:

- Two or more years of experience with deploying and managing Windows Server 2012 or Windows Server 2008 environments.
- Experience with day-to-day Windows Server 2012 or Windows Server 2008 system-administration management and maintenance tasks.
- Experience with Windows networking technologies and implementation.
- Experience with Active Directory technologies and implementation.
- Experience with Windows Server virtualization technologies and implementation.
- Knowledge equivalent to the MCSA credentials of Windows Server 2008 or Windows Server 2012.

Course Objectives

After completing this course, students will be able to:

- Install and configure Windows Server 2016.
- Describe storage in Windows Server 2016.
- Implement directory services.
- Implement Active Directory Federation Services (AD FS).
- Describe networking.
- Implement Hyper-V.
- Configure advanced networking features.
- Implement software-defined networking.
- Implement remote access.
- Deploy and manage Windows and Hyper-V containers.
- Implement failover clustering.
- Implement failover clustering by using virtual machines.

Course Outline

The course outline is as follows:

Module 1, "Installing and configuring Windows Server 2016," explains how to install and perform post-installation configuration of Windows Server 2016 servers.

Module 2, "Overview of storage in Windows Server 2016," explains how to configure storage in Windows Server 2016.

Module 3, "Implementing Directory Services," explains how to implement the Directory Services feature. This module also introduces Microsoft Azure Active Directory.

Module 4, "Implementing AD FS," explains how to implement an AD FS deployment.

Module 5, "Implementing network services," explains how to configure advanced features for Dynamic Host Configuration Protocol (DHCP) and how to configure IP Address Management (IPAM).

Module 6, "Implementing Hyper-V," explains how to install and configure Hyper-V virtual machines.

Module 7, "Configuring advanced networking features," explains how to implement an advanced networking infrastructure.

Module 8, "Implementing Software Defined Networking," explains how to implement software-defined networking.

Module 9, "Implementing remote access," explains how to configure connectivity for remote users by using the DirectAccess feature and virtual private networks (VPN).

Module 10, "Deploying and managing Windows and Hyper-V containers," provides an overview of Windows Server 2016 containers. Additionally, it explains how to deploy, install, configure, and manage containers in Windows Server 2016.

Module 11, "Implementing failover clustering," explains how to implement failover clustering to provide high availability for network services and applications.

Module 12, "Implementing failover clustering with Windows Server 2016 Hyper-V," explains how to deploy and manage Hyper-V virtual machines in a failover cluster.

Course Materials

The following materials are included with your kit:

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



Additional Reading: Course Companion Content on the <http://www.microsoft.com/learning/en/us/companion-moc.aspx> Site: searchable, easy-to-browse digital content with integrated premium online resources that supplement the Course Handbook.

- **Modules:** include companion content, such as questions and answers, detailed demo steps and additional reading links, for each lesson. Additionally, they include Lab Review questions and answers and Module Reviews and Takeaways sections, which contain the review questions and answers, best practices, common issues and troubleshooting tips with answers, and real-world issues and scenarios with answers.
- **Resources:** include well-categorized additional resources that give you immediate access to the most current premium content on TechNet, MSDN, or Microsoft Press.
- **Course evaluation:** at the end of the course, you will have the opportunity to complete an online evaluation to provide feedback on the course, training facility, and instructor.
 - To provide additional comments or feedback on the course, send an email to mcspprt@microsoft.com. To inquire about the Microsoft Certification Program, send an email to mcphelp@microsoft.com.

Virtual Machine Environment

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

Virtual Machine Configuration

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



Note: Pay close attention to the steps at the end of each lab that explain what you need to do with the virtual machines. In most labs, you will revert the virtual machine to the checkpoint that you create during classroom setup. In some labs, you will not revert the virtual machines, but will keep them running for the next lab.

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

Virtual machine	Role
20743A-LON-DC1 (-B, -C)	Domain controller running Windows Server 2016 in the Adatum.com domain
20743A-LON-SVR1	Windows Server 2016 server in the Adatum.com domain
20743A-LON-SVR2	Windows Server 2016 server in the Adatum.com domain
20743A-LON-SVR3	Windows Server 2016 server in the Adatum.com domain
20743A-LON-CORE	Windows Server 2016 server running the Server Core installation
20743A-NANO-SVR1	Windows Server 2016 Nano Server
20743A-LON-CL1	Client computer running Windows 10 and Microsoft Office 2016
20743A-LON-CL2	Client computer running Windows 10 and Office 2016
20743A-LON-CL3	Client computer running Windows 10 and Office 2016
20743A-LON-RTR	Stand-alone Windows Server 2016 computer used as a router
20743A-TREY-DC1	Domain controller running Windows Server 2016 in the TreyResearch.net domain
20743A-INET1	Stand-alone Windows Server 2016 computer used as a simulated Internet host
20743A-LON-HOST1	Virtual hard disk with Windows Server 2016 used for boot from virtual hard disk exercises
20743A-LON-HOST2	Virtual hard disk with Windows Server 2016 used for boot from virtual hard disk exercises

Virtual machine	Role
20743A-LON-NVHOST3	Nested host Windows Server 2016 virtual machine
20743A-LON-NVHOST4	Nested host Windows Server 2016 virtual machine

Software Configuration

The following software is installed on each VM:

- Either Windows Server 2016 or Windows 10
- For client virtual machines, Office 2016

A Microsoft Azure trial subscription is required.

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.

Hardware Level 7

- Intel Virtualization Technology (Intel VT) or AMD Virtualization (AMD-V) processor
- Dual 120-gigabyte (GB) hard disks 7,200 RPM Serial ATA (SATA) or better*
- 16 GB of RAM
- DVD drive
- Network adapter
- Super VGA (SVGA) 17-inch monitor
- Microsoft mouse or a compatible pointing device
- Sound card with amplified speakers

* Striped

Additionally, the instructor's computer must be connected to a projection display device that supports SVGA 1024 × 768 pixels and 16-bit colors.

Module 1

Installing and configuring Windows Server 2016

Contents:

Module Overview	1-1
Lesson 1: Introducing Windows Server 2016	1-2
Lesson 2: Installing Windows Server 2016	1-7
Lesson 3: Configuring Windows Server 2016	1-15
Lab: Installing and configuring Nano Server	1-21
Lesson 4: Preparing for upgrades and migrations	1-26
Lesson 5: Migrating server roles and workloads	1-34
Lesson 6: Windows Server activation models	1-37
Module Review and Takeaways	1-40

Module Overview

To help ensure that your organization obtains the maximum benefit from implementing the Windows Server 2016 operating system, you should be familiar with its many new and improved features. This module introduces you to Windows Server 2016 and describes how to install it, how to perform post-installation configuration tasks, and how to configure it to support remote management. This module also introduces the concept of enabling and configuring Windows container support in Windows Server 2016. You will also learn how to plan a server and migration strategy, and how to migrate server roles and workloads. Finally, you will learn how to choose the most appropriate activation model for your organization.

Objectives

After completing this module, you will be able to:

- Describe the new and improved features in Windows Server 2016.
- Install Windows Server 2016.
- Configure and manage Windows Server 2016.
- Plan a server upgrade and migration strategy.
- Perform a migration of server roles and workloads within a domain and across domains.
- Choose an appropriate activation model.

Lesson 1

Introducing Windows Server 2016

Knowing the capabilities of the Windows Server 2016 operating system enables you to use it effectively and take full advantage of what it can offer your organization. Some of the many improvements to Windows Server 2016 include increased scalability and performance; improved virtualization; improved management tools; and additional deployment options, including Nano Server. This lesson explores these new features and capabilities in Windows Server 2016.

Lesson Objectives

After completing this lesson, you will be able to:

- Select a suitable Windows Server 2016 edition.
- Describe the new and improved features of Windows Server 2016.
- Explain how Windows Server 2016 eases integration with Microsoft cloud services.

Selecting a suitable Windows Server edition

You can choose one of several editions of Windows Server 2016. These editions allow organizations to select the Windows Server 2016 version that best meets their needs, rather than pay for features that they do not require. When system administrators deploy a server for a specific role, they can reduce costs by selecting the appropriate edition. The following table describes the Windows Server 2016 editions.

- Windows Server 2016 Standard
- Windows Server 2016 Datacenter
- Windows Server 2016 Foundation
- Windows Server 2016 Essentials
- Microsoft Hyper-V Server 2016
- Windows Storage Server 2016 Workgroup
- Windows Storage Server 2016 Standard

Edition	Description
The Windows Server 2016 Standard edition	Provides many of the roles and features available for the Windows Server 2016 operating system. This edition supports up to 64 sockets and up to 4 terabytes (TB) of random access memory (RAM). It includes two virtual machine licenses.
The Windows Server 2016 Datacenter edition	Provides all of the roles and features available for the Windows Server 2016 operating system. This edition supports up to 64 sockets, up to 640 processor cores, and up to 4 TB of RAM. It includes unlimited Windows Server–based virtual machine licenses for virtual machines that run on the same hardware.
The Windows Server 2016 Foundation edition	Designed for small businesses. This edition allows only 15 users, precludes being joined to a domain, and includes limited server roles. It supports one processor core and up to 32 gigabytes (GB) of RAM.
The Windows Server 2016 Essentials edition	<p>Corresponds to Windows Small Business Server from earlier versions of Windows Server. This edition is available in two forms:</p> <ul style="list-style-type: none"> • As an installable server role in an existing domain. • As a core Windows Server edition on a virtual machine, by using a wizard. <p>This edition cannot function as a virtualization server, failover clustering server, Server Core server, or Remote Desktop Services server. It allows 25 users and 50 devices. It supports two processor cores and up to 64 GB of RAM. The new features and improvements include client deployment, user management, storage and data protection, and Microsoft Office 365 integration.</p>

Edition	Description
Microsoft Hyper-V Server 2016	Acts as a standalone virtualization server for virtual machines. The host operating system has no licensing cost (it is free), but virtual machines must be licensed separately. This edition supports up to 64 sockets and up to 4 TB of RAM. It supports domain joining. It does not support Windows Server 2016 roles other than limited file service features. This edition has no GUI, but it does have a UI that displays a menu of configuration tasks.
The Windows Storage Server 2016 Workgroup edition	Acts as an entry-level unified storage appliance. This edition allows 50 users, one processor core, and 32 GB of RAM. It supports domain joining.
The Windows Storage Server 2016 Standard edition	Supports up to 64 sockets but is licensed on a two-socket, incrementing basis. This edition supports up to 4 TB of RAM. It includes two virtual machine licenses. It supports domain joining. It supports some roles, including Domain Name System (DNS) and Dynamic Host Configuration Protocol (DHCP) server roles, but it does not support others, including Active Directory Domain Services (AD DS), Active Directory Certificate Services (AD CS), or Active Directory Federation Services (AD FS).



Note: At the time of this writing, these details had not been finalized.

What is new since Windows Server 2008 was released?

Windows Server 2016 provides many new features and a number of significant improvements over earlier versions of Windows Server. Some of these features and improvements were first introduced in Windows Server 2012 or Windows Server 2012 R2, whereas others are new to Windows Server 2016.

New features and improvements introduced in Windows Server 2012 or Windows Server 2012 R2

The following features and feature improvements in Windows Server 2016 were first introduced in Windows Server 2012 or Windows Server 2012 R2:

New features and improvements introduced in Windows Server 2012 or Windows Server 2012 R2:

- Work Folders
 - DHCP failover
 - IPAM
 - Dynamic Access Control
 - Data Deduplication
- Storage Spaces
Storage tiers
Better domain controller virtualization
Cloning virtual domain controllers

- **Work Folders.** Provides a mechanism for both domain-joined computers and those that are not domain joined to access and synchronize corporate data files.
- **DHCP failover.** Enables you to deploy two DHCP servers containing overlapping DHCP scopes. If a DHCP server goes offline, DHCP client computers can renew their IP configurations from the failover DHCP server.
- **IP Address Management (IPAM).** Provides administrative and monitoring capabilities for the IP address infrastructure within your organization's networks. With IPAM, you can monitor, audit, and manage servers running DHCP and DNS.

- **Dynamic Access Control.** This claims-based authorization platform enables you to control access to file resources within your organization. This is in addition to any folder or shared folder permissions already protecting the resource. Dynamic Access Control enables you to apply access control permissions based on rules that can include the sensitivity of the resources, the user's job or role, and the configuration of the device that is used to access these resources.
- **Data Deduplication.** Involves finding and removing duplication within data. By segmenting files into small, variable-sized pieces, identifying duplicate pieces, and maintaining a single copy of each piece, Data Deduplication enables you to store more data in less space.
- **Storage Spaces.** Enables cost-effective, highly available, scalable, and flexible storage for critical deployments. Storage Spaces are based on virtual disks that are created from free space in a storage pool. Storage pools are collections of physical disks that enable you to aggregate disks, expand capacity in a flexible manner, and delegate administration.
- **Storage tiers.** Automatically moves frequently accessed data to faster storage and less frequently accessed data to slower storage.
- **Better support for domain controller virtualization.** Although many organizations have virtualized domain controllers for several years, potential issues can affect the reliability of this configuration. A feature, known as GenerationID, changes whenever the virtual machine experiences an event that affects its position in time. During startup and normal operations, a virtual domain controller compares the current value of GenerationID against the expected value. A mismatch is interpreted as a rollback event, and the domain controller employs safeguards to prevent the virtual domain controller from creating duplicate security principals.
- **The ability to clone virtual domain controllers.** Enables you to deploy new virtual domain controllers by cloning existing ones.



Note: This is not a complete list of all the new or improved features in Windows Server 2012 or Windows Server 2012 R2.

New features and improvements introduced in Windows Server 2016

The following features and feature improvements were introduced in Windows Server 2016:

- **Nano Server.** Nano Server is a new installation option for Windows Server 2016. Be it has no graphical or command prompt interface, it has a significantly lower hardware requirement than even Server Core. Nano Server is the ideal platform for Hyper-V, Hyper-V cluster, and scale-out file servers, and cloud service apps.
- **Windows Server containers and Hyper-V containers.** Containers enable you to isolate your apps from the operating system environment. This improves security and reliability. Windows containers are isolated from one another but run on the host operating system. Hyper-V containers are further isolated, because they run within a virtual machine.
- **Docker.** Docker is a technology for managing containers. Although Docker is usually associated with Linux, Windows Server 2016 provides support for Docker as a means for managing Windows containers and Hyper-V containers.
- **Rolling upgrades for Hyper-V and storage clusters.** These upgrades enable you to add Windows Server 2016 nodes to an existing Windows Server 2012 R2 failover cluster. The cluster continues to operate at a Windows Server 2012 R2 functional level until all the nodes are upgraded.
- **The ability to hot add and hot remove virtual memory and network adapters from virtual machines.** In Hyper-V in Windows Server 2016, you can now add or remove virtual memory and network adapters while the virtual machines are running.

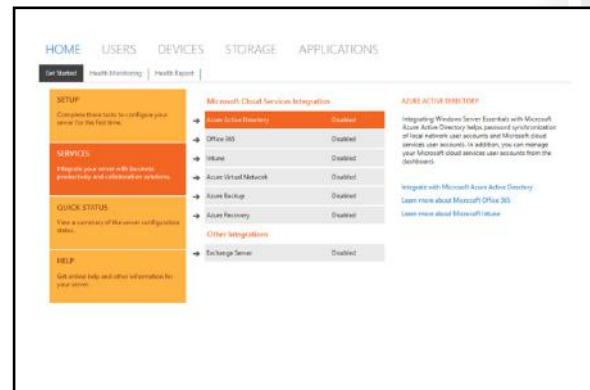
- **Nested virtualization.** In Hyper-V in Windows Server 2016, you can enable nested virtualization, which allows you to run Hyper-V virtual machines within a virtual machine.
- **PowerShell Direct.** This feature enables you to run Windows PowerShell commands against a guest operating system in a virtual machine without handling security policies, host network settings, or firewall settings.
- **Shielded virtual machines.** Shielding your virtual machines enables you to help protect the data on them from unauthorized access.
- **Windows Defender.** Windows Defender is provided to help protect your server against malware. Although the Windows Defender interface is not installed by default, the antimalware patterns are automatically kept up to date.
- **Storage Spaces Direct.** This feature enables you to build highly available storage with directly attached disks on each node in a cluster. The Server Message Block 3 (SMB3) protocol provides resiliency.
- **Storage Replica.** This feature enables you to synchronously or asynchronously replicate volumes at the block level.
- **Microsoft Passport.** This service replaces passwords with two-factor authentication that consists of an enrolled device and a Windows Hello (biometric) authentication or PIN. This helps provide a more secure and convenient sign-in experience.

Integration with Microsoft cloud services

Organizations are increasingly moving some or all of their IT infrastructure and services to the cloud. Whereas large organizations usually have the technical skill set to integrate and migrate their on-premises infrastructure to the cloud, small and medium-sized businesses might struggle to accomplish the same tasks.

Windows Server 2016 seeks to simplify the process for integration with Microsoft cloud services.

One way that Windows Server 2016 accomplishes this is through the provision of the Windows Server Essentials Experience server role. This role provides easy access to cloud service integration features.



The Windows Server Essentials Experience server role

To begin the process of integrating your server running Windows Server 2016 with Microsoft cloud services, start by installing the Windows Server Essentials Experience server role. Use the following procedure:

1. On your server, in **Server Manager**, click **Add roles and features**.
2. Click **Next** twice, select the appropriate server from the list, and then click **Next**.
3. In the **Roles** list, select **Windows Server Essentials Experience**, and then click **Next**.
4. Click **Add Features**, and then click **Next** twice.
5. Click **Next**, and then click **Install**.

6. When prompted, click **Close**.
7. In **Server Manager**, click the notification flag, and then click **Configure Windows Server Essentials**.
8. In the **Configure Windows Server Essentials Wizard**, click **Configure**.
9. After your server is configured, click **Close**.
10. Open the **Windows Server Essentials Experience** console from your server desktop.

You can now configure integration with online services. From the console, click **SERVICES**. You can configure integration with the following Microsoft cloud services:

- Microsoft Azure Active Directory. Enables user name and password synchronization from the local AD DS to Azure Active Directory (Azure AD).
- Office 365. Enables integration with Office 365, although the user name and password integration is handled by Azure AD.
- Microsoft Intune. Enables you to assign Intune licenses to your local directory to enable the management of users' devices.
- Azure Virtual Network. Enables you to set up and configure hybrid deployments, with part of your infrastructure in the cloud and the rest on-premises.
- Azure Backup. Enables you to easily set up Azure Backup as a means of providing for data recovery from the Microsoft cloud platform.
- Azure Recovery. Enables you to replicate recovery information to the Microsoft cloud platform.
- Microsoft Exchange Server. Enables you to configure integration between your local Exchange Server environment and Microsoft cloud services.

To configure integration, click the appropriate service, and then click the appropriate link.



Note: You must have an appropriate Microsoft cloud services subscription to configure these integrations.

Verify the correctness of the statement by placing a mark in the column to the right.

Statement	Answer
Docker is a container that enables you to run an app in an isolated and portable operating environment.	

Lesson 2

Installing Windows Server 2016

When you prepare to install Windows Server 2016, you must understand whether a particular hardware configuration is suitable. You must also select among the installation options: Windows Server 2016 (Desktop Experience), Server Core, or Nano Server. This lesson describes each of these installation options and provides guidance on how to perform an installation of Windows Server 2016.

Lesson Objectives

After completing this lesson, you will be able to:

- List the hardware requirements for Windows Server 2016.
- Describe the installation options for Windows Server.
- Describe Server Core.
- Describe Nano Server.
- Select a suitable installation type.
- Install Nano Server.

Hardware requirements

The hardware requirements define the absolute minimum required to run the server software. The actual hardware requirements depend on the services that the server is hosting, the load on the server, and how responsive you want the server to be. The services and features of each role put a unique load on network, disk I/O, processor, and memory resources.

Virtualized deployments of Windows Server 2016 must match the same hardware specifications as those required for physical deployments. Windows Server 2016 is supported on the Hyper-V virtualization platform and on certain non-Microsoft virtualization platforms.

The following table shows the minimum hardware requirements for Windows Server 2016.

Component	Requirement
Processor architecture	64-bit
Processor speed	1.4 gigahertz (GHz)
RAM	512 MB
Hard disk space	32 GB

Windows Server 2016 has the following minimum hardware requirements:

Hardware	Requirement
Processor architecture	x64
Processor speed	1.4 GHz
RAM	512 MB
Hard disk space	32 GB

Overview of installation options

When you install Windows Server 2016, you can select one of three installation options:

- Windows Server 2016 (Desktop Experience). This is a full server installation and includes a complete graphical management interface. This installation option supports all Windows Server roles.
- Windows Server 2016. This is the equivalent of Server Core in earlier versions of Windows Server and provides for a command-line management interface. This installation option has a reduced hardware footprint but does not support all Windows Server roles.
- Nano Server. This is a new installation option for which Windows Server 2012 and earlier versions have no equivalent. Nano Server is administered remotely, and it is optimized for hosting in private clouds and datacenters and for running applications that are developed by using cloud application patterns.

You can choose among the following installation options when deploying Windows Server 2016:

- Windows Server 2016 (Desktop Experience):
 - Full server installation
- Windows Server 2016:
 - Server Core installation
- Nano Server:
 - Minimal server installation



Note: You cannot directly select the Nano Server installation from the installation media during setup but must manually build it by using Windows PowerShell cmdlets from the installation media.

Nano Server uses fewer hardware resources than even Server Core, but it supports fewer server roles.

What is Server Core?

Server Core is an installation option for Windows Server 2016 that uses fewer hardware resources than the full installation option. Server Core does not come installed with a GUI for management purposes, but you can manage Server Core locally by using Windows PowerShell or a command-line interface, or you can manage it remotely by using one of the remote management options. Remote management is covered later in this module. A Windows Server 2016 Server Core installation offers fewer components and administrative management options than the Windows Server 2016 (Desktop Experience) installation option.

- Server Core is:
 - A more security-enhanced, less resource-intensive installation option
 - An installation that cannot be converted to full graphical shell version of Windows Server 2016
 - The default installation option for Windows Server 2016
 - Managed locally by using `sconfig.cmd` and Windows PowerShell
- With remote management enabled, you rarely need to sign in locally

Server Core is the default installation option when you install Windows Server 2016.



Note: Unlike with some earlier versions of Windows Server, when you choose Server Core or Server with Desktop Experience at the time of installation, you cannot later convert to the other mode.

The Server Core option has the following advantages over the full Windows Server 2016 installation option:

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

To perform most important management tasks on Server Core, you must use remote management. However, when you are connected locally, you can use the tools that are listed in the following table to manage Server Core deployments of Windows Server 2016.

Tool	Function
Cmd.exe	Allows you to run traditional command-line tools, such as ping.exe, ipconfig.exe, and netsh.exe.
PowerShell.exe	Launches a Windows PowerShell session on the Server Core deployment. You then can perform Windows PowerShell tasks normally. Windows Server 2016 comes with Windows PowerShell version 5.0 installed.
Sconfig.cmd	Functions as a command-line, menu-driven administrative tool that enables you to perform most common server administrative tasks, such as configuring networking, workgroups, and domains, and configuring Windows Firewall.
Regedt32.exe	Provides registry access within the Server Core environment.
Msinfo32.exe	Allows you to view system information about the Server Core deployment.
Taskmgr.exe	Launches Task Manager.

Server roles available in Server Core

The following server roles are available on Server Core deployments:

- AD CS
- AD DS
- DHCP Server
- DNS Server
- File Services (including File Server Resource Manager)
- Active Directory Lightweight Directory Services (AD LDS)
- Hyper-V
- Print and Document Services
- Streaming Media Services
- Web Server (including a subset of ASP.NET)
- Windows Server Update Server

- Active Directory Rights Management Server
- Routing and Remote Access Server and the following sub-roles:
 - Remote Desktop Connection Broker
 - Licensing
 - Virtualization

What is Nano Server?

Nano Server is a new installation option for Windows Server 2016 that is similar to Windows Server in Server Core mode. However, although it has a significantly smaller hardware footprint, it has no local sign-in capability and it supports only 64-bit apps, tools, and agents. Setup is significantly faster, and after installation, the operating system requires far fewer updates.



Note: Nano Server is not available for selection through the Windows Server 2016 setup wizard. Instead, you must create a virtual hard disk by using Windows PowerShell. You can then use this virtual hard disk on a virtual machine to support a virtualized Nano Server in Hyper-V, or you can configure your server computer to start from a VHD file for a physical Nano Server deployment option.

Nano Server is ideal for use in the following scenarios:

- Compute host for Hyper-V virtual machines, either in clusters or not
- Storage host for a scale-out file server, either in clusters or not
- DNS server
- Web server running IIS
- Host for apps that are developed by using cloud application patterns and run in a container or virtual machine

Use scenarios

Nano Server is ideal for use in the following scenarios:

- As a *compute* host for Hyper-V virtual machines, either in clusters or not in clusters.
- As a storage host for a scale-out file server, either in clusters or not in clusters.
- As a DNS server.
- As a web server running Microsoft Internet Information Services (IIS).
- As a host for applications that are developed by using cloud application patterns and run in a container or virtual machine guest operating system.

Server roles available in Nano Server

The following table shows the server roles and features that you can either install when you deploy Nano Server or subsequently install by using Windows PowerShell on a previously deployed Nano Server.

Role	Option to install
Hyper-V role	-Compute
Failover clustering	-Clustering
Hyper-V guest drivers for hosting Nano Server as a virtual machine	-GuestDrivers

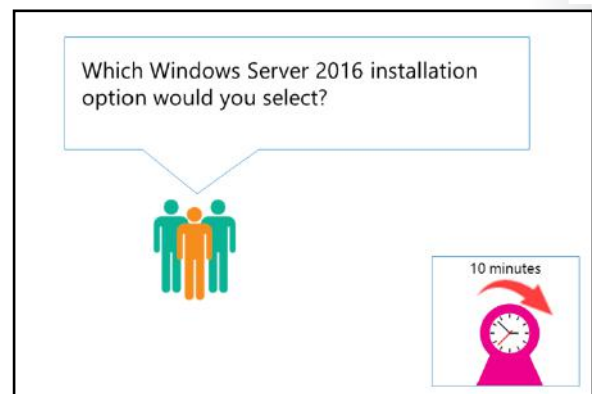
Role	Option to install
Basic drivers for a variety of network adapters and storage controllers. This is the same set of drivers included in a Server Core installation of Windows Server 2016.	-OEMDrivers
File Server role and other storage components	-Storage
Windows Defender Antimalware, including a default signature file	-Defender
Reverse forwarders for application compatibility—for example, common application frameworks, such as Ruby and Node.js	-ReverseForwarders
DNS Server role	-Packages Microsoft-NanoServer-DNS-Package
Desired State Configuration	-Packages Microsoft-NanoServer-DSC-Package
IIS	-Packages Microsoft-NanoServer-IIS-Package
Host support for Windows Containers	-Containers
System Center Virtual Machine Manager agent	-Packages Microsoft-Windows-Server-SCVMM-Package -Packages Microsoft-Windows-Server-SCVMM-Compute-Package
Network Performance Diagnostics Service (NPDS)	-Packages Microsoft-NanoServer-NPDS-Package
Data Center Bridging	-Packages Microsoft-NanoServer-DCB-Package

Discussion: Selecting a suitable Windows Server edition and installation type

Question: Which Windows Server 2016 installation option would you select?

Question: Your customer, a small legal firm, has a requirement for a single server that the firm wants you to deploy at their only office. Which Windows Server 2016 installation option would be the most suitable?

Question: One of your enterprise customers has a new branch office. You must deploy Windows Server 2016 to support the local users at this new branch. The server will be managed remotely from IT staff located in the head office. The server will support the DNS, DHCP, and AD DS server roles. Your customer wants to minimize resource consumption on the server. Which Windows Server 2016 installation option would be best?



Question: Your customer wants to run a web server based on IIS. The server must use as few hardware resources as possible. Which Windows Server 2016 installation option would be best?

Demonstration: Installing Nano Server

Overview of the Nano Server installation process

You can install Nano Server as a guest operating system in Hyper-V or as an operating system on a physical computer. The initial process for these two installation options is broadly the same, and starts with creating a virtual hard disk. Use the following guidance to create a virtual hard disk for Nano Server:

1. Obtain and mount the Windows Server 2016 product DVD (or an ISO image of the DVD) on a Windows Server 2016 or Windows 10 computer.
2. Copy **NanoServerImageGenerator.psm1**, and then **Convert-WindowsImage.ps1** from the **\NanoServer\NanoServerImageGenerator** folder on the Windows Server DVD to a local folder on your computer. The **NanoServerImageGenerator** folder contains a .wim file that stores the Nano Server installation image.
3. Open **Windows PowerShell**, and then run the cmdlet **Import-Module NanoServerImageGenerator.psm1**. This imports the necessary Windows PowerShell modules required for the next step.
4. Create a virtual hard disk (either a VHD or a VHDX file) for the Nano Server. You can set a computer name and include then necessary drivers, optionally including the Hyper-V guest drivers, by running the cmdlet **New-NanoServerImage -Edition Standard -MediaPath <path to root of media> -BasePath .\Base -TargetPath .\NanoServerVM\NanoServerVM.vhd -ComputerName <computer name> -DeploymentType <option>** where:
 - **<path to root of media>** is the path that you provide to the root of the contents of the Windows Server DVD.
 - **-BasePath** specifies a folder that will be created to contain a copy of the Nano Server .wim file and packages.
 - **-TargetPath** specifies the full path, including the file name and extension, where the resulting VHD or VHDX will be created.
 - **-Computer_name** is the computer name you provide for the Nano Server you are creating.
 - **-DeploymentType** Guest or host. Use guest to configure the VHD for use as a guest on another host.

The next steps depend on whether you intend to run Nano Server as a guest in Hyper-V or as an operating system on a physical computer.

Deploying Nano Server as a Hyper-V guest


1. Copy the virtual hard disk you just created (using the **-DeploymentType guest** option) to the appropriate location for virtual hard disks on your Hyper-V host.
2. In Hyper-V, create a virtual machine and configure it to use the virtual hard disk you previously created.
3. Start the virtual machine, and then sign in by using the user name and password you are prompted to create when running the preceding script.

4. Use the Nano Server Recovery Console to configure the basic settings for the computer so that you can enable remote management for the server. These settings typically include the Internet Protocol version 4 (IPv4) and Windows Firewall rule settings.
5. Connect to your Nano Server remotely and complete the configuration process.

Deploying Nano Server on a physical host

1. Copy the virtual hard disk you just created (using the **-DeploymentType host** option) to an appropriate location on the host that you intend to use as a Nano Server.
2. Modify the startup settings of the target computer to use the virtual hard disk:
 - a. In an elevated command prompt, copy the current boot loader entry on the computer where Nano Server will be running, and then use that to create a new entry with the **bcdedit /copy {current} /d "Nano Server"** command.
 - b. Type **bcdedit** again, and then copy the GUID, including the braces ({ }), that appears in the **ID** field of the newly copied boot loader entry.
 - c. Run these commands, replacing {GUID} with the copied GUID, including the braces:


```
bcdedit /set {GUID} device vhd=[c:]\NanoServer\NanoServer.vhd
bcdedit /set {GUID} osdevice vhd=[c:]\NanoServer\NanoServer.vhd
bcdedit /set {GUID} path \windows\system32\boot\winload.exe
```
 - d. Verify the setting is correctly set by going to **System Properties**. Select **Startup and Recovery settings**. You should see that **System startup** lists **Nano Server**.

 **Note:** You can also mount the NanoServer.vhd file by using **Disk Management**. Assuming the drive letter allocated to the mounted NanoServer.vhd file is G, you can run the **bcdboot G:\Windows** command from an elevated command prompt to achieve the same result as the preceding sequence of commands.

3. Start the physical computer that is running Nano Server, and then sign in by using the user name and password you are prompted to create when running the preceding script.
4. Use the Nano Server Recovery Console to configure the basic settings for the computer so that you can enable remote management for the server. These settings typically include the IPv4 and Windows Firewall rule settings.
5. Connect to your Nano Server remotely, and then complete the configuration process.

Demonstration Steps

1. On **LON-DC1**, open an elevated command prompt.
2. Change to the root directory of drive C, and then create a folder called **Nano**.
3. Copy all the files with a **.ps*** extension from the **d:\NanoServer\NanoServerImageGenerator** folder to **C:\Nano**.
4. Open an elevated Windows PowerShell window.
5. Run **Import-Module c:\nano\NanoServerImageGenerator.psm1**. This command imports the required Windows PowerShell module for Nano Server.

6. Run **new-NanoServerImage -Edition Standard -mediapath D:\ -Basepath c:\nano -targetpath c:\nano\nano-svr1.vhdx -DeploymentType Guest -computername NANO-SVR1 -storage -packages Microsoft-NanoServer-IIS-Package**. Type the password **Pa\$\$w0rd** when prompted. This command creates a VHDX file for your Nano Server with the following options:
 - **Mediapath** identifies the source of the installation files.
 - **Basepath** indicates where to create the VHDX file and supplemental files.
 - **Targetpath** identifies the name and location of the VHDX file.
 - **Computername** identifies the name of this instance of Nano Server.
 - **Storage** installs the File Server role.
 - **Packages** enables the additional installation of other roles—in this case, the IIS role.
 - **DeploymentType** configure the VHDX for use as a guest.
7. In C:\Nano, you can see the files that were created, including the Nano-svr1.vhdx file. Ordinarily, you now copy this file to a Hyper-V host, and then create a virtual machine to use the virtual hard disk. You can also reconfigure startup settings on your host so that it can start from this VHDX file. A virtual machine is preconfigured with the VHDX file. Switch to **NANO-SVR1**.
8. Sign in as **Administrator/Pa\$\$w0rd**.
9. By using this console, you can perform basic administration of Nano Server, including making basic changes to IP configuration and firewall settings, enabling the computer to be managed remotely.
10. Observe that the computer name is Nano-Svr1, and that the computer belongs to a workgroup.
11. In **Network Adapter Settings**, notice that DHCP is obtaining the IP configuration. Make a note of the IP address. You will need this in a later demonstration.

Check Your Knowledge

Question	
Which of the following tools can you use to locally manage an installation of Windows Server 2016 Nano Server?	
Select the correct answer.	
	PowerShell.exe
	Sconfig.cmd
	Taskmgr.exe
	All of the above
	None of the above

Lesson 3

Configuring Windows Server 2016

The installation process for Windows Server 2016 requires minimal input from the installer. However, after you complete the installation, you must configure several important settings before you can use your server. In addition, because both Server Core and Nano Server provide no graphical management tools and, in the case of Nano Server, do not provide a command prompt for management, you must know how to enable and perform the remote management of your server infrastructure. This lesson identifies the important post-installation configuration options, and explains how to enable and use the remote management tools.

Lesson Objectives

After completing this lesson, you will be able to:

- List and explain the settings that must be configured following the installation.
- Manage remote installations of Windows Server.
- Manage Windows Server by using Windows PowerShell.
- Manage Nano Server.
- Use remote tools to manage Nano Server.

Post-installation configuration settings

In earlier versions of Windows Server, the installation process required you to configure network connections, the computer name, user accounts, and domain membership information. The Windows Server 2016 installation process reduces the number of questions that you must answer. The only information that you provide during installation is the password that is used by the default local Administrator account.

After you have installed Windows Server 2016, you typically should complete the following:

- Configure the IP address.
- Set the computer name.
- Join an Active Directory domain.
- Configure the time zone.
- Enable automatic updates.
- Add roles and features.
- Enable the Remote Desktop feature.
- Configure Windows Firewall settings.

After you have installed Windows Server 2016, you must complete the following:

- Configure the IP address
- Set the computer name
- Join an Active Directory domain
- Configure the time zone
- Enable automatic updates
- Add roles and features
- Enable Remote Desktop
- Configure Windows Firewall settings

The installation type selected (with Desktop Experience or without) during setup determines which tools you can use to complete these configuration tasks. For example, on Windows Server 2016 (Desktop Experience), you can use Server Manager to complete these post-installation tasks. On Server Core, you

can use Sconfig.cmd or other command-line tools, such as Netsh.exe, or enable remote management and then complete these tasks by using Windows PowerShell. Likewise, with Nano Server, you must use remote management tools to perform these tasks.



Note: You can also use an XML answer file to provide this information during an automated installation.

Managing servers remotely

Performing the interactive management of Windows Server is not considered a best practice. With Server Core and, to a greater extent, Nano Server, your local management options are very limited. After you have configured the network and firewall settings of Server Core or Nano Server, you must perform other management tasks remotely.

You can use the following options to manage a computer that is running Windows Server 2016 remotely, regardless of whether it is running Windows Server 2016 (Desktop Experience), Server Core, or Nano Server:

You can use the following options to remotely manage a computer that is running Windows Server 2016:

- Server Manager
- Windows PowerShell Remoting
- Remote Desktop
- Management console

- **Server Manager.** You can add a remote server to the Server Manager console that is on a server running a full installation of Windows—that is, with Desktop Experience. You can also manage a remote server computer from a Windows 10 workstation with Remote Server Administration Tools installed. You then can use Server Manager to manage the server roles running on the remote server computer.



Note: To download Remote Server Administration Tools, refer to: <http://aka.ms/wzpq0j>

- **Windows PowerShell remoting.** You can use Windows PowerShell to run Windows PowerShell commands or scripts against correctly configured remote servers if the script is hosted on the local server. With Windows PowerShell remoting, where necessary, you also can locally load Windows PowerShell modules, such as Server Manager, and run the cmdlets available in that module against appropriately configured remote servers.
- **Remote Desktop.** You can connect to a remote server computer that is running the Server Core installation or the full installation by using Remote Desktop. On Server Core, you must enable Remote Desktop by using Sconfig.cmd. You cannot use Remote Desktop to remotely manage Nano Server.
- **A management console.** For most server roles, you can add a remote server computer to a management console that is running on another computer. For example, you can add remote computers to Computer Management.

Windows Remote Management


Windows Remote Management (WinRM) is a collection of technologies that enables administrators to manage server hardware when signed in directly or over the network. With WinRM, you can use the following tools to manage a computer remotely:

- Remote Shell. Enables you to run command-line utilities against correctly configured remote servers if the command prompt utility is present on the remote server.
- Windows PowerShell remoting. Enables you to run Windows PowerShell commands or scripts against correctly configured remote servers if the script is hosted on the local server. Windows PowerShell remoting also enables you to locally load Windows PowerShell modules, such as Server Manager, and to run the cmdlets available in that module against suitably configured remote servers.
- Remote management console.

Using Windows PowerShell to manage servers

Windows PowerShell is a scripting language and command-line interface that is designed to assist you in performing day-to-day administrative tasks. Windows PowerShell constitutes cmdlets that you run at a Windows PowerShell command prompt or combine into Windows PowerShell scripts. With the introduction of Nano Server, a headless server environment, the ability to use Windows PowerShell to manage servers remotely is becoming more important.

Windows PowerShell is a scripting language and command-line interface that is designed to assist you in performing day-to-day administrative tasks



Importing modules

Some Windows PowerShell cmdlets are not available in the default Windows PowerShell library. When you enable some Windows features or want to administer particular environments, you must obtain additional Windows PowerShell functions. These additional functions are packaged in modules. For example, to manage Nano Server, Windows Server containers, and Azure AD with Windows PowerShell, you must import the required modules. To do this, use an **import-module** cmdlet:

```
Import-Module NanoServerImageGenerator.psm1
```

The preceding cmdlet imports the required Windows PowerShell module for Nano Server in preparation for performing additional Nano Server management by using Windows PowerShell remoting.

Windows PowerShell remote management

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

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

PowerShell Direct

Many administrators choose to run some or all of their servers running Windows Server in virtualized environments. To enable a simpler administration of Windows Server Hyper-V virtual machines, Windows 10 and Windows Server 2016 both support a new feature called PowerShell Direct.

PowerShell Direct enables you to run a Windows PowerShell cmdlet or script inside a virtual machine from the host operating system without regard to network and firewall configurations and regardless of remote management configuration.



Note: You must still authenticate to the virtual machine by using guest operating system credentials.

To use PowerShell Direct, from your host, run the following Windows PowerShell cmdlet.


```
Enter-PSSession -VMName VMName
```

You can then run the same cmdlets that you normally run in the same way as with any other remote Windows PowerShell situation.

Managing Nano Server

You can perform only the most fundamental management tasks interactively on Nano Server. After you have signed in, the Nano Server Recovery Console displays. This identifies the following:

- The computer name
- The workgroup or domain name
- The installed operating system
- Local data, the local time, and the time zone
- The current network configuration



Configuring networking

You can change the basic network configuration by using the Tab key to navigate to **Networking** and then pressing Enter. You can then select the appropriate network adapter from the list by using the cursor keys to navigate to the correct adapter and then pressing Enter.

The current network settings are displayed. You can press either F11 to configure IPv4 settings or F12 for Internet Protocol version 6 (IPv6) settings. If you choose to configure IPv4, use the F4 key to switch the settings. For example, to enable or disable DHCP, press F4. To enter a manual IPv4 configuration, disable DHCP and then use the number keys to type a suitable IP address, subnet mask, and default gateway. Press Enter twice to update the configuration. Press Esc repeatedly to return to the main menu.

Configuring the firewall

You might need to configure firewall settings to enable remote management. From the main Nano Server Recovery Console, press the Tab key to navigate to **Firewall**, and then press Enter. A list of firewall rules is displayed. Use the cursor keys to navigate up and down the list, and press Enter for a rule that you want to configure.

For example, to enable remote event log management, locate the remote event log management (RPC) rule and press Enter. Press F4 to Enable/Disable. Press ESC, and select the next rule and repeat the procedure. When you have configured all rules, press ESC to return to the main menu.

Ongoing management

After you have configured the networking settings and enabled the appropriate remote management firewall ports for inbound communications, you can manage the Nano Server remotely by using either Server Manager, Windows PowerShell, or any other management tool by using the Connect to option to select the Nano Server. Typical management tasks include:

- Adding the computer to a domain
- Adding roles and features to the server

Demonstration: Configuring Nano Server

In this demonstration, you will see how to:

- Add Nano Server to a domain.
- Configure Nano Server from the command line.
- Configure Nano Server by using Server Manager.

Demonstration Steps

Add Nano Server to a domain

1. On **LON-DC1**, in the **Administrator: Windows PowerShell** window, run **djoin.exe /provision /domain adatum /machine nano-svr1 /savefile .\odjblob**. This creates a file that you will use to complete the process of adding Nano Server to the domain.



Note: Replace the IP address 172.16.0.X in the following commands with the IP address you recorded earlier from your Nano Server installation.

2. The following commands are used to enable Windows PowerShell remoting:
 - **Set-Item WSMAN:\localhost\Client\TrustedHosts "172.16.0.X"**
 - **\$ip = "172.16.0.X"**
 - **Enter-PSSession -ComputerName \$ip -Credential \$ip\Administrator**
3. To enable file sharing through the firewall, run **netsh advfirewall firewall set rule group="File and Printer Sharing" new enable=yes**.
4. Close the Windows PowerShell remoting session by running **Exit-PSSession**.
5. Map a network drive to the C drive on Nano Server.
6. Copy **C:\odjblob** to the root of the C drive on Nano Server.
7. Reestablish a Windows PowerShell remoting session to the Nano Server.
8. Run **djoin /requestodj /loadfile c:\odjblob /windowspath c:\windows /localos** to complete the process of adding the computer to the domain.
9. Run **shutdown /r /t 5** to restart Nano Server.

10. On **NANO-SVR1**, sign in as **Administrator/Pa\$\$w0rd** in the domain **Adatum**.
11. In the Nano Server Recovery Console, observe that the computer is in the Adatum.com domain.

Configure Nano Server from the command line

1. On **LON-DC1**, close Windows PowerShell.
2. Open **Windows PowerShell (Admin)**.
3. Run **get-windowsfeature -comp Nano-svr1** to list the installed roles and features on Nano Server.
4. Run **install-windowsfeature Fs-fileserver -comp Nano-svr1** to add the File Server role to Nano Server.
5. Run **get-windowsfeature -comp Nano-svr1** to verify that the role is installed.
6. Enable a Windows PowerShell remoting session with Nano Server:
 - a. Run **\$ip = "172.16.0.X"**.
 - b. Run **Enter-PSSession -ComputerName \$ip -Credential \$ip\Administrator**.
 - c. Run **get-netipaddress** to view the IP configuration for Nano Server.
 - d. Run **bcdedit /enum** to view the startup environment of Nano Server.
 - e. Run **net share** to view the shared folders. Only default shares exist.

Configure Nano Server by using Server Manager

1. In **Server Manager**, add **Nano-SVR1** to the **Computer** list.
2. In **Server Manager**, expand **File and Storage Services**, click **Shares**, and then in the **TASKS** list, click **New Share**.
3. Create a new shared folder:
 - o Type: **SMB Share - Quick**
 - o Server: **nano-svr1**
 - o Share name: **Data**

Check Your Knowledge

Question	
Which of the following commands do you use to initiate remote Windows PowerShell management?	
Select the correct answer.	
	Enter-PSSession -Name
	Enter-PSRemote -Name
	Enter-PSSession -ComputerName
	Enter-PSRemote -ComputerName

Lab: Installing and configuring Nano Server

Scenario

You are responsible for implementing many of the new features in Windows Server 2016. To become familiar with the new operating system, you decide to install a new server running Windows Server 2016 and complete the post-installation configuration tasks.

Objectives

After completing this lab, you will be able to:

- Install the Nano Server option for Windows Server 2016.
- Configure Nano Server.

Lab Setup

Estimated Time: 60 minutes

Virtual machines: **20743A-LON-DC1** and **20743A-NANO-SVR1**

User name: **Adatum\Administrator**

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

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

1. On the host computer, start **Hyper-V Manager**.
2. In Microsoft Hyper-V Manager, click **20743A-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 and 3 for **20743A-NANO-SVR1**.
6. On **20743A-LON-DC1**, in the virtual machine connection window, click **Media**, point to **DVD Drive**, and then click **Insert Disk**.
7. Browse to **D:\Program Files\Microsoft Learning\20743\Drives**, and then select **WinServer2016_TP5.iso**.
8. Click **Open**.

Exercise 1: Installing Nano Server

Scenario

You determine that Nano Server offers you the best installation option and decide to deploy a web server that uses Nano Server.

The main tasks for this exercise are as follows:

1. Copy the required Windows PowerShell scripts.
2. Import Windows PowerShell modules.
3. Create a virtual hard disk.
4. Sign in to the NANO-SVR1 virtual machine.

► Task 1: Copy the required Windows PowerShell scripts

1. On **LON-DC1**, open an elevated Windows PowerShell prompt.
2. Change to the root directory of drive C, and then create a folder named **Nano**.
3. Copy all the files with a **.ps*** extension from the **D:\NanoServer\NanoServerImageGenerator** folder to **C:\Nano**.

► Task 2: Import Windows PowerShell modules

- Run **Import-Module c:\nano\NanoServerImageGenerator.psm1**. This command imports the required Windows PowerShell module for Nano Server.

► Task 3: Create a virtual hard disk

1. Run **new-NanoServerImage -Edition Standard -mediapath D:\ -Basepath c:\nano -targetpath c:\nano\nano-svr1.vhdx -DeploymentType Guest -computername NANO-SVR1 -storage -packages Microsoft-NanoServer-IIS-Package**, and when prompted, type the password **Pa\$\$w0rd**.
2. Verify that **C:\Nano** contains a file called **nano-svr1.vhdx**.



Note: Normally, you would now create a virtual machine to use the nano-svr1.vhdx file. However, to expedite the process, you will start a virtual machine that has already been created.

► Task 4: Sign in to the NANO-SVR1 virtual machine

- On **NANO-SVR1**, sign in by using the following credentials:
 - User name: **Administrator**
 - Password: **Pa\$\$w0rd**

Results: After completing this exercise, you will have successfully created the required virtual hard disk for Nano Server.

Exercise 2: Completing post-installation tasks on Nano Server

Scenario

You must now complete the installation of Nano Server by configuring the post-installation settings and joining it to the Adatum.com domain.

The main tasks for this exercise are as follows:


1. Use the Nano Server Recovery Console to view basic settings.
2. Add Nano Server to the domain.
3. Use Windows PowerShell to configure the Nano Server settings.
4. Enable remote management with Server Manager.
5. Test the file server and web server on Nano Server.
6. Prepare for the next module.

► Task 1: Use the Nano Server Recovery Console to view basic settings

1. On **NANO-SVR1**, observe that the computer name is **Nano-Svr1** and that the computer is in a workgroup.
2. In **Network Adapter Settings**, notice that DHCP is obtaining the IP configuration. Make a note of the IP address: _____

► Task 2: Add Nano Server to the domain

1. On **LON-DC1**, in the **Administrator: Windows PowerShell** window, run **djoin.exe /provision /domain adatum /machine nano-svr1 /savefile .\odjblob**. This creates a file that you will use to complete the process of adding Nano Server to the domain.

 **Note:** Replace the IP address 172.16.0.X in the following commands with the IP address that you recorded earlier from your Nano Server installation.

2. The following commands are used to enable Windows PowerShell remoting:

```
Set-Item WSMan:\localhost\Client\TrustedHosts "172.16.0.X"
$ip = "172.16.0.X"
Enter-PSSession -ComputerName $ip -Credential $ip\Administrator
```

3. In the **Windows PowerShell credential request** dialog box, in the **Password** box, type **Pa\$\$w0rd**, and then click **OK**.
4. To enable file sharing through the firewall, run **netsh advfirewall firewall set rule group="File and Printer Sharing" new enable=yes**.
5. Close the Windows PowerShell remoting session by running **Exit-PSSession**.
6. Map a network drive to the C drive on Nano Server. (net use z: \\172.16.0.X\c\$)
7. Switch to the Z drive, and then copy **C:\odjblob** to the root of the C drive on Nano Server.
8. Reestablish a Windows PowerShell remoting session to Nano Server.
9. Run **djoin /requestodj /loadfile c:\odjblob /windowspath c:\windows /localos** to complete the process of adding the computer to the domain.
10. Run **shutdown /r /t 5** to restart Nano Server.

11. On **NANO-SVR1**, sign in as **Administrator** with the password of **Pa\$\$w0rd** in the **Adatum** domain.
12. In the Nano Server Recovery Console, observe that the computer is in the **adatum.com** domain.

► **Task 3: Use Windows PowerShell to configure the Nano Server settings**

1. On **LON-DC1**, close Windows PowerShell.
2. Open an elevated Windows PowerShell prompt.
3. Run **get-windowsfeature –comp Nano-svr1** to list the installed roles and features on Nano Server.
4. To add the File Server role to Nano Server, run **install-windowsfeature Fs-fileserver –comp Nano-svr1**. If you see error “Warning: Failed to start automatic updating for installed components. Error: 0x80040154” you can ignore this.
5. To verify that the role is installed, run **get-windowsfeature –comp Nano-svr1**.
6. Enable a Windows PowerShell remoting session with Nano Server. Remember to change X to the last octet of the IP address of your Nano server:
 - a. Run **\$ip = "172.16.0.X"**.
 - b. Run **Enter-PSSession -ComputerName \$ip -Credential \$ip\Administrator**.
7. When prompted, provide **Pa\$\$w0rd** as the password.
8. To view the IP configuration of Nano Server, run **get-netipaddress**.
9. To view the startup environment of Nano Server, run **bcdedit /enum**.
10. To view the shared folders, run **net share**. Only default shares exist.
11. At the command prompt, type the following cmdlet, and then press Enter.

```
Exit-PSSession
```

► **Task 4: Enable remote management with Server Manager**

1. On **LON-DC1**, in **Server Manager**, add **NANO-SVR1** to the **Computer** list.
2. In **Server Manager**, expand **File and Storage Services**, click **Shares**, and then in the **TASKS** list, click **New Share**.
3. Create a new shared folder:
 - Type: **SMB Share - Quick**
 - Server: **nano-svr1**
 - Share name: **Data**

► **Task 5: Test the file server and web server on Nano Server**

1. If necessary, on **LON-DC1**, map drive **Z** to **\\Nano-svr1\c\$**.
2. Start **Notepad**, and then create a file with the following line.

```
<H1> Nano Server Website </H1>
```

3. Save the file called **Default.htm** to **z:\inetpub\wwwroot**.
4. Open **Windows Internet Explorer**, and then navigate to **http://nano-svr1**. Does your webpage display?
5. Map drive **Y** to **\\Nano-svr1\data**.

6. Open **WordPad**, create a file, and then save the file to the root of drive **Y**.
7. Use File Explorer to verify that your file is saved on **Nano-Svr1**.

► **Task 6: Prepare for the next module**

When you have finished the lab, revert the virtual machines to their initial state.

1. On the host computer, start **Hyper-V Manager**.
2. In the **Virtual Machines** list, right-click **20743A-NANO-SVR1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20743A-LON-DC1**.

Results: After completing this exercise, you will have successfully configured the domain and network settings of Nano Server and installed an additional role.

Question: In the lab, you used a virtual machine to run Nano Server. Having created your virtual hard disk, if you want to run Nano Server on a physical host, what commands do you use to configure the startup environment?

Lesson 4

Preparing for upgrades and migrations

One of the key tasks when deploying Windows Server 2016 is identifying when you should upgrade an existing Windows Server deployment by using the existing hardware, or when you should migrate the roles and features to a clean installation of Windows Server 2016 on new hardware.

You should also use available guidance documentation and tools to determine which options are most suitable, and then use tools to automate the process. This lesson describes the considerations for performing an in-place upgrade or migrating to a new server. It also provides scenarios you can compare to your current business requirements and explains the benefits of migrating to a clean installation of Windows Server 2016. The lesson also provides you information about tools and guidance you can use to assess your own environment and help you deploy Windows Server 2016.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the differences between an in-place upgrade and migration.
- Identify scenarios when an in-place upgrade is appropriate.
- Identify the benefits of migrating to Windows Server 2016.
- List solution accelerators available to help with your upgrade or migration.
- Describe best practices for server consolidation.

In-place upgrades vs. server migration

When deploying Windows Server 2016, organizations must make the following choice:


- Use existing hardware and upgrade from supported editions of Windows Server 2008 or later.
- Install Windows Server 2016 on new hardware, and, if required, migrate the roles, features, and settings from servers that are running supported earlier Windows Server editions.

- Upgrading to Windows Server 2016:
 - Can upgrade from Windows Server 2008 R2 or later
 - Can only upgrade to same or newer editions
 - Requires same processor architecture
- Migrating to Windows Server 2016:
 - Must migrate from x86 version of Windows Server
 - Can use the Windows Server Migration Tools feature

When planning whether to upgrade or migrate a server to Windows Server 2016, consider the options in the following table.

Installation option	Description
Upgrade	<p>An upgrade preserves the files, settings, and applications that are installed on the original server. You perform an upgrade when you want to keep all these items and want to continue using the same server hardware. An upgrade requires x64 processor architecture and an x64 edition of the Windows Server operating system.</p> <p>If you are upgrading from Windows Server 2008 R2, you must install Service Pack 1 (SP1).</p> <p>You start an upgrade by running the Windows Server 2016 Setup Wizard from the original Windows Server operating system.</p>

Installation option	Description																				
	<p>You can perform the following upgrades to Windows Server 2016:</p> <table> <tr> <th>Original operating system and edition</th><th>Upgrade edition</th></tr> <tr> <td>Windows Server 2008 R2 Standard or Windows Server 2008 R2 Enterprise</td><td>Windows Server 2016 Standard or Windows Server 2016 Datacenter</td></tr> <tr> <td>Windows Server 2008 R2 Datacenter</td><td>Windows Server 2016 Datacenter</td></tr> <tr> <td>Windows Web Server 2008 R2</td><td>Windows Server 2016 Standard</td></tr> <tr> <td>Windows Server 2008 R2 Datacenter with SP1</td><td>Windows Server 2016 Datacenter</td></tr> <tr> <td>Windows Server 2008 R2 Enterprise with SP1</td><td>Windows Server 2016 Standard or Windows Server 2016 Datacenter</td></tr> <tr> <td>Windows Server 2008 R2 Standard with SP1</td><td>Windows Server 2016 Standard or Windows Server 2016 Datacenter</td></tr> <tr> <td>Windows Web Server 2008 R2 with SP1</td><td>Windows Server 2016 Standard</td></tr> <tr> <td>Windows Server 2012 Datacenter or Windows Server 2012 R2 Datacenter</td><td>Windows Server 2016 Datacenter</td></tr> <tr> <td>Windows Server 2012 Standard or Windows Server 2012 R2 Standard</td><td>Windows Server 2016 Standard or Windows Server 2016 Datacenter</td></tr> </table>	Original operating system and edition	Upgrade edition	Windows Server 2008 R2 Standard or Windows Server 2008 R2 Enterprise	Windows Server 2016 Standard or Windows Server 2016 Datacenter	Windows Server 2008 R2 Datacenter	Windows Server 2016 Datacenter	Windows Web Server 2008 R2	Windows Server 2016 Standard	Windows Server 2008 R2 Datacenter with SP1	Windows Server 2016 Datacenter	Windows Server 2008 R2 Enterprise with SP1	Windows Server 2016 Standard or Windows Server 2016 Datacenter	Windows Server 2008 R2 Standard with SP1	Windows Server 2016 Standard or Windows Server 2016 Datacenter	Windows Web Server 2008 R2 with SP1	Windows Server 2016 Standard	Windows Server 2012 Datacenter or Windows Server 2012 R2 Datacenter	Windows Server 2016 Datacenter	Windows Server 2012 Standard or Windows Server 2012 R2 Standard	Windows Server 2016 Standard or Windows Server 2016 Datacenter
Original operating system and edition	Upgrade edition																				
Windows Server 2008 R2 Standard or Windows Server 2008 R2 Enterprise	Windows Server 2016 Standard or Windows Server 2016 Datacenter																				
Windows Server 2008 R2 Datacenter	Windows Server 2016 Datacenter																				
Windows Web Server 2008 R2	Windows Server 2016 Standard																				
Windows Server 2008 R2 Datacenter with SP1	Windows Server 2016 Datacenter																				
Windows Server 2008 R2 Enterprise with SP1	Windows Server 2016 Standard or Windows Server 2016 Datacenter																				
Windows Server 2008 R2 Standard with SP1	Windows Server 2016 Standard or Windows Server 2016 Datacenter																				
Windows Web Server 2008 R2 with SP1	Windows Server 2016 Standard																				
Windows Server 2012 Datacenter or Windows Server 2012 R2 Datacenter	Windows Server 2016 Datacenter																				
Windows Server 2012 Standard or Windows Server 2012 R2 Standard	Windows Server 2016 Standard or Windows Server 2016 Datacenter																				
Migration	<p>Use migration when you migrate from an x86 edition of Windows Server 2003, Windows Server 2003 R2, or Windows Server 2008. You can use the Windows Server Migration Tools feature in Windows Server 2016 to transfer files and settings from computers that are running the following editions:</p> <ul style="list-style-type: none"> • Windows Server 2003 • Windows Server 2003 R2 • Windows Server 2008 • Windows Server 2008 R2 • Windows Server 2012 • Windows Server 2012 R2 																				

 **Additional Reading:** For more information on migration, refer to: "Install, Use, and Remove Windows Server Migration Tools" at: <http://aka.ms/Drjq4b>

In-place upgrade scenarios

An in-place upgrade involves upgrading a Windows Server operating system on the server that is running an earlier Windows Server edition. One benefit of an in-place upgrade is that you avoid hardware expenses because you install Windows Server 2016 on the existing hardware. Another benefit is that files, settings, and programs are kept intact on the server. Following are scenarios in which you would choose an in-place upgrade of the Windows Server operating system:

Perform an in-place upgrade when:

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

- When the hardware configuration of the existing servers meets the requirements for Windows Server 2016. Because the hardware requirements for Windows Server 2016 do not differ significantly from those for Windows Server 2012 R2, you likely can perform an in-place upgrade on those servers.
- When the software products that run on the existing servers support in-place upgrade of Windows Server 2016. Before performing an in-place upgrade, you must list all of the software products that are running on the server, such as Microsoft SQL Server, Exchange Server, non-Microsoft software, and antivirus software. Next, verify that these products support an in-place upgrade of Windows Server 2016. If so, refer to the specific product's documentation to determine how to perform an in-place upgrade, including any issues or risks that might occur.
- When you want to keep all user data that is on the existing servers, such as data stored on file servers, and the security permissions for accessing those data. When performing an in-place upgrade, user data and security permissions for accessing the data remain unchanged. This scenario is convenient, because after the in-place upgrade, users can continue to access their data on the same file servers.
- When you want to install Windows Server 2016, but you want to keep all roles, features, and settings of the existing server. Before performing an in-place upgrade on a server that has specific roles, features, or settings—such as DHCP, DNS, or AD DS—list those configurations. Then, check if those configurations support an in-place upgrade of Windows Server 2016. If so, refer to the detailed instructions for the specific roles, features, or settings on how to perform the in-place upgrade, including any issues or risks that might occur.

If any of these scenarios do not meet your organization's requirements, then you should perform a migration to Windows Server 2016.

Benefits of migrating to Windows Server 2016

When deploying Windows Server 2016, some organizations should consider migration instead of an in-place upgrade. There can be risks that arise from an in-place upgrade, such as server unavailability or data being inaccessible. Therefore, your organization might choose to perform a migration because of the following benefits:

- You will deploy servers with the Windows Server 2016 operating system installed, and they will not affect the current IT infrastructure. After you install Windows Server 2016, you can perform tests, such as drivers or system performance tests, before you introduce that server to the domain. In this way, the process of installation and testing is less likely to affect your current IT infrastructure.
- You will perform software product migration in a separate environment. For any software solution with an earlier Windows Server edition, you must refer to the product documentation for information about how to migrate that solution to Windows Server 2016. In some scenarios, software products that you are using are not supported for installation on Windows Server 2016, and you will require newer editions of those software products. In this case, by using migration, you can perform systematic installation of the operating system and the software products, in a separate environment. This ensures that the migration does not affect the availability of current services that the software provides.
- You will perform migration of server roles, features, and settings in a separate environment. As with the migration of software products, refer to the documentation on how to migrate the specific roles, features, or settings, such as DHCP, DNS, or AD DS, to Windows Server 2016. Again, migration enables you to perform systematic configuration in a separate environment, which means that the migration should not affect availability of server roles, features, and settings.
- New operating system enhancements are installed by default. When performing an in-place upgrade, for compatibility reasons, Windows Server 2016 is configured with settings for Windows Server 2008 or Windows Server 2008 R2. This means that many enhancements that Windows Server 2016 introduces, such as security, functionality, or performance enhancements, are not enabled by default. When performing migration, Windows Server 2016 deploys as a clean installation with all new enhancements installed. This ensures that the operating system is more secure and has new functionalities installed by default.

When you perform a migration, you:

- Do not affect your current Windows Server 2008 or later IT infrastructure
- Perform software product migration in a separate environment
- Perform migration of server roles, features, and settings in a separate environment
- Ensure that new operating system enhancements are installed by default

Using solution accelerators

Organizations should consider using software tools to help them plan their upgrade and migration to Windows Server 2016. Microsoft provides guidance to help you design and plan your Windows Server 2016 deployment and solution accelerators to assist in the process.

Microsoft Deployment Toolkit

The Microsoft Deployment Toolkit (MDT) is both a process and a lightweight tool for automated server (and desktop) deployments. It is used for deploying standardized images. MDT is based on a variety of Microsoft technologies including preboot execution environment (PXE), Windows Deployment Services (WDS), and System Center Configuration Manager (SCCM). MDT automates the deployment process by configuring unattended Setup files and packaging the files into an image file that you can deploy to a target computer.

- Use MDT to:
 - Automate deployments of Windows Server 2016 or other Windows operating systems
- Use MAP Toolkit for Windows Server 2016 to:
 - Perform inventory of your organization's IT infrastructure
 - Generate a report or proposal based on the Windows Server 2016 Readiness Assessment to plan server consolidation
- Use Windows Server Migration Tools to migrate server roles, features, operating system settings, data, and shares



Additional Reading: For more information about using the MDT as part of a complete deployment solution, refer to: "Automate and manage Windows operating system deployments" at: <http://aka.ms/Ofwaxa>


For more information about the MDT, including the latest updates, refer to: "Microsoft Deployment Toolkit" at: <http://aka.ms/Gqaxp8>

Microsoft Assessment and Planning Toolkit (MAP)

The **Microsoft Assessment and Planning Toolkit** (MAP) is a solution accelerator that analyzes the inventory of an organization's server infrastructure, performs an assessment, and then creates reports that you can use for upgrade and migration plans. MAP is available for Windows Server 2016, Windows Server 2012 R2, Windows 10 and Windows 8.1, and for other products, such as SQL Server 2014 and Hyper-V.


You can use the MAP to perform the following tasks:

- Inventory your organization's IT infrastructure. Based on the inventory, MAP displays a detailed report about which machines are capable of running Windows Server 2016, which machines are capable of running Windows Server 2016 with minimum system requirements, and which machines are not capable of running Windows Server 2016. MAP also recommends specific upgrades that ensure that computers are capable of running Windows Server 2016.
- Generate a report or proposal based on the Windows Server 2016 Readiness Assessment. The report or proposal is a document that contains an Executive Overview, Assessment Results, Next Steps, and a worksheet summarizing Windows Server 2016 readiness for computers that are running Windows Server.
- Capture the performance metrics of the current IT infrastructure, to help plan consolidation and server virtualization. The performance assessment generates reports on performance and presents the server consolidation recommendations.
- Estimate server utilization based on that metric before and after the virtualization. You also can choose which current physical servers are the best candidates for virtualization, and the hosts on which you should place those virtual machines.

 **Reference Links:** For more information about the Microsoft Assessment and Planning (MAP) Toolkit, refer to: <http://aka.ms/vjfbdj>

Win Server migration tools

Windows Server 2016 includes tools to assist you in migrating server roles and features from one computer to another. These Windows PowerShell cmdlets are part of a snap-in that is installed as part of a full installation or Core Server installation. Microsoft also provides detailed migration guides for specific roles.

 **Additional Reading:** For more information about the Windows Server Migration Tools and migration guides for specific roles and features, refer to: "Migrate Roles and Features to Windows Server" at: <http://aka.ms/H31ibv>

Recommendations for server consolidation

When deploying Windows Server 2016, you should plan the placements of server roles, such as AD DS, DNS, and DHCP, to make the best use of hardware and network resources. Organizations should consider cohosting multiple roles, where possible, to achieve the most economical solution. Virtualization is also considered as a consolidation of the server roles. Nano Server is particularly helpful in consolidating multiple server roles to a single machine. However, you should not implement cohosting if it affects server performance or available disk space. Therefore, organizations should evaluate and test whether installing multiple server roles on a server would result in lower overall performance and disk usage. Furthermore, organizations should evaluate the possible security risks of colocating server roles. For example, the server that hosts the root Active Directory Certificate Services role should not be colocated with other server roles and should be offline most of the time.

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

Small organizations should consider the following best practices:

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

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

- If you are cohosting multiple roles on a single server, there might be performance issues because of the large number of client computers that are connecting to that server. In this situation, organizations should consider adding multiple servers that cohost the same multiple roles. They also should consider relocating some of the roles from the first server to the other physical servers.

- High availability configurations of roles have specific requirements and settings, which might not support cohosting of multiple roles. In this situation, organizations could have a high availability solution for one server role, and then locate remaining roles on other servers.

Demonstration: Using MAP

In this demonstration you will see how to:

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

Demonstration Steps

Review the MAP options

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

Perform inventory

1. On **LON-CL1**, in the **Microsoft Assessment and Planning Toolkit** console, in the left pane, select **Overview**, and then in the **Overview** page, create an inventory database named **INVENTORY**.
2. On the **Overview** page, select **Perform an inventory**.
3. In the **Inventory and Assessment Wizard** window, perform the following steps:
 - a. On the **Inventory Scenarios** page, select the following check boxes:
 - **Windows computers**
 - **Exchange Server**
 - **Lync Server**
 - **SQL Server**
 - **Windows Azure Platform Migration**
 - b. On the **Discovery Methods** page, select **Use Active Directory Domain Services**, and **Scan an IP address range**.
 - c. On the **Active Directory Credentials** page, in the **Domain** field, enter **Adatum.com**. In the **Domain Account** field, enter **Adatum\Administrator**, and then in the **Password** field, type **Pa\$\$w0rd**, and on the next two pages accept the default settings.
 - d. On the **Scan an IP Address Range** page, enter the range from **172.16.0.1** to **172.16.0.100**.

- e. On the **All Computers Credentials** page, accept the default settings.
- f. On the **Summary** page, review the inventory options, and then cancel the wizard.



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

Review the MAP inventory from a sample database

1. In the **Microsoft Assessment and Planning Toolkit** console, on the **File** menu, select **Manage Databases**.
2. In the **Microsoft Assessment and Planning Toolkit** dialog box, import the sample database using the following steps:
 - a. Select **Manage**.
 - b. Import the sample database located in following path: In the **File name** field, type **C:\Program Files\ Microsoft Assessment and Planning Toolkit\Sample \MAP_SampleDB.bak**.
 - c. In the **Database Name** field, type **MAPDEMO**.
 - d. In the **Microsoft Assessment and Planning Toolkit** window, choose the **Use an existing database** option, and then select **MAPDEMO** database.
3. In the **Microsoft Assessment and Planning Toolkit** console, review the default window that displays the **Overview** page that includes inventory information from the sample database. Refresh the window in the **Overview** page, if necessary.
4. In the **Microsoft Assessment and Planning Toolkit** console, in the left pane, click **Cloud**, and then review the readiness information on the different cloud scenarios that displays with inventory information from the sample database.
5. In the **Microsoft Assessment and Planning Toolkit** console, in the left pane, click **Desktop**, and then review the readiness information on the different desktop scenarios that appear with inventory information from the sample database.
6. Repeat step 4 for all remaining items in the left pane: **Server**, **Desktop Virtualization**, **Server Virtualization**, **Database**, **Usage Tracking**, and **Environment**.

Question: How does virtualization help in server role consolidation?

Lesson 5

Migrating server roles and workloads

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

You must also determine which roles you can migrate, which you can cohost, and which you can consolidate into a virtual environment. Finally, you must plan for migrating roles within the same domain or across domains.

Lesson Objectives

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

- Explain how to implement server migrations.
- Explain how to migrate servers across domains.

Migrating server roles within a domain

When planning to migrate servers, you must create a list of the server roles that you want to migrate and the steps that each migration involves. For each server role that you plan to migrate, you should refer to the technical documentation and migration guides about how to perform the migration. When you perform the migration, you should use the Windows Server Migration Tools, which are available with Windows Server 2016.

The roles that you can migrate include:

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

The roles that you can migrate from supported earlier editions of Windows Server to Windows Server 2016 include:

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

Installing and preparing the Windows Server Migration Tools consists of the following steps:

1. Install the tools on destination servers as part of Windows Server 2016 setup.
2. Create a deployment folder containing a copy of the tools on the destination server.
3. Copy the deployment folder from destination server to source server.
4. Register Windows Server Migration Tools on the source server by using the **SmigDeploy.exe** tool included in the deployment folder.

After the migration tools are installed, you can run them by using one of the following methods:

- Run **Windows Server Migration Tools** as an administrator from the Windows Start screen.
- Load the Windows Server Migration Tools snap-in into an elevated Windows PowerShell session.
- On source computers running earlier versions of Windows Server, run **Windows Server Migration Tools** under the **Windows PowerShell** folder, which is under the **All Programs** folder of the **Start** menu.



Note: You can migrate roles only from supported earlier Windows Server editions to Windows Server 2016.



Additional Reading: For more information about determining which roles and features to migrate, refer to the migration guides for Windows Server 2016 in "Migrate Roles and Features to Windows Server" at: <https://technet.microsoft.com/en-us/library/jj134039>

Migrating server roles across domains or forests

Organizations could choose to deploy Windows Server 2016 in a new AD DS forest. In this scenario, administrators should plan the migration steps carefully to provide users with seamless access to data and services during the migration process. After the migration is complete, administrators should begin the process of decommissioning and removing the infrastructure of the previous operating system environment.

To migrate a server across domains:

- Create a new Windows Server 2016 AD DS forest that is independent from the forest that is running a previous operating system version.
- Deploy new servers that are running the Windows Server 2016 operating system.
- Deploy Microsoft applications, such as Exchange Server, SQL Server, and Microsoft SharePoint Server in the new AD DS forest.
- Deploy corporate custom applications or third-party applications in the new AD DS forest that the previous infrastructure environment used.
- Configure DNS infrastructure in both forests.
- Establish AD DS trust between both current and new AD DS forests.

When migrating servers across domains:

- Create a new Windows Server 2016 AD DS forest
- Deploy applications on new servers
- Establish AD DS trust between the current and the new AD DS forests
- Migrate AD DS objects
- Migrate application data and settings
- Decommission and remove the old AD DS environment

- Migrate AD DS objects, such as users, computers, groups, and mailboxes.
- Migrate application data and settings for Microsoft applications, corporate custom applications, and third-party applications.
- Ensure that users can connect to corporate IT resources in the new AD DS forest.
- Decommission and remove the environment, based on previous operating system's AD DS forest.



Note: For each product and application that you plan to migrate to a Windows Server 2016 AD DS forest, read the product documentation and best practices, including the supported migration procedures. You can find more information on the individual products' websites.



Note: You must use a tool, such as the Active Directory Migration Tool (ADMT), to migrate resources such as users, computers, and groups across forests or within the same forest. For more information about using ADMT, refer to: "ADMT Guide: Migrating and Restructuring Active Directory Domains" at: <http://aka.ms/Xi5xyk>

Question: What are some reasons that you would do a cross-forest migration instead of a migration within the same domain?

Lesson 6

Windows Server activation models

As part of planning your server upgrade and migration process, you should also consider how you will manage operating system licensing and activation. Your choice of activation model will be based on the characteristics of your environment.

Lesson Objectives

After this lesson, you will be able to:

- Describe the volume licensing and activation options for Windows Server 2016.
- Plan a suitable volume activation process.

Windows Server 2016 licensing and activation

To ensure that your organization has the proper licenses, and to receive notices for product updates, you must activate every copy of Windows Server 2016 that you install. Windows Server 2016 requires that you activate the operating system after installation. This verifies that the products are licensed and that you receive important update information. There is no activation grace period. If you do not activate Windows Server 2016, you cannot customize your operating system. There are two general activation strategies:

Organizations can choose between two activation strategies:

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

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

Manual activation

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

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

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

Automatic activation

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

There are several technologies available that help automate the process of activating Windows Server licenses:

- **Key Management Services (KMS).** KMS is a service that helps you activate licenses on systems within your network from a server where a KMS host has been installed. The KMS host completes the activation process instead of individual computers connecting to a Microsoft online service to complete activation.
- **Volume Activation Services server role.** This server role helps you to automate issuing and managing Microsoft software volume licenses. Volume Activation Services allows you to install and configure KMS and Active Directory–based Activation. KMS requires activating at least five servers and 25 clients. KMS is the default key for volume activation.
- **Active Directory–Based Activation.** This is a service that lets you use AD DS to store activation objects. A computer running Windows Server (or client) automatically contacts AD DS to receive an activation object, without the need to contact Microsoft. You can use Active Directory–based activation when activating servers and clients running Windows Server 2012 or later, and Windows 8 or later. Your AD DS schema must also be Windows Server 2012 or newer.
- **Volume Activation Tools console.** The Volume Activation Tools console is used to install, activate, and manage volume license activation keys in AD DS or KMS.
- **Volume Activation Management Tool (VAMT).** The VAMT is a free tool that you can use to manage volume activation using Multiple Activation Keys (MAKs) or to manage KMS. You can use VAMT to generate license reports and manage client and server activation on enterprise networks.
- **Multiple Activation Key (MAK).** A MAK is a volume license key that you can use for independent activation by connecting with Microsoft or through proxy activation, where a single computer gathers the activation information for multiple computers and contacts Microsoft for them. Use MAK when your systems have poor—or no—connection with your organization’s central network.
- **Automatic Virtual Machine Activation.** This activation option lets you install VMs on a virtualization server with no product key.



Reference Links: For more information on VAMT, refer to: “Introduction to VAMT” at: <http://aka.ms/I0cdnd>

Licensing changes since Windows Server 2008

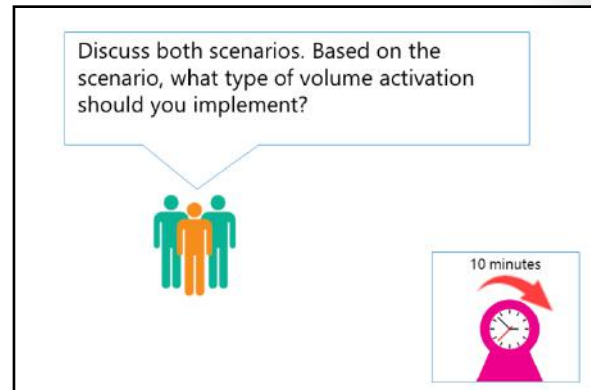
As part of planning your deployment, you must ensure that you have the proper number of licenses for your Windows Server 2016 installation. Windows Server 2016, like Windows Server 2012, is licensed by processor core, not by server. You can purchase additional licenses for two processor cores at a time.

Discussion: Planning volume activation

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

Question: Your organization's IT infrastructure consists of personal computers and servers that are running different editions of Windows client operating systems and Windows Server operating systems. Next month, your organization plans to deploy 500 Windows 10 client computers and 20 Windows Server 2016 servers. Because of a legacy application in the finance department, you must deploy 10 client computers that are running Windows 8.1 and two servers that are running Windows Server 2012 R2. What type of volume activation should you implement?

Question: Your organization's IT infrastructure was upgraded from different editions of Windows client operating systems and Windows Server operating systems to Windows 10 and Windows Server 2016, respectively. What type of volume activation should you implement?



Module Review and Takeaways

Review Questions

Question: When creating a virtual hard disk for Nano Server by using the Windows PowerShell cmdlet **New-NanoServerImage**, when do you use the **-Guestdrivers** switch?

Question: When using the Nano Server Recovery Console, which two fundamental components can you configure?

Question: Which role can you use to manage KMS?

Module 2

Overview of storage in Windows Server 2016

Contents:

Module Overview	2-1
Lesson 1: Overview of storage in Windows Server 2016	2-2
Lesson 2: Implementing Data Deduplication	2-11
Lesson 3: Configuring iSCSI storage	2-30
Lab A: Implementing and managing storage	2-38
Lesson 4: Configuring the Storage Spaces feature in Windows Server 2016	2-45
Lab B: Implementing and managing advanced storage solutions	2-55
Module Review and Takeaways	2-63

Module Overview

The requirements for storage space have been increasing since the invention of server-based file shares. The Windows 10 and Windows Server 2016 operating systems include several new features designed to reduce the required disk space, and to manage physical disks effectively, including Data Deduplication and Storage Spaces. This module provides an overview of these features, and explains the steps you must perform to configure them. Another emerging concern with respect to storage is the connection between storage and remote disks. Internet small computer system interface (iSCSI) storage in Windows Server 2016 is a cost-effective feature that helps create a connection between the servers and the storage. To implement iSCSI storage in Windows Server 2016, you must be familiar with the iSCSI architecture and components. Furthermore, you must be familiar with the tools that Windows Server provides to implement iSCSI-based storage.

Managing physical disks attached directly to a server can be a tedious task for administrators. Therefore, many organizations implement storage area networks (SANs) to address this issue and use storage more efficiently. However, SANs can be expensive, particularly for small businesses. You can use the Storage Spaces feature in Windows Server 2016 to provide similar functionality to hardware-based storage solutions. The Storage Spaces feature pools disks together and presents them to the operating system as a single disk.

This module explains how to configure and implement Storage Spaces. This module also will highlight new storage features in Windows Server 2016, including Storage Spaces Direct.

Objectives

After completing this module, you will be able to:

- Describe the new features in Windows Server 2016 storage.
- Implement Data Deduplication.
- Configure iSCSI storage.
- Configure the Storage Spaces feature.

Lesson 1

Overview of storage in Windows Server 2016

The storage demand on servers is ever increasing, and storage constitutes a large portion of an IT department's budget. Organizations require large volumes on flexible disks that you can add or remove dynamically to ensure business continuity. Windows Server 2016 includes enhanced storage features that provide improved management of physical disks, and technologies that help reduce disk-space consumption.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the new Windows Server 2016 storage features.
- Describe Software Defined Storage.
- Explain Resilient File System (ReFS).
- Describe the features of File Server Resource Manager (FSRM).
- Use FSRM to manage quotas, file screens, and storage reports.
- Explain file classification.
- Configure file classification.

New storage features in Windows Server 2016

File and Storage Services includes technologies that help you set up and manage one or multiple file servers.

New features in Windows Server 2016

The following file and storage services features are new or improved in Windows Server 2016:

- **Storage Spaces Direct.** This feature enables you to build highly available storage systems by using storage nodes with only local storage. You will learn more about this feature later in this module.
- **Storage Replica.** It is a new feature that enables replication—between servers or clusters that are in the same location or different sites—for disaster recovery. Storage Replica includes both synchronous and asynchronous replication for shorter or longer distance between sites. This enables you to achieve storage replication at a lower cost. For more information about Storage Replica, refer to Module 11: “Implementing failover clustering.”
- **Storage Quality of Service (QoS).** With this feature, you can create centralized QoS policies on a Scale-Out File Server and assign them to virtual disks on Hyper-V virtual machines. QoS ensures that performance for the storage adapts to meet policies as the storage load changes.

Windows Server 2016 provides the following new file and storage services features:

- Storage Spaces Direct
- Storage Replica
- Storage QoS
- Data Deduplication (improved)



- Data Deduplication. This feature was introduced in Windows Server 2012 and is improved in Windows Server 2016 in the following areas:
 - Integrated support for virtualized backup workloads.
 - Support for cluster rolling upgrades.
 - Support for Nano Server.
 - Improved optimization throughput for large volumes.
 - Improved deduplication performance of very large files.

New features in Windows Server 2012 and Windows Server 2012 R2

Windows Server 2012 R2 and Windows Server 2012 offer several new and improved file and storage-services features, including:

- Multiterabyte volumes. This feature deploys multiterabyte NTFS file system volumes, which support consolidation scenarios and maximize storage use. NTFS volumes on master boot record (MBR) formatted disks can be up to 2 terabytes (TB) in size. Volumes on a globally unique identifier (GUID) partition table (GPT) formatted disks can be up to 18 exabytes.
- Data Deduplication. This feature saves disk space by storing a single copy of identical data on the volume.
- iSCSI Target Server. The iSCSI Target Server provides block storage to other servers and applications on the network by using the iSCSI standard. Windows Server 2012 R2 includes also VHDX support and end-to-end management via the Storage Management Initiative Specification.
- Storage spaces and storage pools. This feature enables you to virtualize storage by grouping industry standard disks into storage pools, and then create storage spaces from the available capacity in the storage pools. Storage spaces in Windows Server 2012 R2 enables you to create a tiered storage solution that transparently delivers an appropriate balance between capacity and performance that can meet the needs of enterprise workloads.
- Unified remote management of File and Storage Services in **Server Manager**. You can use the **Server Manager** to manage multiple file servers remotely, including their role services and storage.
- Windows PowerShell cmdlets for File and Storage Services. You can use the Windows PowerShell cmdlets for performing most administration tasks for file and storage servers.
- ReFS. The new Resilient File System (ReFS) introduced in Windows Server 2012 offers enhanced integrity, availability, scalability, and error protection for file-based data storage.
- Server Message Block (SMB) 3.0. SMB protocol is a network file-sharing protocol that allows applications to read and write to files and request services from server programs on a network.
- Offloaded Data Transfer (ODX). ODX functionality enables ODX-capable storage arrays to bypass the host computer and directly transfer data within or between compatible storage devices.
- Chkdsk. The new version of Chkdsk runs automatically in the background and monitors the health of the system volume; enabling organizations to deploy multiterabyte NTFS file system volumes without concern about endangering their availability. The Chkdsk tool introduces a new approach. It prioritizes volume availability and allows for the detection of corruption while the volume remains online and its data remains available to the user during maintenance.

What is Software Defined Storage?

The Software Defined Storage is a concept of managing storage independent of the underlying hardware. There, currently, is no defined standard to describe Software Defined Storage, and different vendors might describe Software Defined Storage differently.

Software Defined Storage has the potential to revolutionize the storage area of IT in the same way that virtualization has revolutionized the way you work with servers. Software Defined Storage can lower the overall storage cost by using inexpensive storage arrays like just a bunch of disks (JBOD) and then *virtualize* the storage layer by using the Storage Spaces feature in Windows Server 2016 to view all disks as one virtual disk.

Software Defined Storage also can provide management of options such as deduplication, replication, thin provisioning, tiers, snapshots, and backup. Software Defined Storage features in Windows Server 2016 include:

- **Storage Spaces.** Storage Spaces enables you to virtualize storage by grouping hard disks into storage pools and then creating virtual disks, called storage spaces, from available capacity in the pools.
- **Scale-Out File Servers.** Scale-Out File Server clusters allow you to create continuously available file shares that can balance load between cluster nodes.
- **Storage Replica.** Storage Replica enables replication—between servers or clusters that are in the same location or different sites—for disaster recovery.
- **Data Deduplication.** The Data Deduplication feature optimizes the used storage by removing identical data blocks.
- **Thin provisioning.** You can create storage spaces that are thin provisioned, so you do not need to have the full storage capacity available when you create the storage space. You can add new disks to the storage space when these new disks are available.
- **iSCSI storage.** iSCSI targets can present storage to servers using the same protocol as a hardware-based iSCSI SAN.

Enterprise SANs can provide the same functionality as Software Defined Storage in Windows Server 2016. However, the solution that utilizes Windows Server 2016 often is available at a lower cost.

- Is a concept of managing storage independent of the underlying hardware
- Can revolutionize storage in the same way as virtualization has revolutionized servers
- Includes the following features in Windows Server 2016:

- Storage Spaces
- Scale-Out File Servers
- Storage Replica
- Data Deduplication
- Thin provisioning
- iSCSI storage

Overview of ReFS

ReFS is a file system that Windows Server 2012 introduced. ReFS improves on the NTFS file system and provides several advantages, including:

- Metadata integrity with checksums.
- Integrity streams that provide optional user-data integrity.
- Allocation on a write transactional model for robust disk updates, also known as copy on write.
- Large volume, file, and directory sizes.
- Data striping for performance, which means that you can manage bandwidth, and redundancy for fault tolerance.
- Disk scrubbing for protection against latent disk errors.
- Resiliency to corruptions with salvage for maximum volume availability.

ReFS:

- Is a file system introduced in Windows Server 2012
- Provides the following benefits:
 - Metadata integrity with checksums
 - Integrity streams providing optional user data integrity
 - Allocation on write transactional model
 - Large volume, file, and directory sizes
 - Data striping for performance and redundancy
 - Disk scrubbing for protection against latent disk errors
- Includes a new feature, Accelerated VHDX Operations, that improves performance on Hyper-V storage

ReFS inherits several features from NTFS, including support for BitLocker encryption, access control lists for security, update sequence number (USN) journal, change notifications, symbolic links, junction points, mount points, reparse points, volume snapshots, file IDs, and opportunistic locks.

ReFS uses a subset of NTFS features, and its design maintains backward compatibility with its earlier counterpart. Clients running Windows 10 or earlier operating systems can access file shares stored on ReFS hard-drive partitions, just as they can with those that are stored on NTFS. However, as its name implies, ReFS offers more resiliency. This means it provides better data verification, error correction, and scalability. Clients that are running Windows 10, Windows 8.1, and Windows 8 also can access ReFS drives attached to the local computer. ReFS drives on clients are only available on mirrored storage spaces.

If you are using Windows Server 2016, we recommend using the ReFS file system for storing virtual machines in Hyper-V. Accelerated VHDX Operations is a new feature that improves performance when creating or extending virtual hard disks, merging checkpoints, and creating backups.

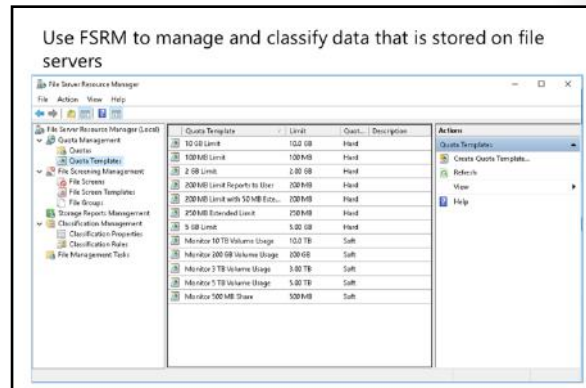
Beyond its greater resiliency, ReFS also surpasses NTFS by offering larger maximum sizes for individual files, directories, disk volumes, and other items, which the following table lists.

Attribute	Limit
Maximum size of a single file	2^{64-1} bytes (18,446,744,073,709,551,616 bytes)
Maximum size of a single volume	2^{78} bytes with 16 kilobytes (KB) cluster size ($2^{64} * 16 * 2^{10}$) Windows stack addressing allows 264 bytes
Maximum number of files in a directory	2^{64}
Maximum number of directories in a volume	2^{64}
Maximum file name length	32K Unicode characters
Maximum path length	32K

Overview of File Server Resource Manager

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

- **Quota management.** You can use this feature to limit the space allowed for a volume or folder. Quotas can apply automatically to new folders that you create on a volume. You also can define quota templates that you can apply to new volumes or folders.
- **File screening management.** You can use this feature to control the types of files that users can store on a file server. You can limit the file types with specific file extensions that users can store on your file shares. For example, you can create a file screen that does not allow users to save files with an MP3 extension in a file server's personal shared folders. File screening only looks for the file extension. It does not examine the file contents.
- **Storage reports.** You can use this feature to identify trends in disk usage and how your data is classified, and monitor attempts by a selected group of users to save unauthorized files.
- **File Classification Infrastructure.** This feature automates the data classification process. You can dynamically apply access policies to files based on their classification. Example policies include Dynamic Access Control for restricting access to files, file encryption, and file expiration. You can classify files automatically by using file classification rules, or you can classify them manually by modifying the properties of a selected file or folder.
- **File management tasks.** You can use this feature to apply a conditional policy or action to files, based on the files' classification. The conditions of a file management task include the file location, the classification properties, the file creation date, the file modification date, or the file access date. The actions that a file management task can take include the ability to expire files, encrypt files, or run a custom command.
- **Access-denied assistance.** You use this feature to customize the access denied error message that users see in Windows 8 and newer clients when they do not have access to a file or a folder.



You can access the FSRM by using the **File Server Resource Manager Microsoft Management Console** (MMC) console or by using Windows PowerShell. You can view all the available cmdlets by running the following command at a Windows PowerShell command prompt:

```
Get-Command -Module FileServerResourceManager
```

Demonstration: Using FSRM to manage quotas, file screens, and storage reports

In this demonstration, you will see how to:

- Create a quota.
- Test a quota.
- Create a file screen.

- Test a file screen.
- Generate a storage report.

Demonstration Steps

Create a quota

1. Sign in to **LON-SVR1** as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. Open **Server Manager**.
3. Install the **File Server Resource Manager** role service.
4. Open the **File Server Resource Manager** console.
5. Create a quota based on the 100 megabyte (MB) limit on the **E:\Labfiles\Mod02** folder.

Test a quota

1. Open **Windows PowerShell**.
2. Attempt to create a new large file in the **E:\Labfiles\Mod02** folder by using the following command:

```
fsutil file createnew largefile.txt 123456789
```

3. Close Windows PowerShell.

Create a file screen

- In File Server Resource Manager, create a new file screen based on the **Block Image Files** file-screen template for **E:\Labfiles\Mod02**.

Test a file screen

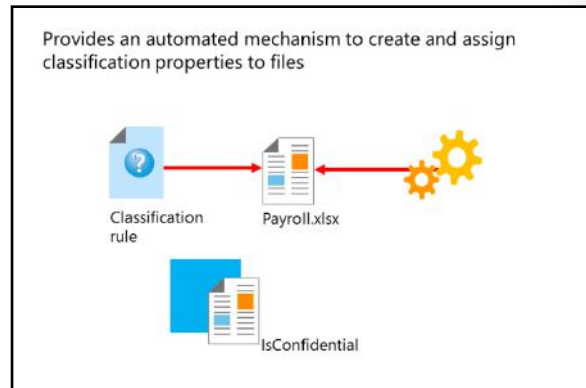
1. Open **File Explorer**.
2. Navigate to **E:\Labfiles**.
3. Create a new bitmap (.bmp) image named **testimage**.
4. Copy the **testimage**, and then paste it into the **E:\Labfiles\Mod02** folder.
5. View and cancel the error window.
6. Close the **File Explorer** window.

Generate a storage report

1. Execute the **E:\Labfiles\Mod02\CreateDemoFiles.cmd** batch file to create some demo files.
2. Generate an on-demand report for **Large Files** on **E:\Labfiles**.
3. View, examine, and then close the html report.
4. Close all open windows except Server Manager.

File classification

The File Classification Infrastructure in Windows Server 2016 allows administrators to classify files and create policies, and then apply those policies based on those classifications. This can lower costs and help you manage the security risks that are inherent with managing data. When you implement the File Classification Infrastructure feature, data management requirements do not affect the storage layout, and the organization can adapt more easily to a changing business and regulatory environment.



Classification Management in FSRM is a tool that can ease the burden of data management in your organization. When you use Classification Management, you can classify files in many different ways. In most scenarios, you perform classification manually. In Windows Server 2016, the File Classification Infrastructure feature allows organizations to automate the way in which they classify data. You then can assign file-management policies based on a file's classification, and you can enforce corporate requirements based on the data's business value. You also can modify the policies easily by using the built-in tools. You can use file classification to:

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

File classification rules

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

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

- Is the rule enabled? On the classification rule's **Properties** page, on the **Rule Settings** tab, you can select or clear the **Enabled** check box to enable or disable the classification rule.
- What is the scope of the rule? On the **Rule Settings** tab, the **Scope** option allows you to select a folder or folders to which the classification rule will apply. When the rule runs, it processes and attempts to classify all file-system objects within this location.
- What classification mechanism will the rule use? On the classification rule's **Properties** page, on the **Classification** tab, you must choose a classification method that the rule will use to assign the classification property. By default, there are three methods from which you can choose:
 - Folder classifier. This mechanism assigns properties to a file based on the file's folder path.
 - Content classifier. This mechanism searches for strings or regular expressions in files. This means that the content classifier classifies a file based on its textual contents, such as whether it contains a specific word, phrase, or numeric value.
 - Windows PowerShell classifier. This classifier uses Windows PowerShell scripts to evaluate conditions and assign values to classification properties.
- What property will the rule assign? The main function of the classification rule is to assign a property to a file object based on how the rule applies to it. On the **Classification** tab, you must specify a property and the specific value that the rule will assign to that property.

- What additional classification parameters will you use? The additional classification parameters form the core of the rule's logic. Clicking the **Advanced** button on the **Classification** tab will open the Additional Classification Parameters window. In this window, you can specify additional parameters—including strings or regular expressions—that, if found in the file system object, will cause the rule to apply itself. For example, this parameter could be the phrase *Social Security Number* or any number with the format 000-00-000. If the rule finds this parameter, the classification parameter will apply a YES value for a Confidential classification property to the file. You also can use this classification to perform some tasks on the file-system object, such as moving it to a secure location.

A classification parameter can be one of three types:

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



Additional Reading: For more information about the syntax for creating regular expression, refer to: "Regular Expression Syntax" at: <http://aka.ms/lf9d2u>

You can run classification rules in two ways: on demand or based on a schedule. Each time you run a classification, both methods use all rules that you have left in the *enabled* state.

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

Demonstration: Configuring file classification

In this demonstration, you will see how to:

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

Demonstration Steps

Create a classification property

- On **LON-SVR1**, open **File Server Resource Manager**, and create a new classification property named **Confidential**, with the **Yes/No** property type.

Create a classification rule

1. Create a new classification rule named **Confidential Payroll Documents**.
2. Configure the scope of the rule to apply to **E:\Labfiles\Mod02**.
3. Configure the rule to classify documents with a value of **Yes** for the **Confidential** classification property, if the file contains the string expression **payroll**.

Modify the classification schedule

1. Create a classification schedule that runs every **Sunday** at **8:30 AM**.
2. Using the **Classification Rule** node, manually run **Classification With All Rules Now**, and then view the report. Ensure that **Document1.txt** and **Document2.txt** are listed at the bottom of the report.
3. Close all open windows, except **Server Manager** on **LON-SVR1**.

Verify the correctness of the statement by placing a mark in the column to the right.

Statement	Answer
You can use Windows PowerShell scripts to change the file-classification properties.	

Lesson 2

Implementing Data Deduplication

Data Deduplication is a role service of Windows Server 2016. This service identifies and removes duplications within data without compromising data integrity. This achieves the goals of storing more data and using less physical disk space. This lesson explains how to implement Data Deduplication in Windows Server 2016 storage.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the Data Deduplication role service in Windows Server 2016.
- Identify the Data Deduplication components in Windows Server 2016.
- Explain how to deploy Data Deduplication.
- Implement Data Deduplication.
- Describe the common usage scenarios for Data Deduplication.
- Explain how to monitor and maintain Data Deduplication.
- Explain how to troubleshoot the adverse effects of Data Deduplication.
- Describe the backup and restore considerations with Data Deduplication.

What is Data Deduplication?

To manage the growth in data storage in the enterprise, organizations are consolidating servers and making capacity scaling and data optimization the key goals. Data Deduplication provides practical ways to achieve these goals, including:

- **Capacity optimization.** Data Deduplication stores more data in less physical space. It achieves greater storage efficiency than was possible by using features such as Single Instance Store (SIS) or NTFS compression. Data Deduplication uses sub-file variable-size chunking and compression, which deliver optimization ratios of 2:1 for general file servers and up to 20:1 for virtualization data.
- **Scalability and performance.** Data Deduplication is highly scalable, resource efficient, and nonintrusive. Although it can process up to 50 MB of data per second in Windows Server 2012 R2 and about 20 MB of data per second in Windows Server 2012, Windows Server 2016 is staged to perform significantly better through the advancements in the deduplication processing pipeline. In Windows Server 2016, Data Deduplication can run multiple threads in parallel by using multiple I/O queues on multiple volumes simultaneously without affecting other workloads on the server. You can maintain a low impact on the server workloads by throttling the CPU and reducing memory resources that are consumed: if the server is very busy, deduplication can stop completely. In addition, you have the flexibility to run Data Deduplication jobs at any time, set schedules for when Data Deduplication should run, and establish file selection policies.

Data Deduplication:

- Identifies and removes duplications within data without compromising the data's integrity or fidelity
- Stores more data in less space

When you enable Data Deduplication on a volume, a background task runs with low priority that:

1. Segments data into small, variable-sized chunks
2. Identifies duplicate chunks
3. Replaces redundant copies with a reference
4. Compresses chunks

- Reliability and data integrity. When Data Deduplication is applied to a volume on a server, the integrity of the data is maintained. Data Deduplication uses checksum results, consistency, and identity validation to ensure data integrity. Data Deduplication maintains redundancy for all metadata and the most frequently referenced data to ensure that the data is repaired, or at least recoverable, in the event of data corruption.
- Bandwidth efficiency with BranchCache. Through integration with BranchCache, client computers can reduce the amount of data transferred over the wide area network (WAN) to a branch office. The result is faster file download times and reduced bandwidth consumption.



Note: BranchCache is a feature that can reduce WAN utilization and enhance network application responsiveness when users access content in a central office from branch office locations. When you enable BranchCache, it retrieves a copy of the content from the web server or file server and caches it within the branch office. If another client in the branch requests the same content, then the client can download it directly from the local branch network without needing to retrieve the content by using the WAN.

- Optimization management with familiar tools. Data Deduplication has optimization functionality built into **Server Manager** and Windows PowerShell. Default settings provide savings immediately, or you can fine-tune the settings to achieve more gains. By using Windows PowerShell cmdlets, you can start an optimization job or schedule one to run in the future. You also can install the Data Deduplication feature and enable deduplication on selected volumes by using the **Unattend.xml** file that calls a Windows PowerShell script and uses Sysprep to deploy deduplication when a system first starts.

The Data Deduplication process finds and removes duplication within data without compromising its fidelity or integrity. More data is stored in less space by segmenting files into small variable-sized chunks (32–128 KB), identifying duplicate chunks, and maintaining a single copy of each chunk.

After deduplication, files are no longer stored as independent streams of data, and they are replaced with stubs that point to data blocks that are stored within a common chunk store. The chunk store is where the identical parts of the deduplicated file are stored. The chunk store consists of container files that grow to about 1 gigabyte (GB) in size and then Data Deduplication creates a new container file. Because these files share blocks, those blocks are only stored once, which reduces the disk space needed to store all files. During file access, Data Deduplication transparently assembles the correct blocks to serve the data without the application or the user having any knowledge of the on-disk transformation to the file. This enables you to apply deduplication to files without having to worry about any change in behavior to the applications or impact to users who are accessing those files. Data Deduplication works best in storage scenarios with large amounts of data that are not modified frequently.

Enhancements to the Data Deduplication role service

Windows Server 2016 includes several improvements to the Data Deduplication feature, including:

- Support for volume sizes up to 64 TB. Data Deduplication in Windows Server 2012 R2 does not perform well on volumes greater than 10 TB in size (or less for workloads with a high rate of data changes). However, in Windows Server 2016, the deduplication processing pipeline is now multithreaded and can utilize multiple CPUs per volume to increase optimization throughput rates on volume sizes up to 64 TB. The 64 TB limit is a limitation of Volume Shadow Copy Service (VSS), which Data Deduplication depends on.
- Support for file sizes up to 1 TB. In Windows Server 2012 R2, very large files are not good candidates for Data Deduplication. However, with the use of the new stream map structures and other improvements to increase the optimization throughput and access performance, deduplication in Windows Server 2016 performs well on files up to 1 TB.

- Simplified deduplication configuration for virtualized backup applications. Although Windows Server 2012 R2 supports deduplication for virtualized backup applications, it requires manual configuration of the deduplication settings. In Windows Server 2016, however, a predefined usage type option significantly simplifies the configuration of deduplication for virtualized backup applications.
- Support for Nano Server. Nano Server is a new deployment option in Windows Server 2016 that has a smaller system resource footprint, starts significantly faster, and requires fewer updates and restarts than the Server Core deployment option. Additionally, Nano Server fully supports Data Deduplication.
- Support for cluster rolling upgrades. Failover clusters running deduplication can include a mix of deduplication nodes running Windows Server 2012 R2 and deduplication nodes running Windows Server 2016. This major enhancement provides full data access to all of your deduplicated volumes during a cluster rolling upgrade. This means that you can gradually upgrade each deduplication node in an existing cluster to Windows Server 2016 without incurring downtime to upgrade all the nodes at once.



Note: Although both the Windows Server versions of deduplication can access the optimized data, the optimization jobs will only run on the Windows Server 2012 R2 deduplication nodes. Optimization jobs will not run on the Windows Server 2016 deduplication nodes until the cluster rolling upgrade is complete.

Effectively, Data Deduplication in Windows Server 2016 allows you to efficiently store, transfer, and back up fewer bits.

Data Deduplication components

The Data Deduplication role service consists of several components. These components include:

- A filter driver. This component monitors local or remote I/O and interacts with various jobs to handle the chunks of data on the file system. There is one filter driver for every volume.
- The Deduplication service. This component manages four job types:
 - Optimization. Consisting of multiple jobs, this job type performs both deduplication and compression of files according to the data deduplication policy for the volume. After initial optimization of a file, if the file is modified and meets the data deduplication policy threshold for optimization, then the file will be optimized again.
 - Garbage collection. Garbage collection jobs process deleted or modified data on the volume so that any data chunks no longer belonging to optimized files. The garbage collection jobs delete the no longer referenced chunks from the chunk store. This job processes previously deleted or logically overwritten optimized content to create usable volume free space. When you delete or modify an optimized file, the old data in the chunk store is not deleted right away. Although garbage collection runs weekly, you might consider running garbage collection only after large deletions have occurred to optimize the chunk store.

The Data Deduplication role service consists of several components:

- A filter driver, which monitors local or remote I/O
- The Deduplication service, which controls the four available job types:
 - Optimization
 - Garbage collection
 - Scrubbing
 - Unoptimization

- Scrubbing. Data Deduplication has built-in data integrity features such as checksum validation and metadata consistency checking. It also has built-in redundancy for critical metadata and for the most popular data chunks. As data is accessed or deduplication jobs process data, these features might encounter corruption, and they will record the corruption in a log file. Scrubbing jobs use these features to analyze the chunk store corruption logs and, when possible, to make repairs. Possible repair operations include using three sources of redundant data:
 - Deduplication keeps backup copies of popular chunks when they are referenced more than 100 times in an area called the hotspot. If the working copy is corrupted, deduplication will use its own redundant copy in the case of soft corruptions such as bit flips or torn writes.



Note: *Bit flipping* is a term that refers to when data supposed to be a 0 is stored as a 1. *Torn writes* is when only a portion of the data is successfully stored on the volume.

- If using mirrored Storage Spaces, deduplication can use the mirror image of the redundant chunk to read the correct data and fix the corruption.
- When processing a corrupted chunk, the corrupted chunk is eliminated, and the Data Deduplication uses the new incoming chunk to fix the corruption.



Note: Because additional validations are built into deduplication, the deduplication subsystem is often the first system to report any early signs of data corruption in the hardware or file system.

- Unoptimization. This job undoes deduplication on all of the optimized files on the volume. Some of the common scenarios for using this type of job include decommissioning a server with volumes enabled for Data Deduplication, troubleshooting issues with deduplicated data, or migration of data to another system that does not support Data Deduplication. Prior to starting this job, you should use the **Disable-DedupVolume** Windows PowerShell cmdlet to disable further data deduplication activity on one or more volumes. After you disable Data Deduplication, the volume remains in the deduplicated state and the existing deduplicated data remains accessible; however, the server stops running optimization jobs for the volume and new data is not deduplicated. Later, you would use the unoptimization job to undo the existing deduplicated data on a volume. The unoptimization job deletes all of the data deduplication metadata from the volume after a successful run.

Data Deduplication process

In Windows Server 2016, Data Deduplication transparently removes duplication, without changing access semantics. When you enable Data Deduplication on a volume, a post-process, or *target*, deduplication optimizes the file data on the volume by performing the following actions:

1. Optimization jobs, which are background tasks, run with low priority on the server to process the files on the volume.
2. Deduplication segments all file data on the volume into small, variable-sized chunks that range from 32 KB to 128 KB. It does this by using an algorithm.
3. Deduplication identifies chunks that have one or more duplicates on the volume.
4. Deduplication inserts chunks into a common chunk store.
5. Deduplication replaces all duplicate chunks with a reference, or stub, to a single copy of the chunk in the chunk store.

6. Deduplication replaces the original files with a reparse point, which contains references to its data chunks.
7. Deduplication compresses chunks and organizes them in container files in the **System Volume Information** folder.
8. Deduplication removes the primary data stream of the files.

The Data Deduplication process works through scheduled tasks on the local server, but you can run the process interactively by using Windows PowerShell. You will find more information about this later in the module.

Data deduplication does not have any write performance impact, because the data is not deduplicated while the file is being written. Windows Server 2016 uses post-process deduplication, which maximizes deduplication potential. The other advantage with this type of deduplication process is that it offloads all processing from your application servers and client computers, which means less stress on the other resources in your environment. There is, however, a small performance impact when reading deduplicated files.



Note: The three main types of data deduplication are source, target (or post-process deduplication), and in-line (or transit deduplication).

Data Deduplication potentially can process all of the data on a selected volume, except for files that are less than 32 KB in size and files in folders that you exclude. You must carefully determine if a server and its attached volumes are suitable candidates for deduplication prior to enabling the feature. You should also consider backing up important data regularly during the deduplication process.

After you enable a volume for deduplication and the data is optimized, the volume contains the following elements:

- Unoptimized files. Includes files that do not meet the selected file-age policy setting, system state files, alternate data streams, encrypted files, files with extended attributes, files smaller than 32 KB, or other reparse point files.
- Optimized files. Includes files that are stored as reparse points that contain pointers to a map of the respective chunks in the chunk store that are needed to restore the file when it is requested.
- The chunk store. Location for the optimized file data.
- Additional free space. The optimized files and chunk store occupy much less space than they did prior to optimization.

Deploying Data Deduplication

Planning a Data Deduplication deployment

Prior to installing and configuring Data Deduplication in your environment, you must plan your deployment by using the following steps:

1. Determine target deployments. Data Deduplication is designed to be applied on primary data volumes, without adding any additional dedicated hardware. Data Deduplication is not designed for logically extended data volumes. You can schedule deduplication based on the type of data that is involved and the frequency and volume of changes that occur to the volume or particular file types. Consider using deduplication for the following data types:
 - General file shares. Group content publication and sharing, user home folders, and Folder Redirection/Offline Files.
 - Software deployment shares. Software binaries, images, and updates.
 - VHD libraries. VHD file storage for provisioning to hypervisors.
 - VDI deployments. Virtual Desktop Infrastructure (VDI) deployments using Hyper-V.
 - Virtualized backup. Backup applications running as Hyper-V guests that save backup data to mounted VHDs.
2. Determine which volumes are candidates for deduplication. Deduplication can be very effective for optimizing storage and reducing the amount of disk space consumed. You save 50 to 90 percent of your system's storage space when you apply deduplication to the right data. Use the following considerations to evaluate which volumes are ideal candidates for deduplication:
 - Is duplicate data present? File shares or servers that host user documents, software deployment binaries, or virtual hard disk files tend to have plenty of duplication and will yield higher storage savings from deduplication. More information about the deployment candidates for deduplication and the supported/unsupported scenarios are provided later in this module.
 - Does the data access pattern allow for sufficient time for deduplication? For example, files that change daily and that users or applications often access are not good candidates for deduplication. In these scenarios, deduplication might not be able to process the files, because the constant access and change to the data are likely to cancel any optimization gains made by deduplication. Good candidates allow time for deduplication of the files.
 - Does the server have sufficient resources and time to run deduplication? Deduplication requires reading, processing, and writing large amounts of data, which consumes server resources. Servers typically have periods of high activity and times when there is low resource utilization—the deduplication jobs work more efficiently when resources are available. However, if a server is constantly at maximum resource capacity, it might not be an ideal candidate for deduplication.
3. Evaluate savings with the Deduplication Evaluation Tool. You can use the Deduplication Evaluation Tool, DDPEval.exe, to determine the expected savings that you would get if deduplication is enabled on a particular volume. DDPEval.exe supports evaluating local drives and mapped or unmapped remote shares.

Prior to installing and configuring Data Deduplication in your environment, plan your deployment using the following steps:

1. Determine target deployments
2. Determine which volumes are candidates for deduplication
3. Evaluate savings with the Deduplication Evaluation Tool
4. Plan the rollout, scalability, and deduplication policies





Note: When you install the deduplication feature, DDPEval.exe is automatically installed to the `\Windows\System32\` directory. For more information about this tool, refer to: "Plan to Deploy Data Deduplication" at: <http://aka.ms/Oykadn>

4. Plan the rollout, scalability, and deduplication policies. The default deduplication policy settings are usually sufficient for most environments. However, if your deployment has any of the following conditions, consider altering the default settings:
 - Incoming data is static or expected to be read-only, and you want to process files on the volume sooner. In this scenario, change the **MinimumFileAgeDays** setting to a smaller number of days.
 - You have directories that you do not want to deduplicate. Add a directory to the exclusion list.
 - You have file types that you do not want to deduplicate. Add a file type to the exclusion list.
 - The server has different off-peak hours than the default, and you want to change the garbage collection and scrubbing schedules. Update the schedules by using Windows PowerShell.

Determining the size of the deduplicated volumes

When planning for Data Deduplication in your environment, one important consideration is the size of the deduplicated volumes. Although Windows Server 2016 supports Data Deduplication on volumes up to 64 TB, you must assess the appropriate size of the deduplicated volumes that your environment can support. Usually, the size depends on your hardware specifications and your unique workload. More specifically, it depends primarily on how much and how frequently the data on the volume changes and on the data access throughput rates of the disk storage subsystem.

Data Deduplication in Windows Server 2016 performs intensive I/O and compute operations. In most deployments, deduplication operates in the background or on a daily schedule on each day's new or modified data (that is, the data churn). As long as deduplication is able to optimize the entire data churn on a daily basis, the volume size will work for deduplication. Alternatively, some organizations simply create a 64 TB volume, enable deduplication, and then wonder why they experience low optimization rates. Most likely in this scenario, deduplication is not able to keep up with the incoming churn from a dataset that is too large on a configured volume. Although Data Deduplication in Windows Server 2016 runs multiple threads in parallel using multiple I/O queues on multiple volumes simultaneously, the deduplication environment might require additional computing power.

Consider the following factors when estimating the size of your volumes enabled for Data Deduplication:

- Deduplication optimization needs to be able to keep up with the daily data churn.
- The total amount of churn scales with the size of the volume.
- The speed of deduplication optimization significantly depends on the data access throughput rates of the disk storage subsystem.

Therefore, to estimate the maximum size for a deduplicated volume, you should be familiar with the size of the data churn and the speed of optimization processing on your volumes. You can choose to use reference data, such as server hardware specifications, storage-drive or array speed, and deduplication speed of various usage types, for your estimations. However, the most accurate method of assessing the appropriate volume size is to perform the measurements directly on your deduplication system based on the representative samples of your data, such as data churn and deduplication processing speed.

Installing and configuring Data Deduplication


After the planning phase, perform the following steps to deploy Data Deduplication to a server in your environment:

1. Install Data Deduplication components on the server. Use the following options to install deduplication components on the server:
 - o Server Manager. In **Server Manager**, perform the following steps:
 - i. Open the **Add Roles and Features Wizard**.
 - ii. Under **Server Roles**, select **File and Storage Services**, and then select the **File Services** check box.
 - iii. Select the **Data Deduplication** check box, and then click **Install**.
 - o Windows PowerShell. You can use the following commands to install Data Deduplication:

```
Import-Module ServerManager
Add-WindowsFeature -Name FS-Data-Deduplication
Import-Module Deduplication
```

2. Enable Data Deduplication. Use the following options to enable Data Deduplication on the server:
 - o Server Manager. From the **Server Manager** dashboard:
 - i. Right-click a data volume, and then select **Configure Data Deduplication**.
 - ii. In the **Data deduplication** box, select the workload that you want to host on the volume. For example, select General purpose file server for general data files or Virtual Desktop Infrastructure (VDI) server when configuring storage for running virtual machines.
 - iii. Enter the minimum number of days that should elapse from the date of file creation before files are deduplicated. Enter the extensions of any file types that should not be deduplicated, and then click **Add** to browse to any folders with files that should not be deduplicated.
 - iv. Click **Apply** to apply these settings and return to the Server Manager dashboard, or click the **Set Deduplication Schedule** button to continue to set up a schedule for deduplication.
 - o Windows PowerShell. Use the following command to enable deduplication on a volume:

```
Enable-DedupVolume -Volume VolumeLetter -UsageType StorageType
```

 **Note:** Replace *VolumeLetter* with the drive letter of the volume. Replace *StorageType* with the value corresponding to the expected type of workload for the volume. Acceptable values include:

- HyperV. A volume for Hyper-V storage.
- Backup. A volume that is optimized for virtualized backup servers.
- Default. A general-purpose volume.

Optionally, you can use the **Set-DedupVolume** Windows PowerShell cmdlet to configure additional options, such as the minimum number of days that should elapse from the date of file creation before files are deduplicated, the extensions of any file types that should not be deduplicated, or the folders that should be excluded from deduplication.

3. Configure Data Deduplication jobs. With Data Deduplication jobs, you can run them manually, run them on demand, or use a schedule. There are four types of jobs that you can perform on a volume:

- Optimization. Includes built-in jobs that are automatically scheduled for optimizing the volumes on a periodic basis. Optimization jobs deduplicate data and compress file chunks on a volume in accordance with the policy settings. In addition to **Server Manager**, you also can use the following command to trigger an optimization job on demand:

```
Start-DedupJob -Volume VolumeLetter -Type Optimization
```

- Data scrubbing. Scrubbing jobs are scheduled automatically to analyze the volume on a weekly basis and produce a summary report in the Windows event log. You also can use the following command to trigger a scrubbing job on demand:

```
Start-DedupJob -Volume VolumeLetter -Type Scrubbing
```

- Garbage collection. Garbage collection jobs are scheduled automatically to process data on the volume on a weekly basis. Because garbage collection is a process-intensive operation, consider waiting until after the deletion load reaches a threshold to run this job on demand, or schedule the job for after business hours. You also can use the following command to trigger a garbage collection job on demand:

```
Start-DedupJob -Volume VolumeLetter -Type GarbageCollection
```

- Unoptimization. Unoptimization jobs are available on an as-needed basis and are not scheduled automatically. However, you can use the following command to trigger an unoptimization job on demand:

```
Start-DedupJob -Volume VolumeLetter -Type Unoptimization
```



Note: For more information on the Windows PowerShell cmdlet **Start-DedupJob**, refer to: <http://aka.ms/Yphsar>

4. Configure Data Deduplication schedules. When you enable Data Deduplication on a server, three schedules are enabled by default: optimization is scheduled to run every hour, and garbage collection and scrubbing are scheduled to run once a week. You can view the schedules by using the **Get-DedupSchedule** Windows PowerShell cmdlet. These scheduled jobs run on all the volumes on the server. However, if you want to run a job only on a particular volume, you must create a new job. You can create, modify, or delete job schedules from the **Deduplication Settings** page in **Server Manager**, or by using the following Windows PowerShell cmdlets: **New-DedupSchedule**, **Set-DedupSchedule**, or **Remove-DedupSchedule**.



Note: Data Deduplication jobs only support, at most, weekly job schedules. If you need to create a schedule for a monthly job or for any other custom time period, use Task Scheduler. However, you will be unable to view these custom job schedules created with Task Scheduler by using the **Get-DedupSchedule** Windows PowerShell cmdlet.

Demonstration: Implementing Data Deduplication

In this demonstration, you will see how to:

- Install the Data Deduplication role service.
- Enable Data Deduplication.
- Check the status of Data Deduplication.

Demonstration Steps

Install the Data Deduplication role service

- On **LON-SVR1**, in **Server Manager**, add the **Data Deduplication** role service.

Enable Data Deduplication

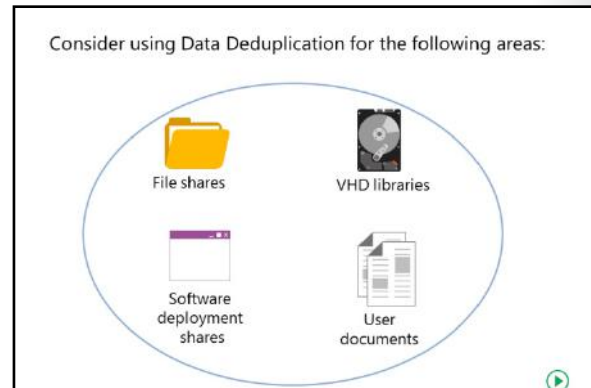
1. In **Server Manager**, click **File and Storage Services**.
2. Click **Disks**.
3. Click the **1** disk, and then click the **E** volume.
4. Enable **Data Deduplication**, and then click the **General purpose file server** setting.
5. Configure the following settings:
 - Deduplicate files older than (in days): **5**
 - Enable throughput optimization

Check the status of Data Deduplication

1. Switch to Windows PowerShell.
2. Execute the following commands to verify Data Deduplication status:
 - **Get-DedupStatus**
 - **Get-DedupStatus | fl**
 - **Get-DedupVolume**
 - **Get-DedupVolume | fl**
 - **Start-DedupJob E: -Type Optimization -Memory 50**
3. Repeat commands 2a and 2c. Notice that the available free space increases on drive E.
4. Close all open windows.

Usage scenarios for Data Deduplication

Your data storage savings will vary by data type, the mix of data, and the size of the volume and the files contained within the volume. Consider using the Deduplication Evaluation Tool to evaluate the volumes before you enable deduplication. The following table highlights typical deduplication savings for various content types.



Content type	Description	Savings
User documents	This includes group content publication or sharing, user home folders (or MyDocs), and profile redirection for accessing offline files.	30–50 percent of your system's storage space
Software deployment shares	This includes software binaries, cab files, symbols files, images, and updates.	70–80 percent of your system's storage space
Virtualization libraries	This includes virtual hard disk files (.vhd and .vhdx files) for provisioning to hypervisors.	80–95 percent of your system's storage space
General file share	This includes a mix of all the types of data identified above.	50–60 percent of your system's storage space

Volume requirements for Data Deduplication

After you install the role service, you can enable Data Deduplication on a per-volume basis. Data Deduplication includes the following requirements:

- Volumes must not be system or boot volumes. Because most files used by an operating system are constantly open, Data Deduplication on system volumes would negatively affect the performance because deduplicated data would need to be expanded again before the files could be used.
- Volumes might be partitioned by using MBR or GPT format, and must be formatted by using the NTFS or ReFS file system.
- Volumes must be attached to the Windows Server drives and cannot appear as non-removable drives. This means that you cannot use universal serial bus (USB) or floppy drives for Data Deduplication, nor can you use remotely mapped drives.
- Volumes can be on shared storage, such as Fibre Channel, iSCSI SAN, or serial-attached SCSI array.
- Files with extended attributes, encrypted files, files smaller than 32 KB, and reparse point files will not be processed for Data Deduplication.
- Data Deduplication is not available for Windows client operating systems.

Data Deduplication deployment candidates

Based on observed savings and typical resource usage in Windows Server 2016, deployment candidates for deduplication are grouped as follows:

- Ideal candidates for deduplication. They are:
 - Folder redirection servers
 - Virtualization depot or provisioning library
 - Software deployment shares
 - Microsoft SQL Server and Microsoft Exchange Server backup volumes
 - Cluster Shared Volumes (CSVs) in Scale-Out File Servers
 - Virtualized backup VHDs (for example, Data Protection Manager)
 - VDI VHDs (only personal VDIs)



Note: In most VDI deployments, special planning is required for the boot storm. *Boot storm* is the phenomenon of large numbers of users trying to simultaneously sign in to their VDI—typically upon arriving to work in the morning. In turn, this puts pressure on the VDI storage system and can cause long delays for VDI users. However, in Windows Server 2016, when the virtual machine reads chunks from the on-disk deduplication store during startup, they are cached in memory. As a result, subsequent reads do not require frequent access to the chunk store because the cache intercepts them. This minimizes the effects of the boot storm because the memory is much faster than the disk.

- Candidates that you should evaluate based on their content. They are:
 - Line-of-business servers
 - Static content providers
 - Web servers
 - High-performance computing (HPC)
- Not ideal candidates for deduplication. They are:
 - Hyper-V hosts
 - Windows Server Update Services (WSUS)
 - SQL Server and Exchange Server database volumes

Data Deduplication interoperability

When Windows Server 2016, you should consider the following related technologies and potential issues when deploying Data Deduplication:

- BranchCache. Access to data over the network can be optimized by enabling BranchCache on Windows servers and clients. When a BranchCache-enabled system communicates over a WAN with a remote file server that is enabled for Data Deduplication, all of the deduplicated files are already indexed and hashed, so requests for data from a branch office are quickly computed. This is similar to preindexing or prehashing a BranchCache-enabled server.

- **Failover clusters.** Windows Server 2016 fully supports failover clusters, which means deduplicated volumes will failover smoothly with no interruption between nodes in the cluster. Effectively, a deduplicated volume is a self-contained and portable unit (that is, all of the data and configuration information is contained on the volume), but it requires that each node in the cluster that accesses deduplicated volumes must be running the Data Deduplication feature. When a cluster is formed, the deduplication schedule information is configured in the cluster. As a result, if another node takes over a deduplicated volume, then the scheduled jobs will be applied on the next scheduled interval by the new node.
- **FSRM quotas.** Although you should not create a hard quota on a volume root folder enabled for deduplication, by using FSRM, you can create a soft quota on a volume root that is enabled for deduplication. When FSRM encounters a deduplicated file, it will identify the file's logical size for quota calculations. Consequently, quota usage (including any quota thresholds) does not change when deduplication processes a file. All other FSRM quota functionality, including volume-root soft quotas and quotas on subfolders, will work as expected when using deduplication.
- **DFS Replication.** Data Deduplication is compatible with Distributed File System (DFS) Replication. Optimizing or unoptimizing a file will not trigger a replication because the file does not change. DFS Replication uses remote differential compression (RDC) and not the chunks in the chunk store for over-the-wire savings. In fact, you can optimize the files on the replica instance by using deduplication if the replica is enabled for Data Deduplication.



Note: Single Instance Store (SIS), a file system filter driver used for NTFS file deduplication, was deprecated in Windows Server 2012 R2 and is completely removed in Windows Server 2016.

Monitoring and maintaining Data Deduplication

After you deploy Data Deduplication in your environment, it is important that you monitor and maintain the systems that you enable for Data Deduplication, and the corresponding data storage, to ensure optimal performance. Although Data Deduplication in Windows Server 2016 includes a lot of automation, including optimization jobs, the deduplication process requires that you verify the efficiency of optimization, make the appropriate adjustments to systems, storage architecture, and volumes, and troubleshoot any issues with Data Deduplication.

By using the **Get-DedupStatus** and **Get-DedupVolume** Windows PowerShell cmdlets, you can:

- Query the progress of a deduplication job
- View the space savings that have been achieved on a volume
- View the current status of the deduplication process

Monitoring and reporting of Data Deduplication

Consider using the following options to monitor deduplication in your environment and to report on its health:

- **Windows PowerShell cmdlets.** After you enable the Data Deduplication feature on a server, you can use the following Windows PowerShell cmdlets:
 - **Get-DedupStatus.** This is the most commonly used option. This cmdlet returns the deduplication status for volumes that have data deduplication metadata. The output of the cmdlet includes the deduplication rate, the number of optimized files, the last run time of the deduplication jobs, and the amount of space saved on the volume.

- **Get-DedupVolume.** This cmdlet returns the deduplication status for volumes that have data deduplication metadata. The metadata includes the deduplication rate, the number of optimized files, and deduplication settings such as minimum file age, minimum file size, excluded files/folders, compression-excluded file types, and the chunk redundancy threshold.
- **Get-DedupMetadata.** This cmdlet returns status information of the deduplicated data store for volumes that have Data Deduplication metadata. The cmdlet returns the number of data chunks in a container, the number of containers in the data store, the number of data streams in a container, the number of containers in the stream map store, the number of hotspots in a container, the number of hotspots in the stream map store, and the number of corruptions on the volume.
- **Get-DedupJob.** This cmdlet returns the deduplication status and information for currently running or queued deduplication jobs.

One common scenario is to assess whether deduplication is keeping pace with the rate of incoming data. You can use the **Get-DedupStatus** cmdlet to monitor the number of optimized files compared with the number of in-policy files. This enables you to see if all the in-policy files are processed. If the number of in-policy files is continuously rising faster than the number of optimized files, you should examine your hardware specifications for appropriate utilization or the type of data on the volume usage type to ensure deduplication efficiency. However, if the output value from the cmdlet for **LastOptimizationResult** is 0x00000000, then the entire dataset was processed successfully during the previous optimization job.



Note: For more information about all the storage cmdlets in Windows Server 2016, refer to: <http://aka.ms/Hdjymw>

- Event Viewer logs. Monitoring the event log can also be helpful with understanding deduplication events and status. To view deduplication events, in **Event Viewer**, navigate to **Applications and Services Logs**, click **Microsoft**, click **Windows**, and then click **Deduplication**. For example, Event ID 6153 will provide you with the elapsed time of a deduplication job and the throughput rate.
- Performance monitor data. In addition to using the counters for monitoring server performance, such as CPU and memory, you can use the typical disk counters to monitor the throughput rates of the jobs that are currently running such as **Disk Read Bytes/sec**, **Disk Write Bytes/sec**, and **Average Disk sec/Transfer**. Depending on other activities on the server, you might be able to use the data results from these counters to get a rough estimate of the savings by examining how much data is being read and how much is being written per interval. You can also use the **Resource Monitor** to identify the resource usage of specific programs or services. To view disk activity, in **Resource Monitor**, filter the list of processes to locate **fsdmhost.exe**, and examine the I/O on the files under the **Disk** tab.



Note: **Fsdmhost.exe** is the executable file for the Microsoft File Server Data Management Host process, which the Data Deduplication process in Windows Server 2016 uses.

- File Explorer. While File Explorer is not the ideal choice for validating deduplication on an entire volume, you can use it to spot check deduplication on individual files. By viewing the properties of a file, you will notice that **Size** displays the logical size of the file, and **Size on Disk** displays the true physical allocation of the file. For an optimized file, **Size on Disk** is less than the actual file size. This is because deduplication moves the contents of the file to a common chunk store and replaces the original file with an NTFS reparse point stub and metadata.

Maintaining Data Deduplication

With the data that is collected during monitoring, you can use the following Windows PowerShell cmdlets to ensure optimal efficiency of deduplication in your environment:

- **Update-DedupStatus.** You can use some of the storage cmdlets, such as **Get-DedupStatus** and **Get-DedupVolume**, to retrieve information from the cached metadata. This cmdlet scans volumes to compute new Data Deduplication information for updating the metadata.
- **Start-DedupJob.** This cmdlet launches ad hoc deduplication jobs, such as optimization, garbage collection, scrubbing, and unoptimization. For example, you might consider launching an ad hoc optimization job if a deduplicated volume is low on available space because of extra churn.
- **Measure-DedupFileMetadata.** This cmdlet measures potential disk space on a volume. More specifically, this cmdlet returns how much disk space you can reclaim on a volume if you delete a group of folders and subsequently run a garbage collection job. Files often have chunks that are shared across other folders. The deduplication engine calculates which chunks are unique and would be deleted after the garbage collection job.
- **Expand-DedupFile.** This cmdlet expands an optimized file into its original location. You might need to expand optimized files because of compatibility with applications or other requirements. Ensure that there is enough space on the volume to store the expanded file.

Troubleshooting adverse effects of Data Deduplication

When an application or access to a file is adversely impacted by Data Deduplication in Windows Server 2016, several options are available:

- Use a different deduplication frequency by changing the schedule or opting for manual deduplication jobs.
- Use job options such as:
 - **StopWhenSystemBusy**, which will halt deduplication if the job interferes with the server's workload.
 - **Preempt**, which will cause the deduplication engine to move specific deduplication jobs to the top of the job queue and cancel the current job.
 - **ThrottleLimit**, which will set the maximum number of concurrent operations that specific deduplication jobs can establish.
 - **Priority**, which will set the CPU and I/O priority for specific deduplication jobs.
 - **Memory**, which will specify the maximum percentage of physical computer memory that the data deduplication job can use.

- Common causes include:
 - Copying data with incompatible Robocopy options
 - Using incompatible backup or restore applications
 - Using a deduplicated volume on a previous version of Windows Server
 - Compressing the root of a volume
 - Hardware issues
 - General corruption
- Troubleshooting includes:
 - Checking event logs
 - Running **CHKDSK** in read-only mode
 - Running deep scrubbing job



Note: Although we recommend that you allow deduplication to manage memory allocation automatically, you might need to adjust the maximum percentage in some scenarios. For most of these scenarios, consider a maximum percentage within a range of 15 to 50 and a higher memory consumption for jobs that you schedule to run when you specify the *StopWhenSystemBusy* parameter. For garbage collection and scrubbing deduplication jobs, which you typically schedule to run after business hours, consider using a higher memory consumption such as 50.

- Use the **Expand-DedupFile** cmdlet to expand or undo deduplication of specific files if needed for compatibility or performance.
- Use the **Start-DedupJob** cmdlet with the **Unoptimization** job type to disable deduplication on a volume.

Troubleshooting Data Deduplication corruptions

Data integrity is highly important in deduplication, because a large number of deduplicated files might be referencing a single popular chunk that, if corrupt, can lead to data loss. Data Deduplication in Windows Server 2016 provides functionality to detect, report, and even repair data corruptions. Although there are a number of features built into deduplication to help protect against corruption, there are still some scenarios where deduplication might not recover automatically from corruption.

Some of the most common causes for deduplication to report corruption are:

- Incompatible Robocopy options used when copying data. Using Robocopy with the /MIR option on the volume root as the target will wipe the deduplication store. To avoid this problem, use the /XD option to exclude the **System Volume Information** folder from the scope of the Robocopy command.




Note: For more information on this issue, refer to: <http://aka.ms/Pkxrmx>

- Incompatible backup or restore program used on a deduplicated volume. You should verify whether your backup solution supports Data Deduplication in Windows Server 2016. Unsupported backup solutions might introduce corruptions after a restore. You will find more information about this in the next topic.
- Migrating a deduplicated volume to an earlier version of Windows Server. You may see file corruption messages if you mount a deduplicated volume on an earlier version of Windows Server than the version of Windows Server that did the actual deduplication. In this scenario, you should verify that the version of the server accessing the deduplicated data is the same version level or newer than the version of the server that optimized the data on the volume. Although you can remount deduplicated volumes on different servers, deduplication is backward compatible but not forward compatible. You can upgrade and migrate to a newer version of Windows Server, but older versions of Windows Server cannot read data deduplicated by a newer version of Windows Server, and the server might report the data as corrupted when trying to read it.
- Enabling compression on the root of a volume that is also enabled with deduplication. Deduplication is not supported on volumes that have compression enabled at the root. As a result, this might lead to the corruption and inaccessibility of deduplicated files.



Note: Deduplication of files in compressed folders is supported in Windows Server 2016 and should function normally.

- Hardware issues. Many hardware storage issues are detectable early by using the deduplication scrubbing job. Refer to the following general corruption troubleshooting steps for more information.
- General corruption. You can use the following steps to troubleshoot most general causes for deduplication to report corruption:
 - a. Check the event logs for details of corruption. Any corruption detected by deduplication is logged to the event log. The scrubbing channel lists any corruptions that the scrubbing job detected and files that the job attempted to fix. The deduplication Scrubbing event logs are located in the **Event Viewer**: under **Application and Services > Microsoft > Windows > Deduplication > Scrubbing**. In addition, searching for hardware events in the System Event logs and Storage Spaces Event logs will often provide additional information about hardware issues.


 **Note:** The potentially large number of events in the deduplication Scrubbing event log might be difficult to parse through the **Event Viewer**. For this reason, a publicly available script is available that generates an easy-to-read HTML that highlights detected corruptions and the results of any attempted corruption fixes from the scrubbing job. For more information about **Get-DedupScrubbingReport**, refer to: <http://aka.ms/mm2y53>

- b. Run **CHKDSK** in the read-only mode. Although this command can repair some data corruption on volumes, running the command without any parameters will initiate a read-only scan.

 **Note:** For more information on **CHKDSK** in Windows Server 2016, refer to: <http://aka.ms/Mftyj3>


- c. Run a deep scrubbing job to investigate detected corruptions. A must for corruption investigations, use a deep scrubbing job to ensure that all corruptions are logged in the deduplication scrubbing channel in the event logs. The scrubbing events will provide a breakdown of the corruptions, including corrupted chunks, affected files, the exact container offsets of the corruption, and the list of affected files (up to 10,000 files). You can use the following command in Windows PowerShell to initiate a deep scrubbing job:

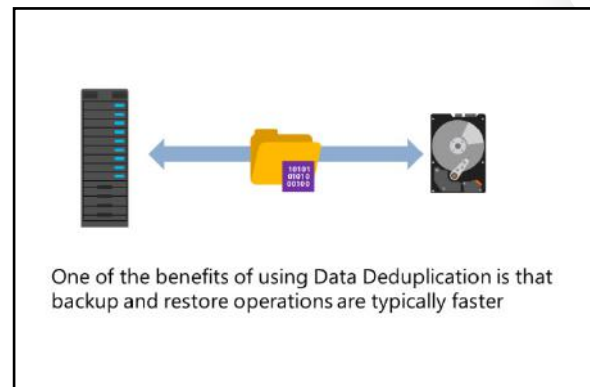
```
Start-DedupJob VolumeLetter -Type Scrubbing -Full
```

 **Note:** Replace *VolumeLetter* with the drive letter of the volume.

Backup and restore considerations with Data Deduplication

One of the benefits of using Data Deduplication is that backup and restore operations will be faster. This is because you have reduced the space used on a volume, which means there is less data to back up. When you perform an optimized backup, your backup is also smaller. This is because the total size of the optimized files, non-optimized files, and data deduplication chunk store files are much smaller than the logical size of the volume.

 **Note:** Many block-based backup systems should work with data deduplication by



maintaining the optimization on the backup media. File-based backup operations that do not use deduplication usually copy the files in their original format.

The following backup and restore scenarios are supported with deduplication in Windows Server 2016:

- Individual file backup and restore
- Full volume backup and restore
- Optimized file-level backup and restore using VSS writer

However, the following backup and restore scenarios are not supported with deduplication in Windows Server 2016:

- Backup or restore of only the reparse points
- Backup or restore of only the chunk store

In addition, a backup application can perform an incremental optimized backup as follows:

- Back up only the changes since your last backup. This includes files created, modified, or deleted.
- Back up the changed chunk store container files.
- Perform an incremental backup at the sub-file level.



Note: Deduplication appends new chunks to the current chunk store container. When its size reaches approximately 1 GB, then deduplication closes that container file and creates a new container file.

Restore operations

Restore operations also can benefit from data deduplication. Any file-level full volume restore operations can benefit because they are essentially a reverse of the backup procedure, and less data means quicker operations. A full volume restore involves the restoration of:

1. The complete set of data deduplication metadata and container files.
2. The complete set of data deduplication reparse points.
3. All non-deduplicated files.

Block-level restore from an optimized backup is automatically an optimized restore because the restore process occurs under Data Deduplication, which works at the file level.

As with any product from a non-Microsoft vendor, you should verify whether the backup solution supports Data Deduplication in Windows Server 2016. Unsupported backup solutions might introduce corruptions after a restore. The common methods on solutions that support Data Deduplication in Windows Server 2016 include:

- Some backup vendors support *unoptimized* backup, which backs up the deduplicated files as normal, full-size files.
- Some backup vendors support *optimized* backup for a full volume backup, which backs up the deduplicated files as they are, as a reparse point stub within the chunk store.
- Some backup vendors support both.

The backup vendor should be able to comment on what their product supports and the method it uses.



Note: For more information about the backup and restore of volumes enabled for Data Deduplication, refer to: <http://aka.ms/G4fm7k>

Question: Can you enable Data Deduplication on ReFS formatted drives?

Question: Can you enable Data Deduplication on volumes where virtual machines are running and also apply it to those virtual machines?

Check Your Knowledge

Question	
Which features of Data Deduplication are present in Windows Server 2016? (Select all that apply.)	
Select the correct answer.	
<input type="checkbox"/>	Deduplication of volumes up to 64 TB
<input type="checkbox"/>	Deduplication of volumes larger than 64 TB
<input type="checkbox"/>	Support for virtualized backup storage
<input type="checkbox"/>	Support for storage on Hyper-V virtual machines
<input type="checkbox"/>	Support for Nano Server

Lesson 3

Configuring iSCSI storage

iSCSI storage is an inexpensive and simple way to configure a connection to remote disks. Many application requirements dictate that remote storage connections must be redundant to provide fault tolerance or high availability. For this purpose, you will learn how to create a connection between servers and iSCSI storage. You also will learn how to create both single and redundant connections to an iSCSI target. You will do this by using the iSCSI initiator software that is available in Windows Server 2016.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe iSCSI.
- Describe the components of iSCSI.
- Describe how to managed iSCSI targets.
- Configure the iSCSI target.
- Configure the iSCSI storage.
- Describe how to implement high availability for iSCSI.

Overview of iSCSI

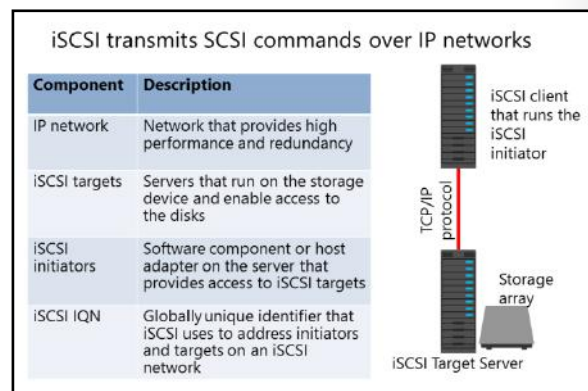
iSCSI is a protocol that supports access to remote, SCSI-based storage devices over a TCP/IP network. iSCSI carries standard SCSI commands over IP networks to facilitate data transfers, and to manage storage over long distances. You can use iSCSI to transmit data over local area networks (LANs), WANs,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

iSCSI relies on standard Ethernet networking architecture, and the use of specialized hardware, such as a host bus adapter (HBA) or network switches, is optional. iSCSI uses TCP/IP (typically, TCP port 3260). This means that iSCSI enables two hosts to negotiate (session establishment, flow control, and packet size, for example) and then exchange SCSI commands by using an existing Ethernet network. By doing this, iSCSI takes a popular, high-performance, local storage bus-subsystem architecture and emulates it over networks, thereby creating a SAN.

Unlike some SAN protocols, iSCSI requires no specialized cabling; you can run it over existing switching and IP infrastructure. However, to ensure performance you should operate an iSCSI SAN deployment on a dedicated network. Otherwise, you might experience severely decreased performance.

An iSCSI SAN deployment includes the following:

- IP network. You can use standard network interface adapters and standard Ethernet protocol network switches to connect the servers to the storage device. To provide sufficient performance, the network should provide speeds of at least 1 gigabit per second (Gbps), and should provide multiple paths to the iSCSI target. We recommend that you use a dedicated physical and logical network to achieve fast, reliable throughput.



- **iSCSI targets.** iSCSI targets present or advertise storage, similar to controllers for hard disk drives of locally attached storage. However, servers 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 devices, implement iSCSI targets by using a software driver and at least one Ethernet adapter. Windows Server 2016 provides the iSCSI Target Server, which is effectively a driver for the iSCSI protocol, as a role service of the File and Storage Services role.
- **iSCSI initiators.** The iSCSI target displays storage to the iSCSI initiator (also known as the client). The iSCSI initiator acts as a local disk controller for the remote disks. All Windows versions since Windows Server 2008 and Windows Vista include the iSCSI initiator and can connect to iSCSI targets.
- **iSCSI qualified name (IQN).** IQNs are unique identifiers that iSCSI uses to address initiators and targets on an iSCSI network. When you configure an iSCSI target, you must configure the IQN for the iSCSI initiators that will be connecting to the target. iSCSI initiators also use IQNs to connect to the iSCSI targets. However, if name resolution on the iSCSI network is a possible issue, you can always identify iSCSI endpoints (both target and initiator) by their IP addresses.

Components of iSCSI

This topic discusses the two main components of iSCSI: an iSCSI Target Server and an iSCSI initiator. This topic also discusses about Internet Storage Name Service (iSNS) and Data Center Bridging.

iSCSI Target Server

The iSCSI Target Server role service provides for a software-based and hardware-independent iSCSI disk subsystem. You can use the iSCSI Target Server to create iSCSI targets and iSCSI virtual disks, and then use **Server Manager** to manage these iSCSI targets and virtual disks. In Windows Server 2016, the iSCSI Target Server is available as a role service under the File and Storage Services role in **Server Manager**.

The iSCSI Target Server:	The iSCSI initiator:
<ul style="list-style-type: none"> • Is available as a role service in Windows Server 2016 • Provides the following functionality: <ul style="list-style-type: none"> • Network or diskless boot • Server application storage • Heterogeneous storage • Lab environments • Has the following features: <ul style="list-style-type: none"> • Authentication • Query initiator computer for ID • Virtual hard disks • Scalability • Manageability 	<ul style="list-style-type: none"> • Runs as a service in the operating system • Is installed by default on Windows Vista and Windows Server 2008 and later operating systems • Requires only to be started and configured to connect the computer to the iSCSI target

The iSCSI Target Server that Windows Server 2016 includes provides the following functionality:

- **Network or diskless boot.** You can rapidly deploy diskless servers by using boot-capable network adapters or a software loader, and you can save as much as 90 percent of the storage space that you use for operating-system images by using differencing virtual hard disks. This is ideal for large deployments of identical operating-system images, such as on virtual machines that are running Hyper-V or in HPC clusters.
- **Server application storage.** Some applications, such as Exchange Server, require block storage. The iSCSI Target Server can provide these applications with continuously available block storage. Because the storage is remotely accessible, it also can combine block storage for central or branch office locations.
- **Heterogeneous storage.** iSCSI Target Server supports iSCSI initiators that are not based on Windows, so you can share storage on servers that are running Windows in mixed environments.
- **Lab environments.** The iSCSI Target Server role enables your Windows Server 2016 computer to be a network-accessible block storage device. This is useful in situations when you want to test applications before deploying on SAN storage.

The features of the iSCSI Target Server in Windows Server 2016 include:

- **Authentication.** You can enable Challenge Handshake Authentication Protocol (CHAP) to authenticate initiator connections or enable reverse CHAP to allow the initiator to authenticate the iSCSI target.
- **Query initiator computer for ID.** To use this, you must use Windows 8 or Windows Server 2012 and newer operating systems.
- **Virtual hard-disk support.** You create iSCSI virtual disks as virtual hard disks. Windows Server 2016 support both VHD and VHDX files. VHDX supports up to 64 TB capacity. You create new iSCSI virtual disks as VHDX files, but you can import VHD files.
- **Scalability.** The maximum number of iSCSI targets per target server is 256 and the maximum number of virtual hard disks per target server is 512.



Additional Reading: For more information on iSCSI Target Server scalability limits, refer to: <http://aka.ms/dfxgja>

- **Manageability.** You manage the iSCSI Target Server by using **Server Manager** or Windows PowerShell. Windows Server 2016 uses the Storage Management Initiative Specification provider with Microsoft System Center 2012 Virtual Machine Manager and later versions to manage an iSCSI Target Server in a hosted and private cloud.

The following Windows PowerShell cmdlets are some examples of managing the iSCSI Target Server:

```
Install-WindowsFeature FS-iSCSITarget-Server
New-IscsiVirtualDisk E:\iSCSIVirtualHardDisk\1.vhdx -size 1GB
New-IscsiServerTarget SQLTarget -InitiatorIds "IQN: iqn.1991-05.com.Microsoft:SQL1.adatum.com"
Add-IscsiVirtualDiskTargetMapping SQLTarget E:\iSCSIVirtualHardDisk\1.vhdx
```



Additional Reading: For more information, refer to: "iSCSI Target Cmdlets in Windows PowerShell" at: <http://aka.ms/j1iomo>

When you enable the iSCSI Target Server to provide block storage, the iSCSI Target Server capitalizes on your existing Ethernet network. You need a dedicated network for iSCSI to ensure performance or you can use quality of service on your existing network. If high availability is an important criterion, you should set up a high-availability cluster. However, when you configure a high-availability cluster, you will need shared storage for the cluster. This storage can be either hardware Fibre Channel storage or a serial attached SCSI storage array. You configure the iSCSI Target Server as a cluster role in the failover cluster. Windows Server 2016 introduces Storage Spaces Direct that can use nonshared storage to make a high availability cluster by using only local nonshared storage and commodity hardware.

iSCSI initiator

The iSCSI initiator was introduced in Windows Server 2008 and Windows Vista, and it is installed by default. To connect your computer to an iSCSI target, you only have to start the service and configure it.

The following Windows PowerShell cmdlets are some examples of managing the iSCSI initiator:

```
Start-Service msiscsi
Set-Service msiscsi -StartupType "Automatic"
New-IscsiTargetPortal -TargetPortalAddress iSCSIServer1
Connect-IscsiTarget -NodeAddress "iqn.1991-05.com.microsoft:netboot-1-SQLTarget-target"
```



Additional Reading: For more information, refer to: <http://aka.ms/ygfwgd>

Internet Storage Name Service (iSNS)

You can use the iSNS protocol when the iSCSI initiator attempts to discover iSCSI targets. The **iSNS Server service** feature in Windows Server 2016 provides storage discovery and management services to a standard IP network. Together with iSCSI Target Server, iSNS functions almost like a SAN. iSNS facilitates the integration of IP networks and manages iSCSI devices.

The iSNS server has the following functionality:

- Contains a repository of active iSCSI nodes.
- Contains iSCSI nodes that can be initiators, targets, or management nodes.
- Allows initiators and targets to register with the iSNS server, and the initiators then query the iSNS server for the list of available targets.
- Contains a dynamic database of the iSCSI nodes. The database provides the iSCSI initiators with iSCSI target discovery functionality. The database is updated automatically using the Registration Period and Entity Status Inquiry features of iSNS. Registration Period allows iSNS to delete stale entries from the database. Entity Status Inquiry is similar to ping. It allows iSNS to determine whether registered nodes are still present on the network, and enables iSNS to delete entries in the database that are no longer active.
- Provides State Change Notification Service. Registered clients receive notifications when changes occur to the database in the iSNS server. Clients keep their information about the iSCSI devices available on the network up-to-date with these notifications.
- Provides Discovery Domain Service. You can divide iSCSI nodes into one or more groups called discovery domains. Discovery domains provide zoning so that an iSCSI initiator can only view and connect to iSCSI targets in the same discovery domain.

Data Center Bridging

Data Center Bridging is a collection of standards-based networking technologies defined by The Institute of Electrical and Electronics Engineers Inc. (IEEE). It allows multiple types of traffic to run on the same physical Ethernet network cables in the datacenter. Data Center Bridging uses hardware-based bandwidth allocation and priority-based flow control instead of the operating system having to handle the traffic.

Windows Server 2012 and newer operating systems support Data Center Bridging by installing the **Data Center Bridging** feature. The advantage of using Data Center Bridging is the ability to run all Ethernet traffic including traffic going to and from your Fibre Channel or iSCSI SANs. This saves cabling, network equipment, space, and power in your datacenter. Data Center Bridging is also called Data Center Ethernet or Converged Enhanced Ethernet or converged networking. For more information about converged networking, refer to Module 5, "Implementing network services."

Data Center Bridging requires compatible network adapters and switches and is currently only configurable via Windows PowerShell. When you install the **Data Center Bridging** feature, you can use the following three Windows PowerShell modules: **netqos**, **dcbbqs**, and **netadapter**.

Managing iSCSI targets

After the initial configuration of an iSCSI target, storage needs might change. The iSCSI Target Server supports resizing the iSCSI virtual disks from their original size. This includes both expanding and compacting the iSCSI virtual disk.

The **Resize-IscsiVirtualDisk** Windows PowerShell cmdlet resizes a virtual disk by expanding or compacting an existing virtual disk. You need to take virtual disks that use the VHD format offline to resize them. However, virtual disks that use the VHDX format can stay online while the resizing occurs. You can resize a virtual disk, even when an iSCSI Target uses it.

Type	Create	Import	Expand	Shrink
Fixed VHD	No	Yes	Offline	No
Fixed VHDX	Yes	Yes	Yes	Yes
Differencing VHD	No	Yes	Offline	No
Differencing VHDX	Yes	Yes	Yes	Yes
Dynamic VHD	No	Yes	Offline	No
Dynamic VHDX	Yes	Yes	Yes	Yes

iSCSI initiators can continue to access the virtual disk over the network while the cmdlet resizes virtual disk. However, the initiators do not automatically use the new size of the virtual disks. You can use the **Resize-Partition** cmdlet to modify the volume hosted on the virtual disk, thereby acquiring access to the new virtual disk's capacity.

To resize a VHDX file from its current size to 100 GB, enter the following Windows PowerShell cmdlet at a command line, and then press Enter:

```
Resize-IscsiVirtualDisk -Path "E:\iSCSIVirtualDisks\iSCSIDisk1.vhdx" -Size 100GB
```



Additional Reading: For more information, refer to: "iSCSI Target Cmdlets in Windows PowerShell" at: <http://aka.ms/j1iomo>

Although most resizing operations will expand the size of a virtual disk, beginning in Windows Server 2012 R2, you also can shrink virtual disks. However, you can shrink virtual disks only that are created in the VHDX format. You also can only shrink the unallocated part of a virtual disk.

You cannot perform all operations on all virtual disks. Which actions you can perform depends on whether a virtual disk is in the VHD or VHDX format. The following table lists the supported operations on iSCSI virtual disks.

Type	Create	Import	Expand	Shrink
Fixed VHD	No	Yes	Offline	No
Fixed VHDX	Yes	Yes	Yes	Yes
Differencing VHD	No	Yes	Offline	No
Differencing VHDX	Yes	Yes	Yes	Yes
Dynamic VHD	No	Yes	Offline	No
Dynamic VHDX	Yes	Yes	Yes	Yes

Demonstration: Configuring the iSCSI target

In this demonstration, you will see how to:

- Add an iSCSI Target Server role service.
- Create two iSCSI virtual disks and an iSCSI target on LON-DC1.

Demonstration Steps

Add an iSCSI Target Server role service

1. On **LON-DC1**, open a **Windows PowerShell** window, and type the following command, and then press Enter:

```
Install-Windowsfeature FS-iSCSITarget-Server -IncludeManagementTools
```

2. Close the **Windows PowerShell** window.

Create two iSCSI virtual disks and an iSCSI target on LON-DC1

1. On **LON-DC1**, in the **Server Manager** window, click the refresh button.
2. Create a new iSCSI virtual disk with the following settings:
 - Storage Location: **C:**
 - Name: **iSCSIDisk1**
 - Disk size: **5 GB**
 - iSCSI target: **New**
 - Target name: **FileServer**
 - Access servers: **172.16.0.11**
3. Create a new iSCSI virtual disk with the following settings:
 - Storage Location: **C:**
 - Name: **iSCSIDisk2**
 - Disk size: **5 GB**
 - iSCSI target: **fileserver**

Demonstration: Configuring iSCSI storage

In this demonstration, you will connect LON-SVR1 to the iSCSI target, and verify the presence of the iSCSI drive.

Demonstration Steps

Connect LON-SVR1 to the iSCSI target

1. On **LON-SVR1**, in **Server Manager**, open the **iSCSI** initiator.
2. Connect to the following iSCSI target:
 - Name: **LON-DC1**
 - Target name: **iqn.1991-05.com.microsoft:lon-dc1-fileserver-target**

Verify the presence of the iSCSI disks

1. In **Server Manager**, in the tree pane, click **File and Storage Services**, and then click **Disks**.
2. Notice the new two **5 GB** disks that are offline. Notice the bus type is **iSCSI**. If you are in the **File and Storage Services** section of **Server Manager**, you might need to click the refresh button to see the two disks.

Implementing high availability for iSCSI

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

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

Two technologies for implementing iSCSI for high availability are:

- **MCS.** In the event of a failure, all outstanding iSCSI commands are reassigned to another connection automatically
- **MPIO.** If you have multiple network interface cards in your iSCSI initiator and iSCSI Target Server, you can use MPIO to provide failover redundancy in the event of network outages

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

MCS is a iSCSI protocol feature that:

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

MPIO provides redundancy differently, as follows:

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

Check Your Knowledge

Question	
What are the required components of an iSCSI solution? (Select all that apply.)	
Select the correct answer.	
<input type="checkbox"/>	IP network
<input type="checkbox"/>	iSCSI targets
<input type="checkbox"/>	iSCSI initiators
<input type="checkbox"/>	iSCSI qualified name
<input type="checkbox"/>	Domain Name System (DNS)

Verify the correctness of the statement by placing a mark in the column to the right.

Statement	Answer
You can use Server Manager to configure both the iSCSI Target Server and the iSCSI initiator.	

Lab A: Implementing and managing storage

Scenario

At A. Datum Corporation, the cost of storage has decreased significantly over the last few years. However, the amount of data that business groups are producing continues to increase. Organization leaders are considering alternate ways to optimize the cost of storing data on the network and data access in the new branch offices. They also would like to ensure that the data stored on the shared folders is limited to company data, and that it does not include personal files. They also want you to test Data Deduplication to save storage. Additionally, the organization is exploring options for making storage highly available and the requirements that it must meet for high availability.

You are responsible for implementing some of the new file-storage technologies for the organization. You will implement FSRM to help optimize file storage at A. Datum. You also will implement Data Deduplication to save storage space.

You will implement iSCSI storage to provide a simpler solution for deploying storage in the organization.

Objectives

After completing this lab, you will be able to:

- Implement FSRM.
- Implement Data Deduplication.
- Configure iSCSI storage.

Lab Setup

Estimated Time: 55 minutes

Virtual machines: **20743A-LON-DC1**, **20743A-LON-SVR2**, and **20743A-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, complete the following steps:

1. On the host computer, start **Hyper-V Manager**.
2. In Microsoft Hyper-V Manager, click **20743A-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, 3, and 4 for **20743A-LON-SVR2**.
6. Repeat steps 2 and 3 for **20743A-LON-CL1**. Do not sign in to **20743A-LON-CL1** until you are instructed to do so.

Exercise 1: Implementing FSRM

Scenario

To control disk-space consumption by users, you are implementing FSRM quotas. You must limit each user's home folder to 500 MB. You also need to configure an email that will be sent to users to alert them when their home folders exceed 75 percent of their storage quota. An event also is written to the event log so that administrators can track it.

Additionally, managers are concerned that employees are storing large media files in home folders, which violates corporate policy. You need to implement file screening to prevent users from storing media files such as video, audio, and graphics files in their home folders.

The main tasks for this exercise are as follows:

1. Create a quota template.
2. Configure a quota based on the quota template.
3. Test that the quota is functional.
4. Create a file screen.
5. Create a file group.
6. Test the file screen.
7. Generate an on-demand storage report.

► Task 1: Create a quota template

1. On **LON-DC1**, in **Server Manager**, add the **File Server Resource Manager** role service.
2. Open the **File Server Resource Manager** tool.
3. In the **File Server Resource Manager** console, create a new quota template with the following values:
 - Template name: **500 MB Limit - Mail to user and log to Event Viewer**
 - Limit: **500 MB**
 - Quota type: **Hard quota. Do not allow users to exceed limit option**
4. Add a threshold with the following values:
 - Generate notifications when usage reaches (%): **75**
 - Send e-mail to the user who exceeded the threshold (E-mail message tab): **Selected**
 - Send warning to event log (Event log tab): **Selected**
5. When the **File Server Resource Manager** dialog box appears, click **Yes**.

► Task 2: Configure a quota based on the quota template

- In the **File Server Resource Manager** console, create a new quota with the following values:
 - Quota path: **E:\Labfiles\Mod02\Home**
 - Auto apply template and create quotas on existing and new subfolders: **Selected**
 - Derive properties from this quota template (Recommended): **500 MB Limit - Mail to user and log to Event Viewer**

► **Task 3: Test that the quota is functional**

1. Switch to **LON-CL1**, and then sign in as **Adatum\Abbi** with the password **Pa\$\$w0rd**.
2. On **LON-CL1**, open **File Explorer**.
3. Notice that the drive P only has 500 MB available space.

► **Task 4: Create a file screen**

1. Switch to **LON-DC1**.
2. In the **File Server Resource Manager** console tree, create a new file screen with the following values:
 - o File screen path: **E:\Labfiles\Mod02\Home**.
 - o Derive properties from this file screen template (recommended): **Block Audio and Video Files**.

► **Task 5: Create a file group**

1. On **LON-DC1**, right-click **File Server Resource Manager (Local)**, and then click **Configure Options**.
2. On the **File Screen Audit** tab, select the **Record file screening activity in auditing database** check box, and then click **OK**.
3. In **File Server Resource Manager**, create a new file group with the following values:
 - o File group name: **Adatum Media Files Group**.
 - o Files to include: ***.mp***
 - o Files to exclude ***.mpp**



Note: For the purposes of this course, only a sample of digital-media extensions is being added here. To limit all digital media files, you need to add more extensions to the file group.

4. In **File Server Resource Manager**, modify the **Block Audio and Video Files** template, and then change the template to use the **Adatum Media Files Group** file group instead of the **Audio and Video Files** file group.

► **Task 6: Test the file screen**

1. Switch to **LON-CL1**. In File Explorer, open the **P** drive.
2. Create a new text document.
3. Configure **File Explorer** to show file name extensions.
4. Rename **New Text Document.txt** to **musicfile.mp3**.
5. Click **Yes** to change the file-name extension.
6. Notice the **File Access Denied** dialog box that opens.
7. Cancel the rename operation.

► **Task 7: Generate an on-demand storage report**

1. Switch to **LON-DC1**. In the File Server Resource Manager, generate a **File Screening Audit** report for the **E:\Labfiles\Mod02\Home** folder.
2. In the **Generate Storage Reports** dialog box, verify that **Wait for reports to be generated and then display them** is selected, and then click **OK**.

3. Double-click the HTML document that starts with **FileScreenAudit**, and then review the generated html reports.
4. Close all open windows on **LON-DC1**.

Results: After completing this exercise, you should have successfully configured FSRM quotas and file screening, and generated a storage report.

Exercise 2: Implementing Data Deduplication

Scenario

You decide to install the Data Deduplication role service by using **Server Manager**. You determine that drive E is heavily used and you suspect it contains duplicate files in some folders. You decide to enable and configure the Data Deduplication role to reduce the consumed space on this volume.

The main tasks for this exercise are as follows:

1. Install the Data Deduplication role service.
2. Enable and configure Data Deduplication.
3. Test Data Deduplication.

► Task 1: Install the Data Deduplication role service

- On **LON-DC1**, in **Server Manager**, add the **Data Deduplication** role service.

► Task 2: Enable and configure Data Deduplication

1. On **LON-DC1**, execute the **E:\Labfiles\Mod02\CreateLabFiles.cmd** file.
2. Open **File Explorer** and notice that **Allfiles (E:)** has less than **50%** free space.
3. In **Server Manager**, click **File and Storage Services**.
4. Click **Disks**.
5. Click disk **1**, and then click the **E** volume.
6. Enable **Data Deduplication** for the **General purpose file server** setting.
7. Configure the following settings:
 - Deduplicate files older than (in days): **0**
 - Enable **throughput optimization**.

► Task 3: Test Data Deduplication

1. On **LON-DC1**, open **Windows PowerShell**, and then run the following commands:

```
Start-DedupJob E: -Type Optimization -Memory 50
Get-DedupJob -Volume E:
```

2. Switch to the **File Explorer** window. In File Explorer, navigate to **E:\Labfiles\Mod02\Data** and observe the following values from the **report.docx** and **song1.mp3** files' properties: **Size** and **Size on disk**.
3. Wait for five minutes to allow the deduplication job to run.
4. Switch to the **Windows PowerShell** window.

5. To verify the Data Deduplication status, run the following commands:

```
Get-DedupStatus -Volume E: | fl
Get-DedupVolume -Volume E: | fl
Get-DedupMetadata -Volume E: | fl
```

6. In **Server Manager**, click **File and Storage Services**, select Disk **1**, and then select Volume **E**.
7. Observe the values for **Deduplication Rate** and **Deduplication Savings**.
8. Close all open windows except **Server Manager**.

Results: After completing this exercise, you should have successfully installed and configured the Data Deduplication role service for the appropriate data volume on LON-DC1.

Exercise 3: Configuring iSCSI storage

Scenario

Executives at A. Datum are exploring the option of using iSCSI to decrease the cost and complexity of configuring centralized storage. Therefore, you must install and configure the iSCSI targets, and configure the iSCSI initiators to provide access to the targets.

The main tasks for this exercise are as follows:

1. Install the iSCSI target role service.
2. Configure the iSCSI targets.
3. Configure Multipath I/O (MPIO).
4. Connect to, and configure, the iSCSI targets.
5. Verify the presence of iSCSI disks.
6. Prepare for the next lab.

► Task 1: Install the iSCSI target role service

1. Switch to **LON-SVR2**, and then start **Server Manager**.
2. In **Server Manager**, add the **iSCSI Target Server** role service.

► Task 2: Configure the iSCSI targets

1. On **LON-SVR2**, in **Server Manager**, navigate to **File and Storage Services**, and then navigate to **iSCSI**.
2. Create an iSCSI virtual disk with the following values:
 - Storage location: **C:**
 - Disk name: **iSCSIDisk1**
 - Size: **5 GB**

3. Create a new iSCSI target with the following values:
 - o Target name: **lon-dc1**
 - o Two iSCSI initiators with the following IP addresses:
 - IP Address: **172.16.0.10**
 - IP Address: **10.10.0.10**
4. Repeat step 4 to create four more iSCSI virtual disks with disk names as **iSCSIDisk2**, **iSCSIDisk3**, **iSCSIDisk4**, and **iSCSIDisk5**.

► **Task 3: Configure Multipath I/O (MPIO)**

1. Switch to **LON-DC1**.
2. Open **Network Connections** and enable **Ethernet 2**.
3. Switch to **LON-SVR2** and open **Network Connections**. Enable **Ethernet 2**.
4. Switch to **LON-DC1**, and then start **Server Manager**.
5. In **Server Manager**, add the **Multipath I/O** feature.
6. Restart the server.
7. After the computer restarts, sign in as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
8. In **Server Manager**, start **iSCSI initiator**.
9. Configure the Microsoft iSCSI service to start automatically.
10. In the **iSCSI Initiator Properties** dialog box, make a quick connection to **LON-SVR2**.
11. In **Server Manager**, start **MPIO**.
12. In the **MPIO Properties** dialog box, on the **Discover Multi-Paths** tab, select the **Add support for iSCSI devices** check box, and then click **Add**. When you are prompted to restart the computer, click **Yes**.
13. After the computer restarts, sign in as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
14. In **Server Manager**, start **MPIO**.
15. In the **MPIO Properties** dialog box, on the **MPIO Devices** tab, notice that **Device Hardware ID MSFT2005iSCSIBusType_0x9** appears on the list.
16. Close the **MPIO Properties** dialog box.

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

1. On **LON-DC1**, in **Server Manager**, start **iSCSI Initiator**.
2. In the **iSCSI Initiator Properties** dialog box, disconnect from all targets.
3. In the **iSCSI Initiator Properties** dialog box, on the **Targets** tab, click **Connect**.
4. In the **Connect to Target** window, select the **Enable multi-path** check box, verify that the **Add this connection to the list of Favorite Targets** check box is selected, and then click **Advanced**.
5. In the **Advanced Settings** dialog box, on the **General** tab, configure the following:
 - o Local Adapter: **Microsoft iSCSI Initiator**
 - o Initiator IP: **172.16.0.10**
 - o Target Portal IP: **172.16.0.12 / 3260**.

6. Click **OK** twice to return to the **iSCSI Initiator Properties** dialog box.
 7. In the **iSCSI Initiator Properties** dialog box, on the **Targets** tab, click **Connect**.
 8. In the **Connect to Target** window, select the **Enable multi-path** check box, verify that the **Add this connection to the list of Favorite Targets** check box is selected, and then click **Advanced**.
 9. In the **Advanced Settings** dialog box, on the **General** tab, configure the following:
 - o Local Adapter: **Microsoft iSCSI Initiator**
 - o Initiator IP: **10.10.0.10**
 - o Target Portal IP: **10.10.0.12 / 3260**.
 10. Click **OK** twice to return to the **iSCSI Initiator Properties** dialog box.
 11. In the **iSCSI Initiator Properties** dialog box, on the **Volumes and Devices** tab, click **Auto Configure**.
 12. On the **Targets** tab, in the **Targets** list, select **iqn.1991-05.com.microsoft:lon-svr2-lon-dc1-target**, and then click **Devices**.
 13. In the **Devices** dialog box, click **MPIO**.
 14. Verify that in the Load balance policy, **Round Robin** is selected. In the **This device has the following paths** section, notice that two paths are listed. Select the first path, and then click **Details**.
 15. Note the IP address of the Source and Target portals, and then click **OK**.
 16. Select the second path, and then click **Details**.
 17. Verify that the Source IP address is the address for the second network adapter, and then click **OK**.
 18. Close all open windows except **Server Manager**.
- **Task 5: Verify the presence of iSCSI disks**
1. In **Server Manager**, click **File and Storage Services**.
 2. Click **Disks**.
 3. Notice that the five disks of iSCSI bus type are present in the list.

Results: After completing this exercise, you should have successfully installed iSCSI Target Server, configured MPIO, and connected to the iSCSI target by using iSCSI initiators.

► **Task 6: Prepare for the next lab**

- Leave the virtual machines running for the next lab in this module.

Question: Why would you implement MPIO with iSCSI? What problem does this solve?

Question: What is the purpose of the iSCSI initiator component?

Lesson 4

Configuring the Storage Spaces feature in Windows Server 2016

Managing physical disks attached directly to a server can be a tedious task. Therefore, to address this issue, and make more-efficient use of storage, many organizations have implemented SANs.

However, SANs require special configuration and hardware in some scenarios. They can be expensive, particularly for small businesses. One alternative is to use the Storage Spaces feature to provide some of the same functionality as hardware-based storage solutions. Storage Spaces is a feature in Windows Server 2016 that pools disks together and presents them to the operating system as a single disk. This lesson explains how to configure and implement Storage Spaces.

In addition, Windows Server 2016 introduces Storage Spaces Direct that enables you to create a highly available storage solution using nonshared storage.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the use of storage spaces.
- Describe the features of Storage Spaces Direct.
- Explain how to provision a storage space.
- Describe redundancy in the Storage Spaces feature.
- Describe tiering in storage.
- Describe how to monitor storage spaces.

What are storage spaces?

A storage space is a storage-virtualization capability built into Windows Server 2016 and Windows 10.


The Storage Spaces feature consists of two components:

- **Storage pools.** A storage pool is a collection of physical disks aggregated into a single logical disk so that you can manage multiple physical disks as a single disk. You can use storage spaces to add physical disks of any type and size to a storage pool.
- **Storage spaces.** These are virtual disks created from free space in a storage pool. Storage spaces have attributes such as resiliency level, storage tiers, fixed provisioning, and precise administrative control. The primary advantage of storage spaces is that you no longer need to manage single disks. Instead, you can manage them as one unit. Virtual disks are the equivalent of a logical unit number (LUN) on a SAN.

Use storage spaces to add physical disks of any type and size to a storage pool and create highly-available virtual disks from it

To create a virtual disk, you need:

- One or more physical disks
- Storage pool that includes the disks
- Virtual disks (or storage spaces) that are created with disks from the storage pool
- Disk drives that are based on virtual drives



You can manage storage by using the:

- Windows Storage Management application programming interface (API) in Windows Management Instrumentation (WMI) and Windows PowerShell.
- File and Storage Services role in **Server Manager**.

To create a highly available virtual disk, you need the following:

- Physical disks. Physical disks are disks such as SATA or serial-attached SCSI disks. If you want to add physical disks to a storage pool, the disks must satisfy the following requirements:
 - One physical disk is required to create a storage pool, and a minimum of two physical disks is required to create a resilient mirror virtual disk.
 - A minimum of three physical disks are required to create a virtual disk with resiliency through parity.
 - Three-way mirroring requires at least five physical disks.
 - Disks must be blank and unformatted. No volume can exist on disks.
 - You can attach disks by using a variety of bus interfaces, including serial attached SCSI, SATA, SCSI, Non-Volatile Memory Express (NVMe), and USB. If you want to use failover clustering with storage pools, you cannot use SATA, USB, and 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 available, nonformatted physical disks to a storage pool. You can attach a physical disk to only one storage pool. However, there can be several physical disks in a storage pool.
- Virtual disk (or storage space). This is similar to a physical disk from the perspective of users and applications. However, virtual disks are more flexible because they include both thick and thin provisioning, and just-in-time (JIT) allocations. They include resiliency to physical disk failures with built-in functionality such as mirroring and parity. These resemble Redundant Array of Independent Disks (RAID) technologies, but storage spaces store the data differently than RAID.
- Disk drive. For example, you can access this volume from your Windows operating system by using a drive letter.

You can format a storage space virtual disk with FAT32, NTFS file system, and ReFS. You have to format the virtual disk with NTFS to use it with data deduplication or with FSRM.

Windows Server 2012 introduced the Storage Spaces feature. Windows Server 2016 includes the Storage Spaces Direct feature, which is an enhancement that allows you to create a highly available storage space using local, unshared disks as direct-attached storage (DAS) or JBOD.

The new features in Windows Server 2012 R2 Storage Spaces are:

- Tiered Storage Spaces. The Tiered Storage Spaces feature allows you to use a combination of disks in a storage space, such as very fast, but small-capacity, hard disks, such as solid-state drives (SSDs), used with slower, but large-capacity hard disks. When you use this combination of disks, Storage Spaces automatically moves data that is accessed frequently to the faster hard disks, and then moves data that is accessed less frequently to the slower disks. By default, the Storage Spaces feature moves data once a day at 01:00 AM. However, you also can configure where files are stored. The advantage is that if you have files that are accessed frequently, you can pin them to the faster disk. The goal of tiering is to balance capacity against performance. Windows Server 2012 R2 recognizes only two levels of disk tiers: SSD and non-SSD.


- **Write-back caching.** The purpose of write-back caching is to optimize writing data to the disks in a storage space. Write-back caching typically works with Tiered Storage Spaces. If the server that is running the storage space detects a peak in disk-writing activity, it automatically starts writing data to the faster disks. By default, write-back caching is enabled but also limited to 1 GB of data.

The new feature in Windows Server 2016 Storage Spaces is Storage Spaces Direct, which you will learn about in the next topic.

The Storage Spaces Direct feature

Storage Spaces introduced in Windows Server 2012 might have shortcomings in big-scale installations depending on the configuration. There is a limit to how many JBODs you can interconnect to physical servers. Storage Spaces also only supported highly available hard disks through the serial attached SCSI interface. However, Storage Spaces Direct in Windows Server 2016 overcomes these limitations.

The Storage Spaces Direct feature, which can use local, unshared storage to create highly available storage for storing virtual-machine files. You can use both DAS and JBODs for Storage Spaces Direct. Additionally, with Storage Spaces Direct, you connect the JBODs to a single storage node only. This eliminates the need for a shared storage fabric, and enables you to use SATA disks to lower costs and NVMe devices to improve performance.



Configuring the Storage Spaces Direct feature by using Windows PowerShell

You use Windows PowerShell to configure Storage Spaces Direct. Both storage and failover cluster cmdlets have been improved in Windows PowerShell to handle the configuration of the Storage Spaces Direct feature. Examples of cmdlets that you can use to configure Storage Spaces Direct include:

- **Test-Cluster.** This cmdlet tests the suitability of a hardware configuration before you create a cluster.
- **Enable-ClusterStorageSpacesDirect.** This cmdlet configures a cluster for the Storage Spaces Direct feature.
- **Enable-ClusterS2D.** This cmdlet configures a cluster for the Storage Spaces Direct feature for use with NVMe devices and SATA SSDs.
- **Optimize-StoragePool.** This cmdlet rebalances storage optimization if a disk or storage node changes.
- **Debug-StorageSubsystem.** This cmdlet displays any faults that are affecting the Storage Spaces Direct feature.



Note: For more information about clusters and the Failover Clustering feature, refer to Module 11, "Implementing failover clustering" and Module 12, "Implementing failover clustering with Windows Server 2016 Hyper-V."

Configuring the Storage Spaces Direct feature by using System Center Virtual Machine Manager

Although there is no graphical user interface to configure Storage Spaces Direct in Windows Server 2016, you can use workflows in System Center Virtual Machine Manager to deploy new clusters that are Storage Spaces Direct enabled.

To do so, select the **Enable Storage Spaces Direct** check box when you run the **Create Hyper-V Cluster wizard**. The wizard performs the following:

- Installs the relevant Windows Server roles.
- Runs cluster validation.
- Installs and configures failover clustering.
- Enables storage features.

You must create the storage pool and the volumes on the cluster and then deploy the virtual machines on the cluster.

Storage Spaces Direct components

The Storage Spaces Direct feature consists primarily of known components that you put together to form a Storage Spaces Direct solution, including:

- **Network.** Storage Spaces Direct needs a network for the hosts to communicate. The network interface card must be capable of Remote Direct Memory Access (RDMA) or two network interface cards must be present to ensure performance and minimize latency.
- **Servers.** Each server or storage node has internal disks or an JBOD that connects externally. There is a requirement to have a minimum of three servers in a Storage Spaces Direct solution. Depending on the resiliency you want to achieve, you might need four servers.
- **Software Storage Bus.** Storage Spaces Direct uses SMB for intranode communication by using a new Software Storage Bus. The Software Storage Bus is the software component that combines the storage of each node so they are visible for the Storage Spaces layer.
- **Storage pools.** The storage pool uses local storage from all servers.
- **Storage spaces.** You create storage spaces, also known as virtual disks, from the storage pool. The virtual disks you create provide resiliency against both disk and server failure, because data is stored on disks on different servers.
- **ReFS.** In Windows Server 2016, ReFS is the primary file system to store virtual-machine files because of accelerated VHD/VHDX operations, which provide superior performance, compared to NTFS. ReFS also provides error detection and automatic correction.
- **CSVs.** CSVs consolidates all volumes into a single namespace that is accessible through the file system on any cluster node.
- **Scale-Out File Server.** Scale-Out File Server provides access to the storage system by using SMB 3.0. You only need Scale-Out File Server in disaggregated configurations in which the Storage Spaces Direct feature only provides storage), and is not implemented in hyper-converged configurations, in which Hyper-V runs on the same cluster as the Storage Spaces Direct feature.

The Scale-Out File Server or Hyper-V scenarios

When you use the Storage Spaces Direct feature, you decide if you want to separate the virtualization and storage layers. You can use Storage Spaces Direct in two scenarios.

You can configure a Hyper-V cluster with local storage on each Hyper-V server, and scale this solution by adding extra Hyper-V servers with extra storage. You use this implementation for small and medium businesses. This also is known as a *hyper-converged solution*.

If you want the flexibility to scale the virtualization layer independent of the storage layer, and vice versa, you can implement two clusters. One cluster is for Hyper-V and one as a Scale-Out File Server. This solution lets you add extra processing power for the virtualization layer and extra storage capacity for the storage layer. You use this for large-scale deployment, and this is known as a *disaggregated solution*.

Other uses for Storage Spaces Direct are storage of Hyper-V Replica files, or backup or archival of virtual machine files.

Storage Spaces Direct in the current implementation requires a minimum of four servers, and it supports up to 12 servers. You can add up to 240 disks to a storage pool.

Provisioning a storage space

You can create virtual disks from storage pools. To configure virtual disks or storage spaces in **Server Manager** or Windows PowerShell, you need to consider disk-sector size, drive allocation, and your provisioning scheme.

Disk-sector size

You set a storage pool's sector size when you create it. If you use only 512 and/or 512e drives, then the pool is defaulted to 512e. A 512 disk uses 512 byte sectors. A 512e drive is a hard disk with 4,096 byte sectors that emulates 512 byte sectors. If the list contains at least one 4 KB drive, then the pool sector size is 4 KB by default. Optionally, an administrator can explicitly define the sector size that all contained spaces in the pool inherit. After an administrator defines this, the Windows operating system only permits you to add drives that have a compliant sector size, that is: 512 or 512e for a 512e storage pool, and 512, 512e, or 4 KB for a 4 KB pool.

Feature	Options
Disk sector size	• 512 or 512e
Drive allocation	• Automatic • Manual • Hot Spare
Provisioning schemes	• Thin provisioning • Fixed provisioning

Drive allocation

You can configure how a pool allocates drives, and options include:

- **Automatic.** This is the default allocation when you add any drive to a pool. Storage Spaces can automatically select available capacity on data-store drives for both storage-space creation and JIT allocation.
- **Manual.** You can specify **Manual** as the usage type for drives that you add to a pool. A storage space will not use a manual drive automatically unless you select it at the creation of that storage space. This usage property makes it possible for administrators to specify that only certain storage spaces can use particular types of drives.
- **Hot Spare.** Drives that you add as hot spares to a pool are reserve drives that the storage space will not use when creating a storage space. If a failure occurs on a drive that is hosting columns of a storage space, a reserve drive replaces the failed drive.

Provisioning schemes

You can provision a virtual disk by using one of two schemes:

- **Thin provisioning space.** Thin provisioning is a mechanism that enables the Storage Spaces feature to allocate storage, as necessary. The storage pool organizes storage capacity into provisioning slabs, and does not allocate them until datasets grow to the required storage. As opposed to the traditional fixed-storage allocation method, in which you might allocate large pools of storage capacity that remain unused, thin provisioning optimizes utilization of available storage. Organizations also are able to save on operating costs, such as electricity and floor space, which are associated with keeping unused drives operating. The downside of using thin provisioning is lower disk performance.
- **Fixed provisioning space.** With Storage Spaces, fixed provisioned spaces also employ the flexible provisioning slabs. Unlike thin provisioning, in a fixed provisioning space, Storage Spaces allocates the storage capacity at the same time that you create the storage space.

Redundancy in storage spaces

If your storage pool contains more than one disk, you can also create redundant virtual disks. To configure virtual disks or storage spaces, you need to consider the redundancy functionalities of the Storage Spaces feature.

Storage layout


Depending on your redundancy requirement, you can select from three different storage layouts that might provide data availability even if one or more disks fail. Valid options include:

- **Simple**
 - * No resiliency
- **Two-way mirror**
 - * Keeps functioning even with one failed disk
- **Three-way mirror**
 - * Keeps functioning even with two failed disks
- **Parity**
 - * Keeps functioning even with one failed disk
- **Clustered storage spaces**
 - * Keeps functioning in case of both server and disk failure

- **Simple.** A simple space has data striping but no redundancy. In data striping, the disks segment logically sequential data so that access to these sequential segments are to different physical storage drives. Striping makes it possible to access multiple segments of data concurrently. Do not host important data on a simple volume, because it provides no failover capabilities when the disk that is storing the data fails.
- **Two-way and three-way mirrors.** Mirror spaces maintain two or three copies of the data that they host. There are two data copies for two-way mirrors and three data copies for three-way mirrors. Data duplication happens with every write to ensure that all data copies remain current. Mirror spaces also stripe the data across multiple physical drives. Mirror spaces provide the benefit of greater data throughput and lower access latency. They also do not introduce a risk of corrupting at-rest data, and do not require the extra journaling stage when writing data.
- **Parity.** A parity space is similar to RAID 5. Storage Spaces stripe data, along with parity information, across multiple physical drives. Parity enables Storage Spaces to continue to service read and write requests even when a drive has failed. Parity rotates across available disks to enable I/O optimization. Storage spaces require a minimum of three physical drives for parity spaces. Parity spaces have increased resiliency through journaling.

Clustered storage spaces

When you use the Storage Spaces feature with the Failover Clustering feature, you provide a resilient and highly available solution. You configure a clustered storage spaces deployment with a small number of servers, and a set of shared serial-attached SCSI JBOD enclosures. The JBOD enclosures should have connections to all the servers, and each server should have redundant paths to all the disks in each JBOD enclosure. This limits the size of the deployment. For more about Failover Clustering, refer to Module 11, "Implementing failover clustering."

 **Additional Reading:** For more information, refer to: "Deploy Clustered Storage Spaces" at: <http://aka.ms/k5vwda>

Tiering storage

When you buy storage, you have to decide whether you want capacity, in the form of regular hard disk drives (HDD), or performance disks (or SSDs). The Storage Spaces feature creates tiers of storage, which provides both capacity and performance.

You create a tiered storage space by adding both SSDs and HDDs to a storage space. Windows Server 2016 can register which type each hard disk is and automatically create two tiers.

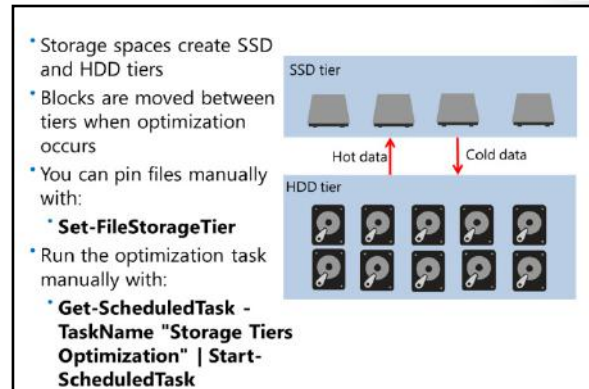
Optimization of the storage tiers optimization is a two-part process. During daily operations, Storage Spaces creates a map of how often users and application access each data block. Each night at 1:00 AM, by default, the Storage Tiers Optimization task runs. It moves the most frequently accessed (hot) data blocks to the SSD tier and the less frequently accessed (cold) data blocks to the HDD tier. The optimization happens on the block level, not file level. Storage Spaces could move a part of a VHDX file for a virtual machine to the SSD tier, if that part has more disk I/O requests than the rest of the file.

To run the optimization task manually, you can run the following command in a Windows PowerShell window:

```
Get-ScheduledTask -TaskName "Storage Tiers Optimization" | Start-ScheduledTask
```

If you want to create a report every time the optimization task runs, you should add the following arguments to the scheduled task: `/c /h /g /# >> C:\Windows\StorageReports\StorageTiersOutput.txt`.

You should not allocate all your SSD capacity for your storage spaces. That way, you can increase the size of an SSD tier when you need to. Although you can pin files to a specific storage tier, wait until you see how well Storage Tiers Optimization can optimize storage performance. When you pin files, you do it on a file-by-file basis by using the **Set-FileStorageTier** Windows PowerShell cmdlet. If you are using VDI, you should consider pinning the image that VDI uses to clone users' desktops to the SSD tier for performance reasons.



Monitoring Storage Spaces

It is important to monitor the health of your storage spaces to ensure that you have a healthy and running storage spaces environment. You can use the following tools to monitor your storage spaces:

- **Server Manager**
- **Windows PowerShell**
- **Event Viewer**
- **System Center Operations Manager**

- **Server Manager:**
 - Displays warning signs, which indicate degraded health or performance
- **Windows PowerShell:**
 - Use the **Get-StoragePool**, **Get-VirtualDisk**, and **Get-PhysicalDisk** cmdlets
 - Check the **OperationalStatus** and **HealthStatus** properties
- **Event Viewer:**
 - System log event IDs 100–104, 200–203, 300–308
 - Check the Storage Spaces Events event log
- **System Center Operations Manager Management Pack**

Server Manager

You can use the **Storage Pools** node in **Server Manager** to view the status of your storage pools, and virtual and physical disks. A warning sign appears if storage spaces' health is not okay.

Windows PowerShell

You can use the Windows PowerShell cmdlets in the Storage module to retrieve information regarding storage pool, physical and virtual disks, enclosures, and storage tiers. Most of the objects have both operational status and health status attributes that you can query to monitor the health of your storage spaces infrastructure.

For example, you can use the following commands in Windows PowerShell to query the status of your storage spaces:

```
Get-StoragePool -FriendlyName Pool1
Get-StoragePool -FriendlyName Pool1 | Get-VirtualDisk
Get-PhysicalDisk | Where-Object CannotPoolReason -match 'In a pool'
```

Event Viewer

You can also monitor the event logs for events regarding storage spaces. Storage spaces generate events in the System log. All events have **StorageSpaces-Drivers** as the event source. The following table lists the relevant events.

Event ID	Event level	Description
100	Error	Failed to read the storage pool configuration.
102	Error	Majority of the physical drives of storage pool failed a configuration update.
103	Error	The capacity consumption of the storage pool has exceeded the threshold limit set on the pool.
104	Information	The capacity consumption of the storage pool is now below the threshold limit set on the pool.
200	Error	Windows was unable to read the drive header for physical drive.
201	Error	Physical drive has invalid meta-data.
202	Error	Physical drive has invalid meta-data.

Event ID	Event level	Description
203	Error	An IO failure has occurred on Physical drive
300	Error	Physical drive failed to read the configuration or returned corrupt data for storage space.
301	Error	All pool drives failed to read the configuration or returned corrupt data for storage space.
302	Error	Majority of the pool drives hosting space meta-data for storage space failed a space meta-data update.
303	Error	Drives hosting data for storage space have failed or are missing.
304	Warning	One or more drives hosting data for storage space have failed or are missing.
306	Error	The attempt to map, or allocate more storage for, the storage space has failed.
307	Error	The attempt to unmap or trim the storage space has failed.
308	Information	A repair attempt for storage space was initiated by the driver.

The **New-StorageSpacesEventLog** Windows PowerShell cmdlet creates an event log named **Storage Spaces Events** under **Application and Services** in **Event Viewer**. It also creates a scheduled task to monitor for events and logs them in the new event log. You can remove the event log again if you run the **Remove-StorageSpacesEventLog** cmdlet.

System Center Operations Manager Management Pack

You can use **System Center Operations Manager** to monitor the health of your storage spaces by installing and configuring a management pack. The management pack requires that you also have System Center Virtual Machine Manager installed. The management pack can monitor up to:

- 16 storage nodes.
- 12 storage pools.
- 120 file shares.



Additional Reading: To download the Microsoft System Center Operations Manager Management Pack for Windows Server Storage Spaces, refer to: <http://aka.ms/Uzb16z>

Check Your Knowledge

Question	
Which resiliency types can you configure in the Storage Spaces feature? (Select all that apply.)	
Select the correct answer.	
<input type="checkbox"/>	Simple
<input type="checkbox"/>	Advanced
<input type="checkbox"/>	Two-way mirror
<input type="checkbox"/>	Three-way mirror
<input type="checkbox"/>	Parity

Lab B: Implementing and managing advanced storage solutions

Scenario

At A. Datum, you will implement the Storage Spaces feature on the Windows Server 2016 servers to simplify storage access and provide redundancy at the storage level.

You also want to test the feasibility of using highly available storage. You decide to test Storage Spaces Direct.

Objectives

After completing this lab, you will be able to:

- Configure redundant storage spaces.
- Implement Storage Spaces Direct.

Lab Setup

Estimated Time: 50 minutes

Virtual machines: **20743A-LON-DC1, 20743A-LON-SVR1, 20743A-LON-SVR2, 20743A-LON-SVR3**

User name: **Adatum\Administrator**

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

For Exercise 1 of this lab, you need to use the available virtual machine environment. The virtual machines are running from the last lab. After Exercise 1, you will revert all virtual machines and then for Exercise 2, start the machines that are specifically required for the exercise.

Exercise 1: Configuring redundant storage spaces

Scenario

To meet some of the requirements for high availability, you decided to evaluate redundancy options in the Storage Spaces feature. Additionally, you want to test the provisioning of new disks to the storage pool.

The main tasks for this exercise are as follows:

1. Create a storage pool by using the iSCSI disks attached to the server.
2. Create a three-way mirrored disk.
3. Copy a file to the volume, and verify visibility in Windows Explorer.
4. Disconnect an iSCSI disk.
5. Verify that the file still is accessible, and check the virtual disk's health.
6. Add a new iSCSI virtual disk.
7. Add the new disk to the storage pool, and extend the virtual disk.
8. Prepare for the next exercise.
9. Prepare for the next module.

► **Task 1: Create a storage pool by using the iSCSI disks attached to the server**

1. On **LON-DC1**, in **Server Manager**, click **Storage Pools**.
2. Create a new storage pool with the following values:
 - Storage pool name: **StoragePool1**
 - Available disks group: **LON-DC1 – Primordial**
 - Physical disks: All five physical disks

► **Task 2: Create a three-way mirrored disk**

1. In **Server Manager**, select **StoragePool1**.
2. Create a new virtual disk with the following values:
 - Storage pool: **StoragePool1**
 - Virtual disk name: **Mirrored vDisk**
 - Enable enclosure awareness: Not selected
 - Storage layout: **Mirror**
 - Resiliency settings: **Three-way mirror**
 - Provisioning type: **Thin**
 - Size of the virtual disk: **10**
3. After creation of the virtual disk, ensure that **Create a volume when this wizard closes** is selected, and then click **Close**.
4. In the **New Volume Wizard** window, create a volume with the following values:
 - Virtual disk: **Mirrored vDisk**
 - Volume size: default value
 - Drive letter: **F**
 - File system: **ReFS**
 - Volume label: **Mirrored Volume**

► **Task 3: Copy a file to the volume, and verify visibility in Windows Explorer**

1. Open a **Windows PowerShell** window.
2. At the Windows PowerShell prompt, type the following command, and then press Enter:

```
Copy C:\windows\system32\write.exe F:\
```

3. Open **File Explorer**, and then navigate to **Mirrored Volume (F:)**. You should now see Write.exe in the file list.
4. Close **File Explorer**.

► **Task 4: Disconnect an iSCSI disk**

1. Switch to **LON-SVR2**.
2. Open **Server Manager**, and in the **iSCSI VIRTUAL DISKS** pane, right-click **iSCSIDisk1.vhd**, and then click **Disable iSCSI Virtual Disk**.
3. Accept the warning message box.

► **Task 5: Verify that the file still is accessible, and check the virtual disk's health**

1. Switch to **LON-DC1**.
2. Open **File Explorer**, and then navigate to **Mirrored Volume (F:)**.
3. Double-click **write.exe** to ensure access to the volume still is available.
4. Close the **Document - WordPad** window.
5. Close **File Explorer**.
6. In **Server Manager**, on the menu bar, click the **Refresh** button. Wait until all panes are refreshed. Notice the warning that appears next to **Mirrored vDisk**. The result may vary slightly. A warning may also appear in the **physical disks** section. If the status for **StoragePool1** does not change:
 - a. Restart **LON-DC1**.
 - b. Sign in as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
 - c. In **Server Manager**, click **File and Storage Services**, and then click **Storage Pools**.
7. In the **VIRTUAL DISKS** pane, right-click **Mirrored vDisk**, and then in the drop-down list, select **Properties**.
8. In the **Mirrored vDisk Properties** window, click **Health**.
9. Click **OK**.

► **Task 6: Add a new iSCSI virtual disk**

1. Switch to **LON-SVR2**.
2. In **Server Manager**, navigate to **iSCSI**.
3. Create a **New iSCSI Virtual Disk** with the following values:
 - Storage location: **C:**
 - Disk name: **iSCSIDisk6**
 - Size: **5 GB**
 - iSCSI target: **lon-dc1**

► **Task 7: Add the new disk to the storage pool, and extend the virtual disk**

1. Switch to **LON-DC1**.
2. In **Server Manager**, click the **Refresh** button.
3. Wait for all of the panes to refresh.
4. In the **STORAGE POOLS** pane, right-click **StoragePool1**, and then add the new physical disk to the storage pool.
5. In the **VIRTUAL DISKS** pane, right-click **Mirrored vDisk**, and then extend the virtual disk to 15 GB.

Results: After completing this exercise, you should have successfully created a storage pool, added a new disk to the storage pool, and extended the disk.

► Task 8: Prepare for the next exercise

When you finish Exercise 1, revert the virtual machines to their initial state. To do this, perform the following steps:

1. On the host computer, start **Hyper-V Manager**.
2. In the **Virtual Machines** list, right-click **20743A-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20743A-LON-SVR2** and **20743A-LON-CL1**.

Exercise 2: Implementing the Storage Spaces Direct feature

Scenario

You want to test if using local storage as highly available storage is a viable solution for your organization. Previously, your organization only has used SANs for storing virtual machines. The new features in Windows Server 2016 make it possible to use only local storage, so you will implement Storage Spaces Direct as a test implementation.

You will use a Windows PowerShell script in the exercise to create the Storage Spaces Direct cluster and configure Storage Spaces Direct. The script will perform the following steps:

1. Install the Remote Server Administration Tools to enable configuration from one server:

```
Install-WindowsFeature RSAT -IncludeAllSubFeature
```

2. Install the required roles and features on the servers that will be a part of the Storage Spaces Direct cluster:

```
Install-WindowsFeature -Name File-Services, Failover-Clustering -  
IncludeManagementTools
```

3. Validate if you can use the hardware configuration for Storage Spaces Direct:

```
Test-Cluster-Node <computer names for the nodes in the cluster> -Include "Storage  
Spaces Direct"
```

4. Create the cluster, but not add any storage:

```
New-Cluster -Name <Cluster name> -Node <computer names for the nodes in the cluster> -  
NoStorage -StaticAddress <IP address>
```

5. Configure the cluster properties to enable the Software Storage Bus and add storage:

```
Enable-ClusterStorageSpacesDirect -Cluster S2DCluster
```

6. Create the storage pool in the cluster:

```
New-StoragePool -StorageSubSystemName <Cluster name> -FriendlyName <Storage Pool  
name>
```

7. Create the volume and add it to the CSVs:

```
New-Volume -StoragePoolFriendlyName <Storage Pool name> -FriendlyName <Friendly name>  
-FileSystem CSVFS_ReFS -Size <Size of volume>
```

8. Create the Scale-Out File Server on the cluster:

```
New-StorageFileServer -StorageSubSystemName <Cluster name> -FriendlyName <Name of file server in cluster> -HostName <Virtual host name> -Protocols SMB
```

9. Create a folder on the file server and create a new share:

```
md "C:\ClusterStorage\Volume1\<Folder name>"
New-SmbShare -Name <Share name>-Path "C:\ClusterStorage\Volume1\<Folder name>" -
FullAccess <Users with Full Control>
Set-SmbPathAcl -ShareName <Share name>
```

The main tasks for this exercise are as follows:

1. Install the Windows Server roles and features.
2. Validate cluster configuration.
3. Create a cluster.
4. Enable the Storage Spaces Direct feature.
5. Create a storage pool.
6. Create a virtual disk.
7. Create a file server and file share.
8. Test high availability for the storage.

► Task 1: Install the Windows Server roles and features

1. On the host computer, start **Hyper-V Manager**.
2. In Hyper-V Manager, click **20743A-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 and 3 for **20743A-LON-SVR1**, **20743A-LON-SVR2**, and **20743A-LON-SVR3**.
6. Repeat step 4 for **LON-SVR1**.
7. Switch to **LON-DC1**.
8. In **Server Manager**, click **All Servers**.
9. Add **LON-SVR1**, **LON-SVR2**, and **LON-SVR3** to **All Servers**.
10. Verify that all three servers have a **Manageability** of **Online – Performance counters not started** before continuing.
11. Open **Windows PowerShell ISE**.
12. Open the **Implement-StorageSpacesDirect.ps1** file located at **E:\Labfiles\Mod02**.
13. Select the first line in step 0 starting with **Install-Windowsfeature**, and then press F8. Wait until the installation finishes.

14. If the column **Restart Needed** has a value of **Yes**, select the line starting with **Restart-Computer**, and then press F8. After the server restarts, sign in as **Adatum\Administrator** with the password **Pa\$\$w0rd**. Repeat steps 11 and 12. If a dialog box opens, click **OK**.
15. Select the line in step 1 starting with **Invoke-Command**, and then press F8. Wait until the installation finishes.
16. Verify that the output of the command includes four lines with **Success** as **True**.
17. Switch to **LON-SVR1**. In **Server Manager**, open **Failover Cluster Manager**.
18. Notice that the tool opens. This verifies that the command was successful.

► **Task 2: Validate cluster configuration**

1. Switch to **LON-DC1**.
2. In the **Administrator: Windows PowerShell ISE** window, select the line in step 2, starting with **Test-Cluster**, and then press F8. Wait until the installation finishes.
3. Verify that the output of the command only includes warnings, and that the last line is a validation report in html format. This validates that the command was successful.

► **Task 3: Create a cluster**

1. In the **Administrator: Windows PowerShell ISE** window, select the line in step 3, starting with **New-Cluster**, and then press F8. Wait until the installation finishes.
2. Verify that the output of the command only includes warnings, and that the last line has a **Name** column with the value **S2DCluster**. This validates that the command was successful.
3. Switch to **LON-SVR1**, and then in the **Failover Cluster Manager** window, connect to the **S2DCluster**.
4. Notice that the cluster opens. This verifies that the command was successful. If you receive an error message stating **Cluster S2DCluster not found**, perform the following steps:
 - a. Restart **LON-SVR1**.
 - b. Sign in as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
 - c. Click **Start**, and then click **Server Manager**.
 - d. In **Server Manager**, click **Tools**, and then click **Failover Cluster Manager**.
 - e. **S2DCluster.Adatum.com** should now appear in the navigation pane. If necessary, perform steps 3 and 4.

► **Task 4: Enable the Storage Spaces Direct feature**

1. Switch to **LON-DC1**.
2. In the **Administrator: Windows PowerShell ISE** window, select the line in step 4, starting with **Invoke-Command**, and then press F8. Wait until the installation finishes.
3. In the **Confirm** dialog box, click **Yes**.
4. Verify that there is no output of the command. This validates that the command was successful.

► **Task 5: Create a storage pool**

1. In the **Administrator: Windows PowerShell ISE** window, select the line in step 5, starting with **Invoke-Command**, and then press F8. Wait until the installation finishes.
2. In the output of the command, verify that the **FriendlyName** attribute has a value of **S2DStoragePool**. This validates that the command was successful.
3. Switch to **LON-SVR1**, and then in the **Failover Cluster Manager** window, verify the existence of **Cluster Pool 1**. This verifies that the command was successful.

► **Task 6: Create a virtual disk**

1. Switch to **LON-DC1**. In the **Administrator: Windows PowerShell ISE** window, select the line in step 6, starting with **Invoke-Command**, and then press F8. Wait until the installation finishes.
2. Verify that in the output of the command is an attribute **FileSystemLabel**, with a value of **CSV**. This validates that the command was successful.
3. Switch to **LON-SVR1**, and then in the **Failover Cluster Manager**, verify the existence of **Cluster Virtual Disk (CSV)**. This verifies that the command was successful.

► **Task 7: Create a file server and file share**

1. Switch to **LON-DC1**. In the **Administrator: Windows PowerShell ISE** window, select the line in step 7, starting with **Invoke-Command**, and then press F8. Wait until the installation finishes.
2. Verify that in the output of the command is an attribute **FriendlyName**, with a value of **S2D-SOFS**. This validates that the command was successful.
3. Switch to **LON-SVR1**, and then in the **Failover Cluster Manager** window, verify the existence of the **S2D-SOFS** role. This verifies that the command was successful.
4. Switch to **LON-DC1**. In the **Administrator: Windows PowerShell ISE** window, select the three lines in step 8, starting with **Invoke-Command**, and then press F8. Wait until the installation finishes.
5. Verify that there in the output of the command is an attribute **Path** with a value of **C:\ClusterStorage\Volume1\VM01**. This validates that the command was successful.
6. Switch to **LON-SVR1**, and then in the **Failover Cluster Manager** window, verify the existence of the **VM01** share. This verifies that the command was successful.

► **Task 8: Test high availability for the storage**

1. Switch to **LON-DC1**.
2. Open a **File Explorer** window and navigate to **\\s2d-sofs\VM01**.
3. Create a new folder named **VMFolder**.
4. Shut down **LON-SVR2** by running the following Windows PowerShell cmdlet:

```
Stop-Computer -computername LON-SVR2
```

5. In **Server Manager**, verify that **LON-SVR2** is no longer accessible.
6. Switch to the **File Explorer** window and create a new text document in the **VMFolder** folder. This verifies that the storage is still available with one server turned off.

7. Switch to **LON-SVR1**. In **Failover Cluster Manager**, click **Disks** and then click **Cluster Virtual Disk (CSV)**. Verify that for the **Cluster Virtual Disk (CSV)**, the **Health Status** is **Warning** and **Operational Status** is **Degraded**. **Operational Status** may also display as **Incomplete**.

Results: After completing this exercise, you should have implemented Storage Spaces Direct successfully.

► **Task 9: Prepare for the next module**

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

1. On the host computer, start **Hyper-V Manager**.
2. In the **Virtual Machines** list, right-click **20743A-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20743A-LON-SVR1**, **20743A-LON-SVR2**, and **20743A-LON-SVR3**.

Question: How many disk failures can you have while remaining functional with a three-way mirrored storage space?

Question: How many servers do you need as a minimum to set up the Storage Spaces Direct feature?

Module Review and Takeaways

Review Questions

Question: Is the Storage Spaces feature also available in Windows 10?

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

Question: Can you use both local and shared storage with the Storage Spaces Direct feature?

MCT USE ONLY. STUDENT USE PROHIBITED

Module 3

Implementing Directory Services

Contents:

Module Overview	3-1
Lesson 1: Deploying Active Directory domain controllers	3-2
Lesson 2: Implementing service accounts	3-14
Lab: Implementing and managing AD DS	3-18
Lesson 3: Azure AD	3-22
Module Review and Takeaways	3-31

Module Overview

Active Directory Domain Services (AD DS) is the central location for configuration information, authentication requests, and information about all of the objects that are stored in your Active Directory forest. Using AD DS, you can efficiently manage users, computers, groups, printers, and other directory-enabled objects from one secure, central location. Providing authentication for cloud-based applications and services is becoming increasingly important for many organizations. The ability to integrate and sync a local AD DS environment with cloud-based authentication services is also becoming more common. This module discusses deployment and configuration of domain controllers, service accounts in AD DS, and integration with Microsoft Azure Active Directory (Azure AD).

Objectives

After completing this module, you will be able to:

- Deploy Windows Server 2016 Active Directory domain controllers.
- Implement service accounts in AD DS.
- Integrate on-premises AD DS with Azure AD.

Lesson 1

Deploying Active Directory domain controllers

To establish the Active Directory forest and the first domain in the forest, you must deploy at least one domain controller. In this lesson, you will learn about the new features of AD DS on Windows Server 2016 and the various methods for deploying domain controllers.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the new Active Directory features on Windows Server 2016.
- Explain how to deploy Windows Server 2016 domain controllers.
- Explain how to deploy Active Directory domain controllers on Server Core.
- Explain how to use the **Install from Media** option to deploy a domain controller.
- Describe Active Directory read-only domain controllers (RODCs).
- Deploy an RODC.
- Explain how to clone virtual domain controllers.
- Explain how to upgrade an Active Directory forest to Windows Server 2016.

What's new in AD DS on Windows Server 2016?

Windows Server 2016 has several new features for AD DS. Some of these features and improvements were first introduced in Windows Server 2012 or Windows Server 2012 R2, while others are new to Windows Server 2016.

New features introduced in Windows Server 2012 or Windows Server 2012 R2

The following features in AD DS on Windows Server 2016 were first introduced in Windows Server 2012 or Windows Server 2012 R2:

- Improved support for running virtualized domain controllers. Improvements in the virtual environment include:
 - Cloning domain controllers. Cloning domain controllers is now a supported option, which enables automated deployment and rollback protection.
 - Domain controller restoration. Restoration of a domain controller checkpoint does not disrupt the Active Directory environment.
- Multi-factor authentication. When using Active Directory Federation Services (AD FS), you can grant authentication based on multiple factors. The first factors are the Active Directory credentials, while other possible factors include the device and whether it is workplace joined.

- New features introduced in Windows Server 2012 or Windows Server 2012 R2:
 - Improved support for running virtualized domain controllers
 - Multi-Factor Authentication
 - Active Directory-based activation
- New features and improvements in Windows Server 2016:
 - PAM
 - Group member expiration
 - Microsoft Passport
 - Azure AD Connect




Note: *Workplace Join* is a Windows Server 2012 R2 feature that enables users to join a personal computing device to your Active Directory domain in order to gain access to apps and data. In Windows Server 2016, this is referred to as *device registration*.

- Active Directory-based activation. Key Management Servers (KMSs) are no longer required to activate computers that run Windows Server 2012 and Windows 8 or later Windows operating systems. Activating the initial customer-specific volume license key (CSVLK) requires one-time contact with Microsoft activation over the Internet.


New features and improvements introduced in Windows Server 2016

The following features and improvements were introduced in AD DS for Windows Server 2016 since Windows Server 2012 R2:


- Privileged Access Management (PAM). The PAM feature helps to counter security issues arising from credential theft, and it helps you restrict privileged access within your existing Active Directory environment. Enabling PAM includes establishing a shadow bastion forest that is known to be secure and does not contain artifacts from earlier attempts at privileged access and credential theft.

 **Additional Reading:** For more information on PAM for AD DS, refer to: <http://aka.ms/dav4pu>

- Group member expiration. This feature enables you to configure automatic expiration for group membership. This is useful when you have temporary staff such as interns or when you have a short-term project within your organization.

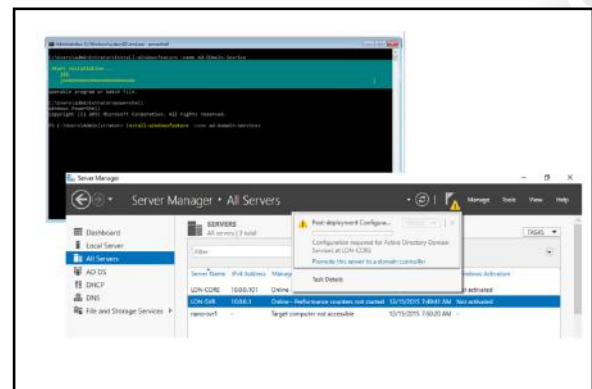
 **Note:** Group member expiration requires that you raise your Active Directory functional level to a Windows Server 2016 functional level. Group membership expiration also requires Microsoft Identity Manager (MIM) 2016.

- Microsoft Passport. Windows Server 2016 supports Microsoft Passport, a key-based authentication system. Users sign in by using a PIN or biometric information that is linked to a digital certificate or an asymmetrical key pair. This provides a more secure sign-in experience.
- Azure AD Connect. This feature enables you to integrate on-premises AD DS with Azure AD, thereby enabling users to connect to **Microsoft Office 365** and Azure services, which are part of your organization's subscription.

 **Note:** You can use Azure AD Connect with Windows Server 2012 R2 deployments of Active Directory.

Deploying domain controllers on Windows Server 2016

With Windows Server 2008, you install the Active Directory role to add the binary files, and then you use the **Active Directory Installation Wizard** to install AD DS. In Windows Server 2012 and Windows Server 2016, you deploy a domain controller by first adding the Active Directory role. You then use a separate wizard to configure AD DS within **Server Manager**.



Use the following methods to add the Active Directory role binaries:

- **Add Roles and Features Wizard** in **Server Manager**
- **Windows PowerShell**
- **Dism.exe**
- **Active Directory Domain Services Installation Wizard**

Using Server Manager

You can use the **Add Roles and Features Wizard** in **Server Manager** to install the binary files and to perform all the required configuration of a new domain controller. The wizard uses a single expanding dialog box, and it can:

- Install AD DS remotely.
- Install Domain Name System (DNS) by default.
- Configure the domain controller as a global catalog by default.
- Display advanced mode settings.
- Prepare schema extension and domain preparation automatically in the background.



Note: These new features are not backward compatible with Windows Server 2008 R2 or older Windows Server versions.

Using Windows PowerShell

Use the Windows PowerShell **Install-WindowsFeature** cmdlet to add Active Directory binaries. After you add the binaries, you can use the **Install-ADDSDomainController** cmdlet to complete your deployment of a new domain controller.

Using DISM

Using the **Deployment Image Servicing and Management** (DISM) tool is more complex and less flexible when adding the binaries than it is with Windows PowerShell. DISM is usually associated with creating deployment images, for example, with Windows Deployment Services.

Using the Active Directory Installation Wizard

The **Active Directory Installation Wizard** no longer has a GUI and is only supported with the **Unattend** option. **DCPromo** is primarily used when deploying a domain controller in a Server Core installation of Windows Server.

Deploying domain controllers on Server Core

You must install AD DS on a Server Core installation by using Windows PowerShell—either locally or remotely. Alternatively, you can use **Server Manager** on a remote system to perform the installation.

Installing the Active Directory role locally

To install the Active Directory role locally, use the following procedure:

1. Install the Active Directory binary files. At the local Windows PowerShell command prompt, type the following command, and then press Enter:

```
Install-WindowsFeature -name AD-Domain-Services
```

2. Configure AD DS. At the Windows PowerShell command prompt, type the following command with other arguments as required, and then press Enter:

```
Install-ADDSDomainController -domainname "Adatum.com"
```

For example, the following cmdlet prompts the user for a valid user name and password to perform the promotion, and it then deploys the Active Directory domain controller role in the **Adatum.com** domain to the local server:

```
Install-ADDSDomainController -Credential (Get-Credential) -DomainName "Adatum.com"
```

Windows PowerShell remote installation

You can run Windows PowerShell commands against remote servers. First, install the Active Directory binary files. Then use the **Invoke-Command** cmdlet. For example:

```
Invoke-Command {Install-ADDSDomainController -DomainName Adatum.com -Credential (Get-Credential) -Computername LON-DC3}
```

Server Manager remote installation

To use **Server Manager** to install the Active Directory role remotely on Server Core, perform these high-level steps:

1. Add the computer with the Server Core installation as another computer to manage in **Server Manager**.
2. Use the **Add Roles and Features Wizard** to install AD DS.
3. Complete the configuration by running the **Active Directory Domain Services Configuration Wizard**.

Deploying domain controllers by using the Install from Media option

Another method for installing AD DS is to install from installation media that you create by using the **Ntdsutil.exe** tool. Installation media is created from an existing domain controller in the form of a backup. The advantage of installing from media is that it reduces the directory replication traffic that is necessary to synchronize the new domain controller. By default, a new domain controller replicates all the data for all of the directory partitions that it hosts from other domain controllers. When you use the **Install from Media** option, the new domain controller has most of the Active Directory data. It only replicates updates that have occurred since the backup media was created.



Note: Installation from media does not work across different operating system versions. For example, you must generate media from an existing Windows Server 2016 domain controller to install AD DS on a computer that is running Windows Server 2016.

Deploying Active Directory domain controllers on Server Core


A Windows Server 2016 server that is running Server Core does not have the Server Manager GUI interface, so you need to use alternative methods to install the files for the domain controller role and to install the domain controller role itself. You can use **Server Manager**, **Windows PowerShell**, or **Remote Server Administration Tools (RSAT)** installed on a client computer that has the Windows 8.1 operating system or a later version.

To install the Active Directory files on a server, you can do one of the following:

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

After you install the Active Directory files, you can complete everything except the hardware installation and configuration in one of the following ways:

- Use **Server Manager** to start the **Active Directory Domain Services Configuration Wizard** as described in the previous topic.
- Run the **Install-ADDSDomainController** Windows PowerShell cmdlet and supply the required information on the command line.

 **Note:** In Windows Server 2016, running a cmdlet loads the cmdlet's module automatically, if it is available. For example, running the **Install-ADDSDomainController** cmdlet loads the **ADDSDeployment** module automatically into your current Windows PowerShell session. If a module is not loaded or available, you will receive an error when you run the cmdlet. The error message reads that it is not a valid cmdlet.

You can still import the module that you need manually. However, you do not need to do this in Windows Server 2016 unless there is an explicit need to do so, such as pointing to a particular source to install the module.

Additional Reading:

- For more information on using the **Install-ADDSDomainController** Windows PowerShell cmdlet, refer to: <http://aka.ms/mvkc3u>
- For more information, refer to: <http://aka.ms/WivzdV>

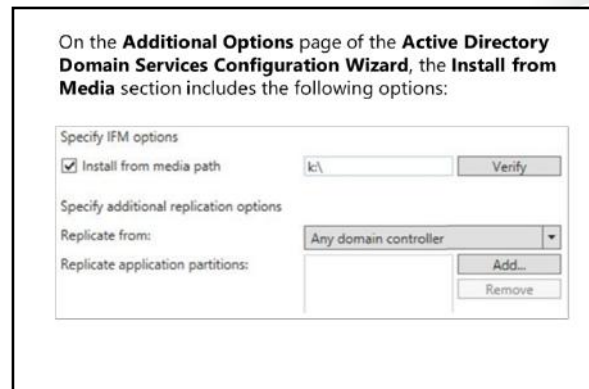
Installing AD DS is a two-step process regardless of which installation method you use:

- Method 1: Use Server Manager to connect to the target server:
 1. Install the files by installing the Active Directory role
 2. Install the domain controller role by running the **Active Directory Domain Services Configuration Wizard**
- Method 2: Use Windows PowerShell:
 1. Install the files by running the **Install-WindowsFeature AD-Domain-Services** cmdlet
 2. Install the domain controller role by running the **Install-ADDSDomainController** cmdlet

Deploying Active Directory domain controllers by using the Install from Media method

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

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



To install a domain controller by using **Install from Media**, browse to a domain controller that is not an RODC. Use the **ntdsutil** command-line tool to create a snapshot of the Active Directory database, and then copy the snapshot to the server that you will promote to a domain controller. Use **Server Manager** to promote the server to a domain controller by selecting the **Install from Media** option, and then provide the local path to the **Install from Media** directory that you created previously.

The procedure is as follows:

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

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

2. On the server that you are promoting to a domain controller, perform the following steps:

- a. Add the Active Directory role by using **Server Manager**.
- b. Wait while the Active Directory files install.
- c. In **Server Manager**, click the **Notification** icon, and then under **Post-Deployment Configuration**, click **Promote this server to a domain controller**.

The **Active Directory Domain Services Configuration Wizard** runs.

- d. On the appropriate page of the wizard, select the option to **Install from Media**, and then provide the local path to the snapshot directory.

AD DS then installs from the snapshot. When the domain controller restarts, it contacts other domain controllers in the domain and updates AD DS with any changes that occurred after the snapshot's creation.



Additional Reading: For more information on the steps that are necessary to install AD DS, refer to: "Install Active Directory Domain Services (Level 100)" at: <http://aka.ms/nmus1d>

Active Directory read-only domain controllers

An RODC hosts read-only partitions of an Active Directory database. This means that Active Directory change requests are not made directly to the database copy that the RODC stores. Instead, Active Directory modifications are forwarded to RODCs through replication with a writable domain controller. All RODC Active Directory replication uses a one-way, incoming-only connection from a domain controller that has a writable Active Directory database copy.

RODCs provide:

- Unidirectional replication
- Credential caching
- Administrative role separation
- Read-only DNS
- The RODC filtered attribute set

By design, RODCs are primarily for branch office deployments where you cannot guarantee the physical security of domain controllers. By deploying an RODC in a branch office, you can give users a local domain controller to facilitate efficient Active Directory sign in and Group Policy application even if the WAN link to the main office (where read/write domain controllers are located) is not available. A locally based RODC that is configured to cache local user passwords ensures faster sign-ins compared to signing in across a slow network connection to authenticate with a remote domain controller.

Characteristics of RODCs

RODCs have the following characteristics:

- Server Core installations support RODCs.
- An RODC cannot hold an operations master role.
- An RODC cannot be a site bridgehead server.
- RODCs only support incoming replication.
- You can explicitly enable or deny the caching of user and computer credentials in the **Active Directory Configuration Wizard**. By default, user credentials do not cache.
- Users can be delegated administrative rights to a specific RODC without being granted administrative rights to AD DS. You can configure this in the **Active Directory Configuration Wizard**.
- RODCs support read-only DNS.
- RODCs can use the **Install from Media** feature for initial deployment.
- The RODC filtered attribute set is a set of attributes that contain sensitive data. The attributes are marked as confidential and do not replicate to RODCs for security purposes.

Installing an RODC

Several prerequisites must be met before you install an RODC. They are:

- The forest functional level must be at least at the Windows Server 2003 level.
- A writable domain controller that is running Windows Server 2008 or newer must be in the same domain.
- You must have prepared the domain with the **Adprep.exe /rodcprep** command.

After ensuring these prerequisites, you can install an RODC through the **Active Directory Configuration Wizard**. On the **Additional Domain Controller Options** page, select the check box for RODC.

Demonstration: Deploying an RODC

In this demonstration, you will see how to:

- Add a server that you will manage.
- Create a new server group.
- Install the RODC role remotely.
- Configure the password replication policy.

Demonstration Steps

Add a server that you will manage

1. Switch to **LON-DC1**.
2. Use the **Server Manager** dashboard to add **LON-SVR3** as a server to manage.

Create a new server group

1. Use the **Server Manager** dashboard to create a server group named **DCs**.
2. Add **LON-SVR3** and **LON-DC1** to the group.

Install the RODC role remotely

1. Use the **Server Manager** dashboard to add the **Active Directory Domain Services** role to **LON-SVR3**.
2. Open the notifications, and then complete the post-deployment configuration to promote **LON-SVR3** to be a **Read only domain controller (RODC)** in the existing domain.
3. Set the **Directory Services Restore Mode (DSRM) password** to **Pa\$\$w0rd**.
4. Accept the defaults for all other settings.

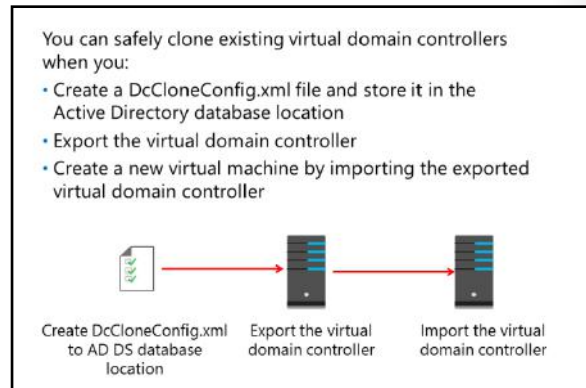
Configure the password replication policy

1. Use **Active Directory Users and Computers** to configure the **LON-SVR3** password caching options in such a way that passwords cache on the RODC for members of the **Managers** group.
2. Configure the **IT** group to have administrative access to **LON-SVR3**.

Cloning virtual domain controllers

Cloning a virtualized domain controller presents its own challenges. For example, two domain controllers cannot coexist in the same forest with the same name, InvocationID, and security identifier. In versions of Windows Server operating systems prior to Windows Server 2012, you must create virtualized domain controllers by deploying a Sysprepped base server image and then promote the domain controller manually. However, both Windows Server 2012 and Windows Server 2016 include specific virtualization capabilities for AD DS to resolve those issues. These virtualization capabilities are as follows:

- You can safely clone domain controllers to deploy additional capacity and to save configuration time.
- Accidental restoration of domain controller checkpoints does not disrupt the Active Directory environment.



Safe cloning

A cloned domain controller automatically Syspreps, based on settings in **DefaultDCCloneAllowList.xml**, and promotes with the existing local Active Directory data as installation media.

Safe backup and restore

For Windows Server versions prior to Windows Server 2012, rolling back to a previous checkpoint of a virtual domain controller is problematic. This is because Active Directory uses multimaster replication that relies on transactions being assigned numeric values called *update sequence numbers* (USNs). These USNs, together with a database identifier named InvocationID, uniquely identify transaction events. The virtual domain controller tries to assign USNs to prior transactions that are already assigned to valid transactions. This occurs because after a virtual domain controller reverts to a previous state, it uses USNs that are associated with other transactions. Other domain controllers discard the transactions based on the USN value. This causes inconsistencies in the Active Directory database.

Windows Server 2012 and Windows Server 2016 implement a process known as *USN rollback protection*. With this process in place, USN rollback is detected, and replication is stopped before divergence in the forest is created. Since Windows Server 2012, Windows Server detects rollbacks automatically, and it non-authoritatively syncs the delta of changes between the domain controller and the replication partners for AD DS and SYSVOL. You can now use checkpoints without risk of corrupting domain controllers.

Creating a virtual domain controller clone

Before you begin to create a virtual domain controller clone, ensure that the primary domain controller (PDC) emulator role is on a domain controller that runs Windows Server 2012 or newer. Additionally, ensure that the source domain controller is not holding the PDC emulator master or the relative identifier (RID) operations master—also known as the *Flexible Single Master Operations* (FSMO)—role. To create a virtual domain controller clone, perform the following high-level steps:

1. Ensure that the PDC emulator master FSMO role holder and the RID operations master FSMO role holder are online and available.
2. Grant the source virtualized domain controller permission to be cloned by adding it to the Cloneable Domain Controllers group.
3. Run the Windows PowerShell **Get-ADDCCloningExcludedApplicationList** cmdlet to identify any programs or services that are not evaluated for cloning.

4. Run the Windows PowerShell **New-ADDCCloneConfigFile** cmdlet to generate a clone configuration file.
5. Export the virtual domain controller.
6. Create a new virtual machine by importing the exported virtual domain controller. This virtual machine is promoted automatically as a unique domain controller.



Note: You should delete any checkpoints before exporting the source domain controller.

Upgrading an Active Directory forest to Windows Server 2016

Although you can perform an in-place upgrade of existing Active Directory domain controllers, the procedure carries certain risks. Generally, a more suitable, and certainly safer, approach is to add a Windows Server 2016 domain controller to an existing Active Directory forest.

Deploying the first Windows Server 2016 domain controller in an existing forest

To add a Windows Server 2016 domain controller to an existing Active Directory domain, you must use the following high-level procedure:

- Ensure that forest and domain functional levels are at least Windows Server 2008
- Prepare the Active Directory forest
- Prepare the Active Directory domain
- Install Windows Server 2016 and add the Active Directory server role
- Promote the computer as a new domain controller in an existing domain
- Transfer operations master roles
- Decommission older domain controllers
- Optionally, raise the forest and domain functional level



1. Ensure that the current forest and domain functional levels are at least Windows Server 2008.



Note: Both file replication service (FRS) and Windows Server 2003 functional levels were deprecated in earlier versions of Windows Server. However, as you consider your approach to implementing Windows Server 2016 domain controllers in your Active Directory environment, you should remember that Windows Server 2003 is no longer supported. Consequently, you must remove any of your domain controllers that are still running Windows Server 2003.

2. Prepare the Active Directory forest. On a domain controller in your existing forest, run **adprep /forestprep** from the **Support\Adprep** folder on the Windows Server 2016 installation media. This procedure also prepares the Active Directory schema.
3. Prepare the Active Directory domain. On a domain controller in your existing forest, run **adprep /domainprep**.



Note: Both these steps were required in Windows Server 2008 R2. In Windows Server 2012 and later, if you do not perform these two tasks manually from an elevated command prompt, the **Active Directory Domain Services Configuration Wizard** performs them automatically. Many larger organizations still prefer to separate the process of preparing the Active Directory forest from promoting the first domain controller. This is because often, separate administrative teams are responsible for those distinct configuration changes.

4. Install Windows Server 2016 on a server computer, and then add the **Active Directory Domain Services** server role with **Server Manager** or by using Windows PowerShell.

5. Promote the Windows Server 2016 server computer as a new domain controller in an existing domain, either by using Windows PowerShell or the **Active Directory Domain Services Configuration Wizard**.



Note: If you perform the deployment on Server Core, the forest and domain preparation tasks complete automatically when you run the Windows PowerShell **Install-ADDSDomainController** cmdlet.

Migrating operations master roles

After you deploy the first Windows Server 2016 domain controller, you can migrate the existing Active Directory operations master roles, previously known as *FSMO roles*, that other domain controllers held. These operations master roles are:

- Schema master (forest-wide). Use the Active Directory Schema snap-in tool to transfer.
- Domain naming master (forest-wide). Use the Active Directory Domains and Trusts snap-in tool to transfer.
- PDC emulator (domain-wide). Use the Active Directory Users and Computers snap-in tool to transfer.
- RID master (domain-wide). Use the Active Directory Users and Computers snap-in tool to transfer.
- Infrastructure master (domain-wide). Use the Active Directory Users and Computers snap-in tool to transfer.



Note: You might consider moving these roles from the existing domain controllers if you intend to remove the existing domain controllers from the forest.

You can also move these roles by using the Windows PowerShell **Move-ADDirectoryServerOperationMasterRole** cmdlet. For example, the following cmdlet transfers the PDC emulator master to **LON-DC2**:

```
Move-ADDirectoryServerOperationMasterRole -Identity "LON-DC2" -OperationMasterRole PDCEmulator
```

After you transfer the operations master roles to new Windows Server 2016 domain controllers, you can remove your older domain controllers from AD DS, and then physically decommission them.

Completing the upgrade

Finally, after you decommission all of your older domain controllers, you can consider changing the forest and domain functional levels to Windows Server 2016.



Note: A number of the new features in AD DS on Windows Server 2016 require the Active Directory forest to be at the Windows Server 2016 functional level.

To raise the forest functional level, you can use the Active Directory Domains and Trusts snap-in tool. To raise the domain functional level, use Active Directory Users and Computers. Alternatively, you can use Windows PowerShell cmdlets. To raise the forest functional level, use the following cmdlets:

```
$ad_forest = get-ADforest
Set-ADForestMode -Identity $ad_forest -Server $ad_forest.SchemaMaster -ForestMode Windows2016Forest
```

To raise the domain functional level, use the following cmdlet:

```
Set-ADDomainMode -identity adatum.com -ForestMode Windows2016Domain
```

Windows Server 2016 domain functional level

Active Directory domains can run at different functional levels. Generally, upgrading a domain to a higher functional level introduces additional features. By upgrading to the Windows Server 2016 domain functional level, you enable the following new features in AD DS:

- Privileged access management
- Azure AD Join
- Microsoft Passport

Because Windows Server 2003 is no longer supported, we recommend that you to raise your domain and forest functional levels to a minimum of Windows Server 2008 to ensure SYSVOL replication consistency.

Check Your Knowledge

Question	
Which of the following commands would you use to promote Windows Server 2016 Server Core to a domain controller?	
Select the correct answer.	
<input type="checkbox"/>	Adprep /forestprep
<input type="checkbox"/>	Adprep /domainprep
<input type="checkbox"/>	Install-ADDSDomainController -domainname "Adatum.com"
<input type="checkbox"/>	Install-WindowsFeature -name AD-Domain-Services

Lesson 2

Implementing service accounts

Many organizations face a common issue: how to manage accounts securely that are used for network services. Many applications use services that require an account for service startup and authentication. Like end user accounts, you must effectively manage service accounts to ensure security and reliability.

Lesson Objectives

After completing this lesson, you will be able to:

- Manage service principal names (SPNs).
- Describe managed service accounts and group managed service accounts.
- Explain how to configure Kerberos delegation.
- Configure managed service accounts.

Managing SPNs

SPNs represent the accounts in whose security context a service executes. SPNs support mutual authentication between a client app and a service. SPNs are built from information that a client computer knows about a service or from a trusted third party, such as AD DS. SPNs are associated with accounts, and an account can have a different SPN for each service it is used to authenticate and execute.

- SPNs represent the user accounts under which services run
- SPNs support mutual authentication between apps and services
- An account can have a different SPN for each service it authenticates and executes
- The basic syntax of an SPN is:
`< service type >/< instance name >:
 < port number >/< service name >`

The basic syntax of an SPN is as follows:

```
< service type >/< instance name >:< port number >/< service name >
```

The following table describes the syntax elements and their meanings.

Element	Description
Service type	This is the type of service, such as www for the World Wide Web service.
Instance name	This is the name of the service instance, which is either the host name or the IP address of the server that is running the service.
Port number	This is the port number that the host uses for the service, if it differs from the default.
Service name	This might be the DNS name of the host, a replicated service, or a domain. Alternatively, it can be the distinguished name of a service connection point object or a remote procedure call (RPC) service object.

If the service name and the instance name are the same, as they are for most host-based services, then you can abbreviate an SPN to two components:

```
< service type >/< instance name>
```

Service names in AD DS

The syntax for service names in AD DS includes the distinguished name of the instance of the service. The syntax is:

```
< service type >/< host name >:< port number >/< distinguished name >
```

Creating new SPNs and viewing existing SPNs

Windows Server includes the **SetSPN.exe** command-line tool to manage SPNs. This tool can register a new SPN when you use the following command:

```
setspn -s http/Server1.adatum.com:80 APP1-SVC
```

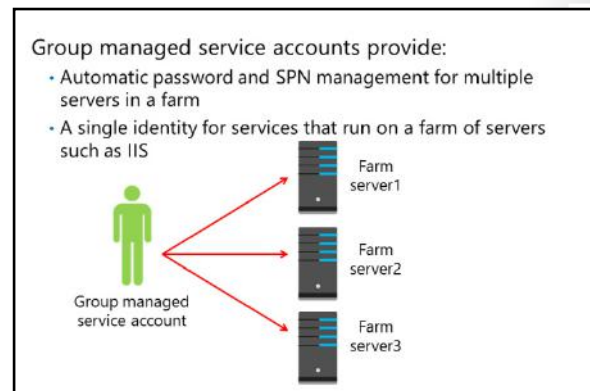
In this example, the server name is **Server1**. It has an Internet Information Services (IIS) application running on port 80. The server is joined to the **adatum.com** domain, and the service account name is **APP1-SVC**. The **-s** switch ensures that there is not already a duplicate SPN before proceeding with the creation.

To view a list of SPNs on **Server1**, use the following command:

```
setspn -l server1.adatum.com
```

What are managed service accounts and group managed service accounts?

Frequently, apps are configured to execute non-interactively on servers that use the security authentication context of the Local Service, Network Service, or Local System accounts. Because many apps and processes typically share these accounts, you cannot isolate their credentials. This means that you cannot customize the security settings of these accounts without also affecting all apps and processes that map to them. A managed service account provides an app its own service account. In Windows Server 2016, you no longer have to administer the credentials for this account manually.



Note: Managed service accounts are also available in Windows Server 2012.

Managed service accounts in Windows Server 2016 offer the following benefits:

- Automatic password management. A managed service account automatically maintains its own password, including password changes. This can better isolate services from other services on a computer.

- Simplified SPN management. SPN management can occur automatically if the Active Directory domain is configured at the Windows Server 2008 R2 domain functional level. For example, if the **samAccountName** property of the computer changes, or if the DNS host name property is modified, the managed service account SPN changes automatically from the old name to the new name for all managed service accounts on the computer.

Stand-alone managed service accounts are managed, domain-based accounts for single servers. Group managed service accounts provide the same functionality, but for multiple servers. When you connect to a service that is hosted on a server farm, such as a web application in an IIS farm, all computers that run an instance of that application must use the same security principal. When a group managed service account is used as the service principal, AD DS on Windows Server 2016 manages the password for the account instead of relying on the administrator to manage the password. We recommend group managed service accounts instead of managed service accounts. While they provide similar functionality, group managed service accounts offer easier expansion of server farms.



Note: You can configure and administer group managed service accounts only on computers that are running Windows Server 2012 or later.

The group managed service account has features to manage hosts that are offline for an extended period. This means that you can deploy a server farm that uses a single group-managed security account identity to which existing client computers can authenticate, without knowing the instance of the service to which they are connecting.



Note: For Windows Server 2012 and Windows Server 2016, the Windows PowerShell cmdlets default to managing group managed service accounts instead of the original stand-alone managed service accounts.

Configuring Kerberos delegation

Kerberos delegation, sometimes referred to as *Kerberos impersonation*, enables a remote computer or service account to act on behalf of a user. Kerberos-constrained delegation adds additional security to the delegation by limiting which resources can be accessed through delegation. Kerberos delegation is widely used in web environments with IIS and Microsoft SQL Server.

A common example is a website on an IIS server with all of the website data stored in a SQL database on a database server. When a user connects to the website, the IIS server must query the SQL database for the data to render the website. With Kerberos delegation, the query occurs based on the user account. Without Kerberos delegation, the query occurs based on a service account. Some of the advantages of using Kerberos delegation in such a scenario include the following:

- The SQL logs show queries from the user account, which is important for auditing and compliance.
- Access to SQL data is based on the user account instead of a service account. This ensures that only data to which the user has access can be returned in a query. Data access varies by user account.

- Kerberos delegation enables a remote computer or service account to act on behalf of a user
- Requirements for Kerberos delegation:
 - A user account cannot be marked as sensitive
 - SPNs must be registered on both sides of the delegation (the service account that is used for delegation and the service account for the target resource)
 - The service account that is used for delegation must be enabled for delegation

- Authentication takes place one time when the user accesses the website. Delegation reduces the number of authentications performed by a user.

To utilize Kerberos delegation, the following requirements must be met:

- The user object must not have the **Account is sensitive and cannot be delegated** option enabled.
- SPNs must be registered for the application service accounts; for example, the IIS and **SQL Server** service accounts.
- The application service account must be trusted for delegation. This enables the account to act on behalf of another user.



Note: Only an account that has a registered SPN can be trusted for delegation.

Demonstration: Configuring managed service accounts

In this demonstration, you will see how to create a group managed service account, and then associate the account with a server.

Demonstration Steps

1. On **LON-DC1**, in **Windows PowerShell (Admin)**, create the Key Distribution Service (KDS) root key by using the **Add-KdsRootKey** cmdlet.
2. Make the effective time **-10** hours so that the key is effective immediately.
3. Create a new service account named **Webservice** for the **LON-DC1** host.
4. Associate the **Webservice** managed account with **LON-DC1**.
5. Use the **Get-ADServiceAccount** cmdlet to verify that the group managed service account was created.

Verify the correctness of the statement by placing a mark in the column to the right.

Statement	Answer
You can only configure and administer group managed service accounts on computers that run Windows Server 2008 R2 or later.	

Lab: Implementing and managing AD DS

Scenario

You are about to deploy additional domain controllers. Your manager asked you to use the clone feature to reduce the administrative effort that is necessary to deploy new domain controllers in the Active Directory forest.

Additionally, A. Datum Corporation wants to centralize management of all accounts that are being used for services and to discontinue using local accounts for that purpose.

Objectives

After completing this lab, you will be able to:

- Clone a domain controller.
- Implement service accounts.

Lab Setup

Estimated Time: 30 minutes

Virtual machine: **20743A-LON-DC1**

User name: **Adatum\Administrator**

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

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

1. On the host computer, start **Hyper-V Manager**.
2. In Hyper-V Manager, click **20743A-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:
 - o User name: **Administrator**
 - o Password: **Pa\$\$w0rd**
 - o Domain: **Adatum**

Exercise 1: Cloning a domain controller

Scenario

You decide to test the procedure for cloning domain controllers and to create a test lab environment. You start the process by preparing the source domain controller, **LON-DC1**.

The main tasks for this exercise are as follows:

1. Prepare a source domain controller to be cloned.
2. Export the source virtual machine.
3. Create and start the cloned domain controller.

► **Task 1: Prepare a source domain controller to be cloned**

1. On **LON-DC1**, from **Server Manager**, start the **Active Directory Administrative Center**, and then add the **LON-DC1** domain controller to the **Cloneable Domain Controllers** Active Directory group.
2. Verify that apps and services on **LON-DC1** support cloning by running the following two cmdlets:

```
Get-ADDCCloningExcludedApplicationList
Get-ADDCCloningExcludedApplicationList -GenerateXML
```

3. Create a **DCCloneConfig.xml** file, and then configure it such that a cloned domain controller is named **LON-DC3**:

```
New-ADDCCloneConfigFile -CloneComputerName "LON-DC3"
```

4. Shut down **LON-DC1**.

► **Task 2: Export the source virtual machine**

1. On the host computer, in **Hyper-V Manager**, export **20743A-LON-DC1** to **D:\Program Files\Microsoft Learning\20743**. Wait until the export finishes. This can take from 10 to 15 minutes.
2. Start and connect to **20743A-LON-DC1** and sign in as **Adatum\Administrator** with the password of **Pa\$\$w0rd**.

► **Task 3: Create and start the cloned domain controller**

1. Import a new virtual machine by using the exported files:
 - a. In the **Import Virtual Machine Wizard**, on the **Before You Begin** page, click **Next**.
 - b. On the **Locate Folder** page, click **Browse**, go to the **D:\Program Files\Microsoft Learning\20743\20743A-LON-DC1** folder, click **Select Folder**, and then click **Next**.
 - c. On the **Select Virtual Machine** page, select **20743A-LON-DC1** if it is not already selected, and then click **Next**.
 - d. On the **Choose Import Type** page, select **Copy the virtual machine (create a new unique ID)**, and then click **Next**.
 - e. On the **Choose Folders for Virtual Machine Files** page, select the **Store the virtual machine in a different location** check box. For each folder location, specify **D:\Program Files\Microsoft Learning\20743** as the path, and then click **Next**.
 - f. On the **Choose Folders to Store Virtual Hard Disks** page, provide the **D:\Program Files\Microsoft Learning\20743** path, and then click **Next**.
 - g. On the **Completing Import Wizard** page, click **Finish**. The machine imports. This can take from 10 to 15 minutes or longer.



Note: You can continue with the next exercise while the import proceeds.

- h. In the **management** list, identify and select the newly imported **20743A-LON-DC1** virtual machine, which has the **State** value as **Off**. In the lower section of the **Actions** pane, click **Rename**.
- i. Type **20743A-LON-DC3** as the name, and then press Enter.

2. In Hyper-V Manager, start **20743A-LON-DC3**.
3. Connect to **20743A-LON-DC3**, and then notice that while the server is starting, a "Domain Controller cloning is at x% completion" message displays.

Results: After completing this exercise, you should have successfully cloned a domain controller.

Exercise 2: Implementing service accounts

Scenario

Until now, there has been no consistent policy about accounts that are used for services. On some servers, local accounts are used, while other servers use domain accounts. Password management for these accounts also has not been consistent. Some accounts have non-expiring passwords, while others are updated manually with new passwords. In this scenario, you decide to implement managed service accounts to replace all of these techniques. You will need to create the account, and then assign the account to a web service.

The main tasks for this exercise are as follows:

1. Create and associate a managed service account
2. Configure the Web server application pool to use the group managed service account.
3. Prepare for the next module.

► Task 1: Create and associate a managed service account

1. On **LON-DC1**, create a KDS root key by using the following Windows PowerShell command:

```
Add-KdsRootKey -EffectiveTime ((get-date).addhours(-10))
```

2. Create a new service account named **Webservice** for the **LON-DC1** host by typing the following command:

```
New-ADServiceAccount -Name Webservice -DNSHostName LON-DC1 -  
PrincipalsAllowedToRetrieveManagedPassword LON-DC1$
```

3. Associate the **Webservice** managed account with **LON-DC1** by typing the following command:

```
Add-ADComputerServiceAccount -identity LON-DC1 -ServiceAccount Webservice
```

4. Verify that the group managed service account was created by using the **Get-ADServiceAccount -Filter *** cmdlet.
5. Install the **Webservice** service account by using the **Install-ADServiceAccount -Identity Webservice** cmdlet.

► **Task 2: Configure the Web server application pool to use the group managed service account**

1. On **LON-DC1**, open **Internet Information Services (IIS) Manager**, and then configure the **DefaultAppPool** to use the **Webservice** account as the identity.
2. Stop and then restart the application pool.



Note: If you did not complete Exercise 1, "Cloning a domain controller," do so before reverting the virtual machines.

Results: After completing this exercise, you should have successfully implemented service accounts.

► **Task 3: Prepare for the next module**

When you have finished the lab, revert the virtual machines to their initial state:

1. On the host computer, start **Hyper-V Manager**.
2. In the **Virtual Machines** list, right-click **20743A-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Shut down **20743A-LON-DC3**.

Question: What are two benefits of using managed service accounts in Windows Server 2016?

Lesson 3

Azure AD

Understanding the benefits of Azure AD is an important part of designing and maintaining an identity infrastructure. In this lesson, you will learn what Azure AD is, where it fits in a typical environment, and what services it offers. You will also learn about Multi-Factor Authentication in Azure AD.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe Azure AD.
- Explain when to use Azure AD.
- List and describe Azure AD authentication protocols.
- Explain Multi-Factor Authentication in Azure AD.
- Describe Azure AD Join.
- Describe the options for integrating Azure and on-premises AD DS.
- Explain how to integrate Azure AD with applications.
- Explain how to deploy Active Directory domain controllers in Azure.

What is Azure AD?

AD DS is an identity provider, a directory service, and an access management solution. Cloud-based identity providers support the same functionality natively. Azure AD is an example of such a provider.

It might be easy to view Azure AD as a cloud-based counterpart of AD DS. However, while they share some common characteristics, there are also several significant differences between them.

Primarily, Azure AD is implemented as a Microsoft managed service that is part of a platform as a service (PaaS) offering. It is not part of a core infrastructure that customers own and manage, nor is it an infrastructure as a service (IaaS) offering. While this implies that you have less control over its implementation, it also means that you do not have to dedicate resources to its deployment or maintenance. You also do not have to develop additional functionality that is unavailable in AD DS, such as support for Multi-Factor Authentication, because this is a part of Azure AD functionality.

Azure AD:

- Microsoft-managed
- A PaaS offering
- Multitenant by design
- Employs Internet-compatible protocols
- Supports users, groups, applications, and devices
- No OUs or computer objects
- Does not support domain join or Group Policy settings



Note: Azure AD is not the same as deploying virtual machines in Azure and adding AD DS to those virtual machines, thereby deploying some domain controllers for a new forest and domain. Azure AD is a small subset of AD DS that is built into Azure and provides authentication and authorization services in the cloud. Azure AD is not meant to replace an on-premises deployment of AD DS.

Types of tiers

Azure AD constitutes a separate Azure service. In its most elementary form, Azure AD is included automatically with any new Azure subscription, and it does not entail any extra cost. This is referred to as a *Free tier*. Advanced identity management features require paid versions of Azure AD. These are offered in the form of *Basic* and *Premium* tiers. Some of these advanced features also are included automatically in instances of Azure AD that are generated as part of **Office 365** or **Microsoft Intune** subscriptions.

While the Free tier has a 500,000-object limit, the Basic and Premium tiers do not have restrictions on the total number of directory objects. Although the Free tier does not include any service level agreements (SLAs), the Basic and Premium tiers are bundled with a 99.9 percent uptime SLA.

Tenants

Unlike AD DS, Azure AD is multitenant by design, and is implemented specifically to ensure isolation between individual directories. It hosts well over 1 million directory services instances, with billions of authentication requests per week from its multitude of tenants. The term *tenant* in this context typically represents a company or organization that has a subscription to a Microsoft cloud-based service such as **Office 365**, **Intune**, or **Azure**, which uses Azure AD.

Directories

When you create your first Microsoft cloud service subscription, you also automatically generate a new Azure AD directory instance, also referred to as a *directory*. The directory is assigned the default DNS domain name, consisting of a unique name of your choice followed by the "onmicrosoft.com" suffix. It is possible and quite common to add at least one custom domain name that utilizes the DNS domain namespace that the tenant owns.

The directory serves as the security boundary and a container of Azure AD objects, such as users, groups, and applications. It is possible for a single directory to support multiple cloud service subscriptions.

The Azure AD schema contains fewer object types than the Active Directory schema. Most notably, it does not include a definition of the computer class, because there is no process of joining computers to Azure AD. However, Azure AD does facilitate device registration, similar to the Workplace Join feature of AD DS. The Azure AD schema is also easily extensible, and its extensions are fully reversible.

The lack of support for domain membership means that you cannot use Azure AD to manage computers or user settings by using Group Policy Objects (GPOs). Instead, its primary strength lies in providing directory services; storing and publishing user, device, and application data; and handling the authentication and authorization of users, devices, and applications. These features are effective and efficient in existing deployments of cloud services such as **Office 365**, which rely on Azure AD as their identity provider.

Azure AD identity models

Applications are represented in Azure AD by objects of the **Application** class and the **servicePrincipal** class, with the former containing an application definition and the latter constituting its instance in the current Azure AD directory. Separating these two sets of characteristics allows you to define an application in one directory and use it across multiple directories by creating a service principal object for the application in each directory. This facilitates deploying applications to multiple tenants.

Delegation model

Because of its operational model as a software as a service (SaaS) offering and its lack both of management capabilities via Group Policy settings and support for computer objects, the Azure AD delegation model is considerably simpler than the same model in AD DS. All three tiers have several built-in roles, including Global Administrator, Billing Administrator, Service Administrator, User Administrator, and Password Administrator. Each of these roles provides different levels of directory-wide permissions to

its objects. By default, the administrators of the subscription that hosts the Azure AD instance are its Global Administrators, with full permissions to all objects in their directory instance.

Some management actions depend on the type of Azure AD tier. For example, in the Azure AD Free tier, users can access a set of designated applications via the Access Panel. With the Azure AD Basic tier, you can grant such access based on group membership. The Premium tier further extends this functionality by offering delegated and self-service group management, allowing users to create and manage their own groups, and allowing requests for membership in groups that others created.

Role-based access control (RBAC)

The previously described delegation model applies to the GUI that is available in the Azure classic portal (<http://manage.windowsazure.com>). The Azure portal (<http://portal.azure.com>) offers a much more flexible and precise way of restricting management of Azure resources by implementing RBAC. This mechanism relies on three built-in roles: owner, contributor, and reader. Each of these roles performs a specific set of actions on Azure resources—such as websites or SQL databases—that are exposed via the Azure Portal. You can grant the intended access by associating an Azure AD object such as a user, group, or service principal with a role and a resource that appears in the Azure Portal. Note that this approach applies only to resources that are available via the Azure Portal.

Azure AD does not include the organizational unit (OU) class, which means that you cannot arrange its objects into a hierarchy of custom containers, which frequently is used in on-premises Active Directory deployments. This is not a significant shortcoming, because OUs in AD DS are primarily for Group Policy scoping and delegation. Instead, you can accomplish equivalent arrangements by organizing objects based on their attribute values or group membership.

Authenticating access to Azure web applications by using Azure AD

The process of implementing Azure AD support for custom applications is rather complex and beyond the scope of this course. However, the Azure portal and **Microsoft Visual Studio 2013** makes the process of configuring such support more straightforward.

In particular, you can enable Azure AD authentication for Azure websites directly from the **CONFIGURE** page in the Azure portal. By designating the Azure AD directory instance, you can ensure that only users with accounts in that directory are able to access the website. It is possible to apply different authentication settings to individual deployment slots.

In the case of **Visual Studio 2013**, when developing web application projects, you can choose to configure authentication based on organizational accounts, you can automatically register the application with Azure AD, and you can assign its access level to directory content. When you use older versions of **Visual Studio**, you must register the application manually. You can do this when you add its unique identifier, referred to as the *App ID Uniform Resource Identifier* (URI), to the target Azure AD instance from the Azure portal.

Azure AD federations

In Azure AD, Active Directory federations have replaced trust relationships between domains and forests. This allows its directories to integrate with cloud services and to interact with directory instances of other Azure AD tenants and other identity providers. For example, such federation trust exists between Azure AD and the Microsoft identity provider that hosts Microsoft accounts (formerly known as *Live ID* accounts). This means that an Azure AD user account can reference an existing Microsoft account directly, making it possible to use the Microsoft account to sign in to Azure AD. You can also use AD FS and Web Application Proxy to establish such federations with on-premises Active Directory deployments. Using federations eliminates dependency on Active Directory protocols, such as Kerberos, which are best suited for on-premises, local area network (LAN)-based communication for which trust relationships were designed. Instead, federation traffic travels over HTTPS, carrying Web Services Trust (WS-Trust), Web Services Federation (WS-Federation), Security Assertion Markup Language (SAML), or Open Authorization

(OAuth) messages. Instead of using Lightweight Directory Access Protocol (LDAP)-based lookups, Azure AD queries rely on the Azure AD Graph application programming interface (API).

Azure AD identity support

Because of its built-in capabilities as an identity provider and its support for federations, Azure AD provides flexibility in designing an identity solution for your organizational or business needs. This gives you three high-level design choices:

- Fully delegate authentication and authorization to Azure AD. Effectively, this means that identity data, including user credentials, resides only in the cloud. You can define the identities directly in Azure AD, or source them from existing Microsoft accounts, based upon the federation with the Microsoft identity provider. You might prefer this choice if you do not have an existing or significant on-premises Active Directory deployment.
- Maintain an on-premises authoritative source of the identity data in AD DS, which synchronizes in regular intervals to Azure AD. This method allows Azure AD to authenticate and authorize users, but you retain on-premises control over their state. This approach simplifies application support of Active Directory users who are not operating on-premises. It also is suitable in scenarios where a large number of Active Directory users rely on Azure Cloud Services such as **Office 365** to access their applications.
- Form a federation between your on-premises AD DS and Azure AD. Authentication requests that are submitted to Azure Cloud Services redirect from the cloud to your on-premises AD DS via the AD FS server. In effect, this allows you to provide authentication and authorization to cloud-based services by using your on-premises AD DS. This approach is similar to the second design choice, but its distinct advantage is support for single sign-on (SSO).

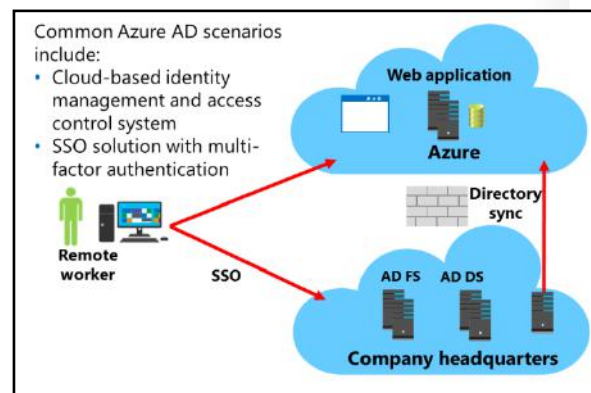
When to use Azure AD

As an identity and access management solution, Azure AD provides a range of features that integrate with other cloud and on-premises services. You can use Azure AD to implement authentication for Azure websites, Azure Cloud Services, and web applications that run on Azure virtual machines. By specifying the Azure AD directory instance for any of these scenarios, you can help ensure that only users with accounts in that directory will gain access.

Similarly, you have the ability to delegate management of any Azure AD service that is accessible via the Azure Portal when you employ RBAC. You also can use Azure AD accounts when designating the service administrator and co-administrator of a subscription.

Azure AD offers additional authentication enhancements, including Multi-Factor Authentication, and SSO for Azure AD provides the following benefits:

- High availability. Azure AD is generally much more highly available than most organizations can provide on-premises.
- Scalable. The Azure environment can scale to meet large workloads on demand.
- Disaster recovery. Azure AD has built-in disaster recovery and the ability to be a disaster recovery site for on-premises AD DS.



- Integration with on-premises AD DS, including directory synchronization and SSO. This includes the Azure AD ability to limit the data that synchronizes to Azure AD.
- APIs. The Azure Service Management representational state transfer (REST) API provides developers the ability to perform management portal tasks programmatically. The Graph API allows developers to query directory data from their applications.

Getting started with Azure AD

The high-level steps for Azure AD without SSO are as follows:

1. Sign up for Azure. This creates an account for managing your Azure subscriptions.
2. Add the Azure AD service from the Microsoft Azure management portal.
3. Add a custom domain name to Azure AD (optional).
4. Add applications that integrate with Azure AD (optional).
5. Add directory integration for your on-premises AD DS (optional).
6. Add users.

To continue the deployment for SSO and directory synchronization, follow these steps:

1. Install Azure AD Connect and configure synchronization.
2. Deploy AD FS and configure federation.

Azure AD authentication protocols

Azure AD authentication protocols differ from Active Directory authentication protocols. Often, Active Directory administrators are less experienced with web-based authentication protocols. Azure AD supports a few different authentication protocols:

- **OAuth 2.0.** Based on RFC 6749, OAuth 2.0 is an open standard for authorization that provides precise access control to destination services. Access can be on a temporary basis. OAuth allows decoupling of the authentication credentials, which allows credentials not to be passed to the destination.
- **SAML 2.0.** SAML is an open standard XML protocol that is made up of security tokens. A security token contains claims, which typically are Active Directory attributes that are used to make decisions for authorization and access.
- **WS-Federation.** WS-Federation is a security mechanism that allows identity federation so that users in one realm, or *directory*, can access resources in another realm.

The supported protocols have some commonalities, including that all are web-based protocols for use on the Internet. Conversely, Active Directory authentication protocols were designed for use on a private network, and initially, without a need for open standards for authentication.

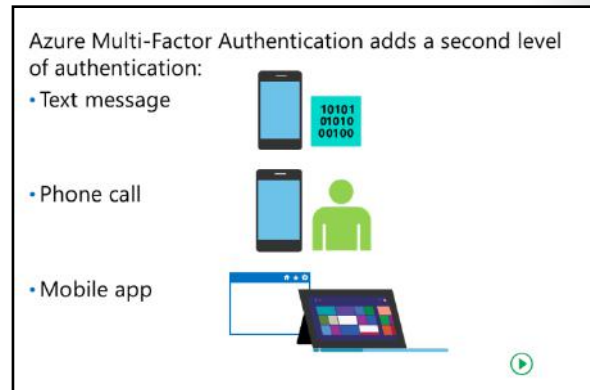
Supported Azure AD authentication protocols include:

- OAuth 2.0
- SAML 2.0
- WS-Federation



Multi-Factor Authentication

The purpose of Multi-Factor Authentication is to increase security. Traditional, standard authentication requires knowledge of sign-in credentials, typically consisting of a user name and an associated password. Multi-Factor Authentication adds an extra verification that relies on having access to a device that is presumably in the possession of the rightful owner or having the correct biometrics. This additional requirement makes it considerably more difficult for an unauthorized individual to compromise the authentication process.



Azure Multi-Factor Authentication

Azure Multi-Factor Authentication is integrated into Azure AD. It allows the use of a phone as the physical device that can provide the means to confirm a user's identity. The process of implementing Multi-Factor Authentication for an Azure AD user account starts when a user with the global administrator role enables the account for Multi-Factor Authentication from the Azure portal. At the next sign-in attempt, the user is prompted to set up the authentication by selecting one of the following options:

- Mobile phone. This requires a user to provide a mobile phone number. The verification can be in the form of a phone call (at the end of which, the user must press the pound key), or a text message.
- Office phone. This requires the specification of the OFFICE PHONE entry of the user's contact info in Azure AD. The administrator must preconfigure this entry, and the user cannot modify or provide this entry at verification time.
- Mobile app. This requires a user to have a smartphone on which they must install and configure the mobile phone app.

App passwords

Because Multi-Factor Authentication is limited to authenticating via a browser, as part of the verification process, the user can also generate app passwords. The user can then assign randomly generated app passwords to individual apps by using their configuration settings. However, app passwords can be a potential security vulnerability. Therefore, as an administrator, you can prevent all directory users from creating app passwords. You also can invalidate all app passwords for an individual user if the computer or device where the apps are installed is compromised.



Note: Effectively, Multi-Factor Authentication does not apply to traditional desktop applications or modern apps such as **Microsoft Outlook**, **Skype for Business**, or mobile apps for email.

After the verification process is complete, the Multi-Factor Authentication status for the user changes from enabled to enforced. The same verification process repeats during every subsequent authentication attempt. The **Additional security verification** option appears in the Access Panel, reflecting the status change. From the Access Panel, you can choose and configure a different verification mechanism, and you can generate app passwords. Generating app passwords is especially important, because without app passwords assigned, desktop apps and modern apps that rely on authenticated access to Azure AD fail to connect to Azure Cloud Services.

What is Azure AD Join?

Joining your organization's devices to your Active Directory domain provides users the best experience for accessing domain-based resources and apps. By using Azure AD Join, you can provide a better experience for users of cloud-based apps and resources. You can also use Azure AD Join to manage your organization's devices from the cloud by using mobile device management (MDM) instead of using GPOs or with Microsoft System Center Configuration Manager.

When determining whether to implement Azure AD Join, consider the following usage scenarios:

- Your organization's apps and resources are mostly cloud-based
- Your organization employs seasonal workers or students
- You want to allow on-premises users to use their own devices

Usage scenarios

When determining whether to implement Azure AD Join, consider the following usage scenarios:

- Your organization's apps and resources are mostly cloud-based. If your organization currently uses or is planning to use SaaS apps (for example, **Office 365**), you should consider using Azure AD Join. Users can join their Windows 10 devices to Azure AD themselves. When they sign in with their Azure AD credentials, they experience SSO to **Office 365** and any other apps that use Azure AD for authentication.
- Your organization employs seasonal workers or students. Many organizations rely on two pools of staff: permanent employees, such as faculty or corporate staff, and students or seasonal workers who do not remain with the organization for long. In this situation, you can continue to manage permanent employees by using your on-premises AD DS (connected to Azure AD). You can manage seasonal and temporary identities in the cloud by using Azure AD. With Azure AD, these cloud-only users get the same SSO on their devices and to **Office 365** and other cloud resources that had previously only been available to on-premises users.
- You want to allow on-premises users to use their own devices. In this scenario, you can provide users with a simplified joining experience for their own personal Windows 10 devices. You can use Azure AD for automatic MDM enrollment and conditional access for these users' devices Azure AD. Users now have SSO to Azure AD resources in addition to on-premises resources.

On-premises AD DS and Azure integration options

Three options exist for integrating Azure with your on-premises AD DS. These three options are:

- Extending on-premises AD DS into Azure. With this option, you host virtual machines in Azure, and you then promote the virtual machines to be domain controllers within your on-premises AD DS.

Because Azure provides IaaS facilities and can host virtual machines in the cloud, by using Azure for hosting domain controllers, you thereby extend the boundaries of your on-premises domains on this platform. Hosting

domain controllers in Azure can provide a range of benefits, both for on-premises users and for those who connect to both on-premises and Azure-based services from around the world.

The following are on-premises AD DS and Azure AD integration options:

- Extend on-premises AD DS into Azure
- Sync AD DS with Azure AD with optional password syncing
- Implement federation and SSO between on-premises AD DS and Azure AD

- Syncing on-premises AD DS with Azure AD. Azure AD Connect propagates user, group, and contact information to AD DS and keeps that information synced. Azure AD Connect can be used with optional password synchronization so that users sign in to Azure by using the same user account and password as their on-premises account, although the authentication processes are still separate.

Azure AD Connect is the engine that provides user, group, and contact synchronization between on-premises AD DS and Azure AD. In its simplest form, you install the Azure AD Connect component on a member server in your on-premises domain, provide an account with Domain Admin and Enterprise Admin access to AD DS, provide another account with administrator access to Azure, and then run the sync. After a few minutes, all of your user accounts, groups, and contacts from AD DS will replicate to Azure AD. Users then can use their accounts to sign in to and access services in Azure.



Note: Unless you activate password synchronization, users must have a different password from their on-premises password for when they sign in to a Azure resource.

- Implementing federation and SSO between on-premises AD DS and Azure AD. This third option supports the largest range of integration features and enables a user to sign in to Azure after being authenticated by the on-premises AD DS. AD FS is the technology that is used, and a typical implementation uses AD FS proxies to manage incoming authentication requests from the Internet.

Integrating Azure AD with applications

While Azure AD integrates with the on-premises AD DS environments, for applications, there is more to do to implement integration across multiple applications. You can use Azure AD as an authentication service for different types of applications, such as:

- On-premises applications.
- Azure applications.
- Applications hosted with another provider.

You can manage authentication by using WS-Federation, SAML, or OAuth. LDAP and Kerberos

authentication are not available. Azure AD can also provide authentication services for multitenant applications. In this scenario, you can have an application that authenticates users from multiple Azure AD tenants. When working with multitenant applications, issues of security and privacy are critical. Azure offers multiple partitioning schemes to meet a variety of requirements.

Azure AD:

- Integrates with three types of applications:
 - On-premises applications
 - Azure applications
 - Applications hosted with another provider
- Offers the ability for multitenant applications:
 - Privacy and security are critical for multitenant deployments
 - Azure offers multiple partitioning schemes
- Uses WS-Federation, SAML, or OAuth:
 - LDAP and Kerberos authentication are not available

Deploying Active Directory domain controllers in Azure

You can extend AD DS into the cloud by deploying Active Directory domain controllers on virtual machines that are based on Azure IaaS. However, ensuring that you protect such domain controllers from unauthorized external access is critical. You can use such deployments to build a disaster recovery solution for an existing on-premises Active Directory environment, to implement a test environment, or to provide local authentication and authorization to Azure-hosted cloud services that are part of the same virtual network. Deploying domain controllers in Azure provides the following benefits:

Deploying domain controllers in Azure:


- Provides resilience for the on-premises directory
- Keeps authentication requests for Azure-based services within the Azure environment
- Extends access to on-premises AD DS to worldwide sites
- Enables additional directory synchronization options such as Azure AD Connect and SSO with AD FS




- Provides resilience for the on-premises directory.
- Keeps authentication requests for Azure-based services within the Azure environment.
- Extends access to on-premises AD DS to worldwide sites.
- Enables additional directory synchronization options such as Azure AD Connect and SSO with AD FS.

Use the following high-level procedure for deploying Azure virtual machines as Active Directory domain controllers:


1. Create an Azure virtual network.
2. Create the virtual machines that are necessary to run the Active Directory domain controller and DNS server roles.

 **Note:** You should deploy at least two virtual domain controllers to provide fault tolerance and redundancy.

3. Install the Windows Server AD DS server role.

 **Note:** Use the same procedure to deploy AD DS that you use for your on-premises domain controllers; for example, you can use **Server Manager**, **Install from Media**, or **Windows PowerShell**.

4. Reset the DNS server for the Azure virtual network.
5. Create virtual machines for domain member computers.

 **Note:** For more information on deploying AD DS on Azure virtual machines, refer to: <http://aka.ms/Qxtizz>

Question: Do you think that your organization might use Azure AD? In what capacity will your organization implement it?

Module Review and Takeaways

Review Questions

Question: To implement Microsoft Passport, which operating systems must be on users' devices?

Question: To enable Kerberos delegation, what requirements must be met?

Question: What are the benefits of implementing Azure AD?

MCT USE ONLY. STUDENT USE PROHIBITED

Module 4

Implementing AD FS

Contents:

Module Overview	4-1
Lesson 1: Overview of AD FS	4-2
Lesson 2: Deploying AD FS	4-12
Lesson 3: Implementing AD FS for a single organization	4-18
Lab A: Implementing AD FS	4-24
Lesson 4: Implementing Web Application Proxy	4-29
Lab B: Implementing Web Application Proxy	4-35
Lesson 5: Implementing SSO with Microsoft Online Services	4-39
Module Review and Takeaways	4-42

Module Overview

Active Directory Federation Services (AD FS) in the Windows Server 2016 operating system enables organizations to provide their users the flexibility to sign in and authenticate to applications that are located on a local network, at a partner company, or in an online service. With AD FS, your organization can manage its own user accounts, and users only have to remember one set of credentials. Those credentials can provide access to a variety of applications, which are located in a variety of places.

This module presents an overview of AD FS and provides details on how to configure AD FS in a single-organization scenario. This module also describes the Web Application Proxy feature in Windows Server 2016 that functions as an AD FS proxy and reverse proxy for web-based applications. Finally, this module describes Microsoft Azure AD FS.

Objectives

After completing this module, you will be able to:

- Describe AD FS.
- Explain how to deploy AD FS.
- Explain how to implement AD FS for a single organization.
- Explain how to extend AD FS to external clients.
- Describe how to implement an SSO to support online services.

Lesson 1

Overview of AD FS

AD FS is the Microsoft implementation of an identity federation framework that enables organizations to establish federation trusts and share resources across organizational and Active Directory Domain Services (AD DS) boundaries. AD FS is compliant with common web services standards, thus enabling interoperability with identity federation solutions that other vendors provide. AD FS addresses a variety of business scenarios in which the typical authentication mechanisms used in an organization do not work.

This lesson provides an overview of the concepts and standards that are implemented in AD FS and the business scenarios that AD FS can address.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe identity federation.
- Describe claims-based identity.
- Describe web services.
- Describe AD FS.
- Describe Web Application Proxy.
- Explain how AD FS enables single sign-on (SSO) in a single organization.
- Describe Device Registration.

What is identity federation?

Identity federation enables you to provide identification, authentication, and authorization across organizational and platform boundaries. You can implement identity federation within a single organization to enable access to diverse web applications or between two organizations that have an established trust relationship.

To establish an identity federation partnership, both partners agree to create a federated trust relationship. This federated trust is based on an ongoing business relationship, and it enables the organizations to implement business processes that are identified in the business relationship.

Identity federation:

- Enables identification, authentication, and authorization across organizational and platform boundaries
- Requires a federated trust relationship between two organizations or entities
- Enables organizations to retain control over who can access resources
- Enables organizations to retain control of their user and group accounts



Note: A federated trust is not the same as a forest trust that organizations can configure between AD DS forests. In a federated trust, the AD FS servers in two organizations never have to communicate directly with each other. In addition, all communication in a federation deployment occurs over Hypertext Transfer Protocol over Secure Sockets Layer (HTTPS), so you do not need to open multiple ports on any firewalls to enable federation.

As a part of the federated trust, each partner defines which of its resources are accessible to the other organization and how access to the resources is enabled. For example, to update a sales forecast, a sales representative might need to collect information from a supplier's database that is hosted on the supplier's network. The administrator of the domain for the sales representative is responsible for ensuring that the appropriate sales representatives are members of the group that requires access to the supplier's database. The administrator of the organization where the database is located is responsible for ensuring that the partner's employees have access only to the data that they require.

In an identity federation solution, user identities and their associated credentials are stored, owned, and managed by the organization where the user is located. As part of the identity federation trust, each organization also defines how user identities are shared securely to restrict access to resources. Each partner must define the services that it makes available to trusted partners and customers and which other organizations and users it trusts. Each partner also must define what types of credentials and requests it accepts, and each partner must define its privacy policies to ensure that private information is not accessible across the trust.

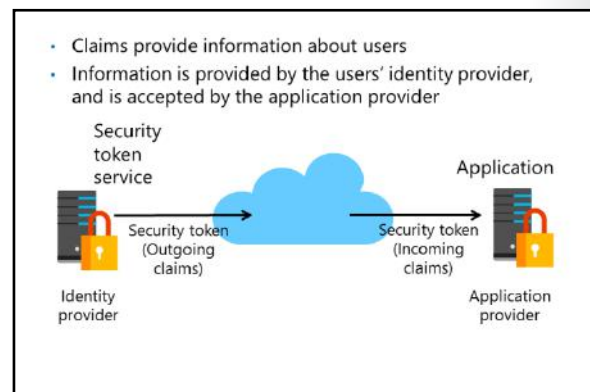
You also can use identity federation within a single organization. For example, an organization might plan to deploy several web-based applications that require authentication. By using AD FS, the organization can implement one authentication solution for all of the applications, making it easy for users in multiple internal domains or forests to access the application. The solution also can extend to external partners in the future, without requiring developers to change the application.

What is claims-based identity?

In most organizations, users sign in to the network and are authenticated by an AD DS domain controller. A user who provides the right credentials to the domain controller is granted a security token. Applications that are running on servers in the same AD DS environment trust the security tokens that the AD DS domain controllers provide, because the servers can communicate with the same domain controllers where the users authenticate.

That type of authentication does not extend easily outside of AD DS forest boundaries. Although trusts based on the Kerberos V5 authentication protocol or NT LAN Manager (NTLM) can be implemented between two AD DS forests, client computers and domain controllers on both sides of the trust must communicate with domain controllers in the other forest to make decisions about authentication and authorization. This communication requires that network traffic is sent on multiple ports, so these ports must be open on all firewalls between the domain controllers and other computers. The problem becomes even more complicated when users have to access resources that are hosted in cloud-based systems, such as Microsoft Azure or Microsoft Office 365.

Claims-based authentication provides a mechanism for separating user authentication and authorization from individual applications. With claims-based authentication, users can authenticate to a directory service that is located within their organization and obtain a claim based on that authentication. The claim is then presented to an application that is running in a different organization. The application allows user access to information or features based on the claims presented. All communication occurs over HTTPS.



The *claim* that is used in claims-based authentication is a statement about a user that is defined in one organization or technology and trusted in another. The claim could include a variety of information. For example, the claim could define the user's email address, user principal name (UPN), and information about specific groups to which the user belongs. This information is collected from the identity store when the user successfully authenticates.

The organization that manages the application defines the types of claims that the application will accept. For example, the application might require the user's email address to verify identity and then use the group membership that is presented inside the claim to determine what level of access the user should have within the application.

Web services overview

For claims-based authentication to work, organizations must agree on the format for exchanging claims. Rather than have each business define this format, a set of specifications broadly identified as web services has been developed. Any organization that is interested in implementing a federated identity solution can use this set of specifications.

Web services are a set of specifications that are used for building connected applications and services. Their functionality and interfaces are exposed to potential users through web technology standards such as Extensible Stylesheet Language (XML), Simple Object Access Protocol (SOAP), Web Services Description Language (WSDL), Hypertext Transfer Protocol (HTTP), and HTTPS. The goal of creating web applications by using web services is to simplify interoperability for applications across multiple development platforms, technologies, and networks.

To enhance interoperability, web services are defined by a set of industry standards, which include the following:

- Most web services use XML to transmit data through HTTP and HTTPS. With XML, developers can create their own customized tags, thereby facilitating the definition, transmission, validation, and interpretation of data between applications and organizations.
- Web services expose useful functionality to web users through a standard web protocol. In most cases, web services use the SOAP protocol, which is the communications protocol for XML web services. SOAP is a specification that defines the XML format for messages; it essentially describes what a valid XML document looks like.
- Web services provide a way to describe their interfaces in enough detail to enable a user to build a client application to communicate with the service. This description usually is provided in an XML document called a WSDL document. In other words, a WSDL file is an XML document that describes a set of SOAP messages and how those messages are exchanged.
- Web services are registered so that potential users can find them easily. This is done with Universal Description, Discovery, and Integration (UDDI). A UDDI directory entry is an XML file that describes a business and the services it offers.

- Web services are a standardized set of specifications used to build applications and services
- Web services typically:
 - Transmit data as XML
 - Use SOAP to define the XML message format
 - Use WSDL to define valid SOAP messages
 - Use UDDI to describe available web services
- SAML is a standard for exchanging identity claims

Web services security specifications

Web services specifications include several components, commonly known as WS-* specifications. However, the most relevant specifications for an AD FS environment are the Web Services Security (WS-Security) specifications. WS-Security includes the following specifications:

- **WS-Security: SOAP Message Security and X.509 Certificate Token Profile.** WS-Security describes enhancements to SOAP messaging that provide quality of protection through message integrity, message confidentiality, and single-message authentication. WS-Security also provides a general-purpose—yet extensible—mechanism for associating security tokens with messages. Additionally, it provides a mechanism to encode binary security tokens (specifically, X.509 certificates and Kerberos tickets) in SOAP messages.
- **Web Services Trust (WS-Trust).** WS-Trust defines extensions that build on WS-Security to request and issue security tokens and to manage trust relationships.
- **Web Services Federation (WS-Federation).** WS-Federation defines mechanisms that WS-Security can use to enable attribute-based identity, authentication, and authorization federation across different trust realms.
- **WS-Federation Passive Requestor Profile (WS-F PRP).** This WS-Security extension describes how passive clients, such as web browsers, can acquire tokens from a federation server, and how the clients can submit tokens to a federation server. Passive requestors of this profile are limited to the HTTP or HTTPS protocol.
- **WS-Federation Active Requestor Profile (WS-F ARP).** This WS-Security extension describes how active clients, such as SOAP-based mobile-device applications, can be authenticated and authorized and how the clients can submit claims in a federation scenario.

Security Assertion Markup Language

The Security Assertion Markup Language (SAML) is an XML-based standard for exchanging claims between an identity provider and a service or application provider. SAML assumes that a user has been authenticated by an identity provider and that the identity provider has populated the appropriate claim information in the security token. When the identity provider authenticates the user, it passes a SAML assertion to the service provider. Based on this assertion, the service provider can make authorization and personalization decisions within an application. The communication between federation servers is based on an XML document that stores the X.509 certificate for token signing and the SAML 1.1 or 2.0 token.

What is AD FS?

AD FS is the Microsoft implementation of an identity federation solution that uses claims-based authentication. AD FS provides the mechanisms to implement both the identity provider and the service provider components in an identity federation deployment.

- AD FS is the Microsoft identity federation product that can use claim-based authentication
- AD FS has the following features:
 - SSO for web-based applications
 - Interoperability with web services on multiple platforms
 - Support for many clients, such as web browsers, mobile devices, and applications
 - Extensibility to support customized claims from third-party applications
 - Delegation of account management to the user's organization



AD FS provides the following features:

- Enterprise claims provider for claims-based applications. You can configure an AD FS server as a claims provider, which means that the AD FS server can issue claims about authenticated users. This configuration enables an organization to provide its users with access to claims-aware applications in another organization by using SSO.
- Federation Service provider for identity federation across domains. This service offers federated web SSO across domains, thereby enhancing security and reducing overhead for information technology (IT) administrators.

AD FS features

The following are some of the key features of AD FS:

- Web SSO. Many organizations deploy AD DS. After authenticating to AD DS through Windows authentication, users can access all other resources that they have permission to access within the AD DS forest boundaries. AD FS extends this capability to intranet or Internet-facing applications, enabling customers, partners, and suppliers to have a similar, streamlined user experience when they access an organization's web-based applications.
- Web services interoperability. AD FS is compatible with the web services specifications. AD FS employs the federation specification of WS-* called WS-Federation. WS-Federation makes it possible for environments that do not use the Windows Identity Foundation (WIF) identity model to federate with environments that use the Windows operating system.
- Passive and smart client support. Because AD FS is based on the WS-* architecture, it supports federated communications between any WS-enabled endpoints, including communications between servers and passive clients such as browsers. AD FS in Windows Server 2016 enables access for SOAP-based smart clients, such as mobile phones, personal digital assistants, and desktop applications. AD FS implements the WS-F PRP and some of the WS-F ARP standards for client support.
- Extensible architecture. AD FS provides an extensible architecture that supports various security token types, including SAML tokens and Kerberos authentication through Windows authentication, and the ability to perform custom claims transformations. For example, AD FS can convert from one token type to another or it can add custom business logic as a variable in an access request. Organizations can use this extensibility to modify AD FS to coexist with their existing security infrastructure and business policies.
- Enhanced security. AD FS also increases the security of federated solutions by delegating responsibility for account management to the organization closest to the user. Each individual organization in a federation continues to manage its own identities, and each is capable of securely sharing and accepting identities and credentials from other members' sources.

New features in AD FS introduced in Windows Server 2012

The AD FS version that ships with Windows Server 2012 includes the following new features:

- Integration with the Windows Server 2012 operating system. In Windows Server 2012, AD FS is included as a server role that you can install by using Server Manager. When you install the server role, all required operating system components install automatically.
- Integration with Dynamic Access Control (DAC). When you deploy DAC, you can configure user and device claims that are issued by AD DS domain controllers. AD FS can consume the AD DS claims that domain controllers issue. This means that AD FS can make authorization decisions based on both user accounts and computer accounts.
- Windows PowerShell command-line interface cmdlets for administering AD FS. Windows Server 2012 provides several new cmdlets that you can use to install and configure the AD FS server role.

New AD FS features introduced in Windows Server 2016

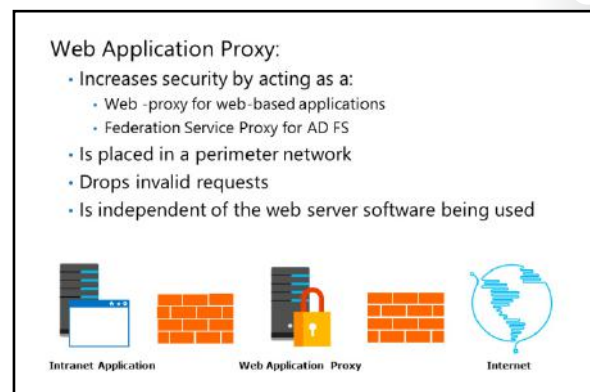
The AD FS version that ships with Windows Server 2016 includes the following new features:

- Support for any LDAP v3 compliant directory. This allows your users to:
 - Sign in to AD FS from any third-party directory that is LDAP v3 compliant.
 - Sign in from untrusted or partially trusted AD DS domains and forests.
- New factors of authentication. AD FS provides additional ways to authenticate users and devices. In addition to AD DS and LDAP v3 directories, you can also configure Azure Multi-Factor Authentication (MFA) as an authentication method.
- Improvements in AD FS management, including:
 - Application policies. In AD FS for Windows Server 2012 R2, you must use claim rules language to create custom AD FS policies. AD FS in Windows Server 2016 provides wizard-based management, making it easier to create custom policies.
 - Delegated service management. AD FS in Windows Server 2016 separates the AD FS server administrators from the AD FS service administrators. This means that the AD FS administrator is no longer required to be a local server administrator.
- Conditional access. AD FS in Windows Server 2016 provides improvements in device registration by working with Azure AD to restrict devices or require multiple factors of authentication based on management or compliance status. For example, with conditional access, you can:
 - Enable access only from your users' devices that are managed and/or compliant with corporate standards.
 - Restrict access to computers joined to the corporate system, including managed devices and computers joined to the domain.
 - Require multi-factor authentication for computers that are not joined to the domain and devices that are not compliant.

Overview of Web Application Proxy

Web Application Proxy in Windows Server 2016 is a remote access role service that you can use to secure remote access to web-based applications on your internal network. Web Application Proxy functions as a reverse proxy for web-based applications and also functions as an AD FS proxy.

You must place Web Application Proxy in a perimeter network, because external clients that access web-based applications or AD FS initiate connections with Web Application Proxy. Web Application Proxy then connects to the web-based application or AD FS on the internal network. You do not need a client-specific configuration to use Web Application Proxy.



When you implement Web Application Proxy, you enhance security for web-based applications or AD FS by isolating them from direct contact with the Internet. This can help protect the internal, web-based application or AD FS from any malformed packets or requests that might result in a security breach. For example, Web Application Proxy can protect against a zero-day vulnerability that uses malformed requests, which could result in a denial-of-service attack on a server that hosts a web-based application.

Web Application Proxy drops invalid requests before they reach the web-based application on an internal network.

Because Web Application Proxy is completely independent of the Web server software being used, it is unlikely that Web Application Proxy is vulnerable to the same denial-of-service attack as a web-based application.



Note: Although Windows Server 2016 and Windows Server 2012 R2 include Web Application Proxy, the role service is not available in Windows Server 2012. Windows Server 2012 does include an AD FS service proxy option that you can install as part of deploying AD FS. This option does not provide reverse proxy functionality for web-based applications. It is a reverse proxy only for AD FS.

AD FS and SSO in a single organization

For many organizations, configuring access to applications and services might not require an AD FS deployment. If all users are members of the same AD DS forest and if all applications run on servers that are members of the same forest, you usually can use AD DS authentication to provide access to the application. However, there are several scenarios in which you can use AD FS to optimize the user experience by enabling SSO. In a single organization, you can use AD FS to enable SSO when:

- Your applications might not be running in Windows-based servers, on any servers that support AD DS authentication, or on servers that are running Windows Server and are not joined to the domain. The applications might require SAML or web services for authentication and authorization.
- You have multiple domains and forests. This might be a result of mergers and acquisitions or due to security requirements. Users in multiple forests might require access to the same applications.
- Users from outside the office might require access to applications that are running on internal servers. External users might log on to applications from computers that are not part of the internal domain.

You can use AD FS to enable SSO in these scenarios. If your organization has a single AD DS forest, you only have to deploy a single federation server. This server can operate as the claims provider so that it authenticates user requests and issues the claims. The same server is also the relying party to provide authorization for application access.

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.
3. The web server redirects the client computer to the Federation Service Proxy.




Note: The Federation Service Proxy server can be a Web Application Proxy server.


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 Windows authentication to collect the user's credentials.

4. The Federation Service Proxy transmits the request and the credentials to the federation server.
5. The federation server uses AD DS to authenticate the user.


6. If authentication is successful, the federation server collects AD DS information about the user. That information is then used to generate the user's claims.

 **Note:** There must be a relying party trust in AD FS for the web app. If the same user tries to access a different web app, different user claims can be included in the security token that is passed to the user and then to the web app.

7. If the authentication is successful, the authentication information and other information is collected in a security token and passed back to the Federation Service Proxy.

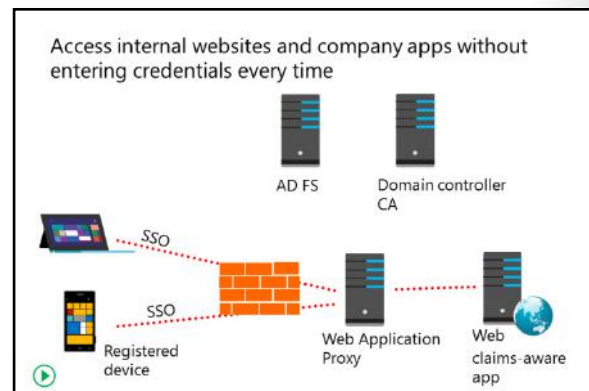
 **Note:** This security token is signed for the specific web app. Consequently, there must be a relying party trust with this web app on AD FS.

8. The Federation Service Proxy then passes the token to the client.
9. The client then presents the token to the web server.
 - o The web resource receives the request and validates the signed tokens.
 - o The web resource uses the claims in the user's token to provide access to the application.

 **Note:** The slide and the preceding description use the terms *Federation Service* and *Federation Service Proxy* to describe AD FS role services. The federation server is responsible for issuing claims and is responsible for consuming the claims in this scenario. The Federation Service Proxy is a proxy component that is recommended for deployments where users outside of the network need access to the AD FS environment. These components are covered in more detail in the next lesson.

What is device registration?

When you join a device to a domain, users can access company resources without having to enter their credentials every time they access company resources. Users can have a similar experience from a device that you enable for Device Registration (previously known as Workplace Join in Windows 8.1) without requiring the device to be a domain member. Device Registration provides an SSO experience when accessing internal company websites and company apps. Users with domain accounts can implement Device Registration on their devices if their company has the appropriate infrastructure in place.



Traditionally, if users wanted to access data transparently from their devices, the devices had to be joined to the domain. If the devices were not joined to the domain, users could still use the devices to access company data, but they had to enter their domain credentials every time they wanted access. Windows 8.1 introduced the Device Registration feature (formerly known as Workplace Join), which enables users to access internal company websites and company apps from devices without having to enter their user credentials each time. Device Registration also enables administrators to have some control over the

devices, such as controlling the web apps that users can access from devices that are enabled for Device Registration.

Device Registration is particularly useful when users access company data with their own devices. If you enable Device Registration, you can register and enroll your devices in the company network.



Note: You must set up AD FS before your users can use the Device Registration feature on their devices. You must configure AD FS with an SSL certificate from a trusted CA, and the SSL certificate must have properly configured Subject Name and Subject Alternative Name attributes.

After you enroll a device, the device is associated with your user account in the company directory. Then the device object is created in AD DS and the user certificate is installed on the device. The device object in AD DS establishes a link between the user and the device. Further communication with company resources that support claims-based authentication (from a device enabled for Device Registration) includes information about the device and the user. Once you properly configure an app to support claim-based authentication, users are not required to enter credentials again.

After you enroll a device, the device is used as a second form of authentication. Administrators can configure which apps users then can access from the device without entering credentials, and they can then ensure that company policies and security apply to those devices by configuring a device policy. Be aware that a company Group Policy applies only to devices joined to the domain and not to devices enabled for Device Registration. If a device enabled for Device Registration is compromised or if a device owner leaves the company, an administrator can remove the device object from the domain; by doing so, the administrator revokes the device's ability to access domain resources through SSO.

Scenarios for using Device Registration

Many devices that employees use to access company data are company owned; those devices usually are joined to the domain. Users also might access company data by using their own devices from inside the company network and over the Internet. The company's IT department can closely monitor and manage computers joined to the domain, but devices that are not domain members can be problematic. Users typically use devices to access virtual desktops, to run company apps, and to access other company resources. Environments that adopt the Bring Your Own Device (BYOD) scenario are particularly suitable for the Device Registration feature. Using this feature, users can access company resources from devices enabled for Device Registration with SSO, and administrators can control access to resources. Administrators also can control the compliance of local copies of company data on such devices while a device is not joined to the domain.

A device that is enabled for the Device Registration feature is used as a second authentication factor when accessing claims-based company apps. For such apps, administrators can control who can access them, from which devices they can be accessed and whether they can be accessed only from the company network or also from the Internet. Devices enabled for Device Registration trust the company certification authority(CA), which makes it easier to configure them for additional features such as Work Folders.

How Device Registration Works

The main purposes of the Device Registration feature are to provide:

- Registration in AD DS for devices that are not joined to the domain.
- SSO for selected application and resources in a company's internal network.

Device Registration works by using Device Registration Service and Active Directory Federation Services (AD FS) with Device Authentication enabled. When a user registers a device through the enrollment process, the Device Registration Service provisions a certificate for the device. This certificate is used to authenticate the device when it accesses internal resources. In addition, the device becomes associated to

the specific user in AD DS, so administrators can configure access policies to apply to users and their registered devices.

By implementing the Web Application Proxy component, you also can enable registered devices to access company resources from external networks such as the Internet. Users can be in a coffee shop or at home, and, if their device is registered, it can access internal applications through Web Application Proxy and AD FS. If the user's device is registered in an internal network, it communicates directly to AD FS and AD DS to authenticate. For devices that are registered, you also can enable SSO for some applications. By doing this, the user does not receive a prompt for credentials each time the user tries to access the resource.

Verify the correctness of the statement by placing a mark in the column to the right.

Statement	Answer
A federated trust is the same as a forest trust that organizations can configure between AD DS forests.	

Lesson 2

Deploying AD FS

After you understand how AD FS works, you can deploy the service. Before you deploy AD FS, you must understand the components that you need to deploy and the prerequisites that you must meet, particularly with regard to certificates. This lesson provides an overview of deploying the AD FS server role in Windows Server 2016.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe AD FS components.
- Describe AD FS prerequisites.
- Describe the public key infrastructure (PKI) and certificate requirements.
- Describe AD FS federation server roles.
- Explain how to install the AD FS server role.

Components in an AD FS deployment

To understand the configuration process for AD FS, you first must understand all of the components that are part of AD FS. These components work together to provide a complete solution for claims-based authentication within an organization or between organizations.

The following table lists the AD FS components.

AD FS components:	
Federation server	Relying parties
Federation server proxy/ Web Application Proxy	Claims provider trust
Claims	Relying party trust
Claim rules	Certificates
Attribute store	Endpoints
Claims providers	

Component	What it does
Federation server	The federation server issues, manages, and validates requests involving identity claims. All implementations of AD FS require at least one Federation Service for each participating party.
Federation server proxy/Web Application Proxy	The federation server proxy is an optional component that you usually deploy in a perimeter network. It does not add any functionality to the AD FS deployment but is deployed to provide a layer of security for connections from the Internet to the federation server. In Windows Server 2016, the federation server proxy functionality is part of Web Application Proxy.
Claims	A claim is a statement that is made by a trusted entity about an object such as a user. The claim could include the user's name, job title, or any other factor that might be used in an authentication scenario.
Claim rules	Claim rules determine how claims are processed by federation servers. For example, a claim rule might state that an email address is accepted as a valid claim or that a group name from one organization is translated into an application-specific role in the other organization. The rules are usually processed in real time as claims are made.

Component	What it does
Attribute store	AD FS uses an attribute store to look up claim values. AD DS is a common attribute store and is available by default because the federation server role must be installed on a server joined to the domain.
Claims providers	The claims provider is the server that issues claims and authenticates users. A claims provider is one side of the AD FS authentication and authorization process. The claims provider manages user authentication and then issues the claims that the user presents to a relying party.
Relying parties	The relying party is the party in which the application is located and is the other side of the AD FS authentication and authorization process. The relying party is a web service that consumes claims from the claims provider. The relying party server must have the Windows Identity Foundation (WIF) installed or use the AD FS 1.0 claims-aware agent.
Claims provider trust	Configuration data that defines rules under which a client might request claims from a claims provider and subsequently submit them to a relying party. The trust consists of various identifiers such as names, groups, and various rules.
Relying party trust	The AD FS configuration data that is used to provide claims about a user or client to a relying party. It consists of various identifiers, such as names, groups, and various rules.
Certificates	AD FS uses digital certificates when communicating over Secure Sockets Layer (SSL) or as part of the token-issuing process, the token-receiving process, and the metadata-publishing process. Digital certificates also are used for token signing.
Endpoints	Endpoints are Windows Communication Foundation mechanisms that enable access to AD FS technologies, including token issuance and metadata publishing. AD FS comes with built-in endpoints that are responsible for specific functionality.

Prerequisites for an AD FS deployment

Before you deploy AD FS, you must ensure that your internal network meets some basic prerequisites. The configuration of the following network services is critical for a successful AD FS deployment:

- Network connectivity. The following network connectivity is required:
 - The client computer must be able to communicate with the web application, the resource federation server or federation server proxy, and the account federation server or federation server proxy by using HTTPS.
 - The federation server proxies must be able to communicate with the federation servers in the same organization by using HTTPS.

Successful AD FS deployment includes the following critical infrastructure:

- TCP/IP network connectivity
- AD DS
- Attribute stores
- DNS

- Federation servers and internal client computers must be able to communicate with domain controllers for authentication.
- AD DS. AD DS is a critical piece of AD FS. Federation servers must be joined to an AD DS domain. The Federation Service Proxy does not have to be joined to the domain.
- Attribute stores. AD FS uses an attribute store to build claims information. The attribute store contains information about users, which is extracted from the store by the AD FS server after the user has been authenticated.
- Domain Name System (DNS). Name resolution allows clients to find federation servers. Client computers must resolve DNS names for all federation servers or AD FS farms to which they connect and the web applications that the client computer is trying to use. If a client computer is external to the network, the client computer must resolve the DNS name for the Federation Service Proxy, not the internal federation server or AD FS farm. The Federation Service Proxy must resolve the name of the internal federation server or farm. If internal users must access the internal federation server directly and external users must connect through the federation server proxy, you will have to configure different DNS records in the internal and external DNS zones.

Public key infrastructure and certificate requirements

AD FS enables computers to communicate securely, even though they might be in different locations. In this scenario, most of the communications between computers pass through the Internet. To provide security for the network traffic, all communications are protected by using SSL. This factor means that it is important to correctly choose and assign SSL certificates to the AD FS servers. To provide SSL security, AD FS servers use certificates as service communication certificates, token-signing certificates, and token-decrypting certificates.

- Certificates used by AD FS:
 - Service communication certificates
 - Token-signing certificates
 - Token-decrypting certificates
- When choosing certificates, ensure that the service communication certificate is trusted by all federation partners and clients

Service communication certificates

AD FS secures all communication by using SSL, which requires a certificate. All computers that communicate with the AD FS server must trust the certificate used for service communication. If all of the computers and devices that contact your AD FS server are joined to the domain, you can consider using an internally generated certificate for AD FS. However, in most cases, at least some communication is between the AD FS server and external computers or partner organizations, in which case you should use a certificate from a third-party certification authority (CA). You can use the certificate's snap-in and the AD FS Management console to manage all certificates.



Note: If you change the service communication certificate after initial configuration, you must change it on all nodes in the server farm and ensure that the AD FS service is granted Read permissions to the private key on the certificate on each node.

Token-signing certificates

AD FS uses a token-signing certificate to sign every token that a federation server issues. This certificate is critical in an AD FS deployment, because the token signature indicates which federation server issued the token. The claims provider uses this certificate to identify itself, and the relying party uses it to verify that the token is coming from a trusted federation partner.

The relying party also requires a token-signing certificate to sign the tokens that it prepares for AD FS-aware applications. For the destination applications to validate these tokens, the relying party's token-signing certificate must validate these tokens.

When you configure a federation server, the server assigns a self-signed certificate as the token-signing certificate. In most cases, you do not need to update this certificate with a certificate from a third-party CA. When AD FS creates a federation trust, it configures the trust of this certificate at the same time. You can configure multiple token-signing certificates on the federation server, but AD FS uses only the primary certificate.

Token-decrypting certificates

AD FS uses token-decrypting certificates to encrypt the entire user token before transmitting the token across the network from the claims provider federation server to the relying party federation server. To provide this functionality, AD FS provides the public key from the relying party federation server certificate to the claims provider federation server. The certificate is sent without the private key. The claims provider server uses the public key from the certificate to encrypt the user token. When the claims provider server returns the token to the relying party federation server, it uses the private key from the certificate to decrypt the token. This provides an extra layer of security when transmitting the certificates across an untrusted network such as the Internet.

When you configure a federation server, the server assigns a self-signed certificate as the token-decrypting certificate. In most cases, you do not have to update this certificate with a certificate from a third-party CA. When AD FS creates a federation trust, it configures the trust of this certificate at the same time.



Note: The federation server proxies require only an SSL certificate. The federation server uses this certificate to enable SSL communication for all client connections.

Choosing a CA

AD FS federation servers can use self-signed certificates, certificates from an internal private CA, or certificates that have been purchased from an external public CA. In most AD FS deployments, the most important factor when choosing certificates is that they are trusted by all parties involved. This means that if you configure an AD FS deployment that interacts with other organizations, you almost certainly will use a public CA for the SSL certificate on a federation server proxy, because the certificates issued by the public CA are trusted by all partners automatically.

If you deploy AD FS just for your organization and all servers and client computers are under your control, you can consider using a certificate from an internal private CA. If you deploy an internal enterprise CA in Windows Server 2016, you can use Group Policy to ensure that all computers in the organization automatically trust the certificates issued by the internal CA. Using an internal CA can decrease the cost of certificates significantly.




Note: Deploying an internal CA by using Active Directory Certificate Services (AD CS) is a straightforward process, but it is critical that you plan and implement the deployment carefully.

AD FS server roles


In Windows Server 2016, the server roles for AD FS are:

- **Claims provider.** A claims provider is a federation server that provides users signed tokens containing claims. Claims provider federation servers are deployed in organizations where user accounts are located. When a user requests a token, the claims provider federation server verifies user authentication by using AD DS and then collects information from an attribute store, such as AD DS or Active Directory Lightweight Directory Services (AD LDS), to populate the user claim with the attributes required by the partner organization. The server issues tokens in the SAML format. The claims provider federation server also protects the contents of security tokens in transit by signing and optionally encrypting them.
- **Relying party.** A relying party is a federation server that receives security tokens from a trusted claims provider. Relying party federation servers are deployed in organizations that provide application access to claims provider organizations. The relying party accepts and validates the claim, and then it issues new security tokens that the web server can use to provide appropriate access to the application.

- **Claims provider federation server:**
 - Authenticates internal users
 - Issues signed tokens containing user claims
- **Relying party federation server:**
 - Consumes tokens from the claims provider
 - Issues tokens for application access
- **Federation service proxy:**
 - Is deployed in a perimeter network
 - Provides a layer of security for internal federation servers

 **Note:** A single AD FS server can operate as both a claims provider and a relying party, even with the same partner organizations. The AD FS server functions as a claims provider when it authenticates users and provides tokens for another organization, but it also can accept tokens from the same or different organizations in a relying party role.

- **Federation service proxy.** A federation service proxy provides an extra level of security for AD FS traffic that comes from the Internet to internal AD FS federation servers. Federation service proxies can be deployed in both claims provider and relying-party organizations. On the claims provider side, the proxy collects the authentication information from client computers and passes it to the claims provider federation server for processing. The federation server issues a security token to the proxy, which sends it to the relying party proxy. The relying party federation server proxy accepts these tokens and then passes them on to the internal federation server. The relying party federation server issues a security token for the web application and then sends the token to the federation server proxy, which then forwards the token to the client. The federation service proxy does not provide any tokens or create claims—it only forwards requests from clients to internal AD FS servers. All communication between the federation service proxy and the federation server uses HTTPS.

 **Note:** You cannot configure a federation service proxy as a claims provider or a relying party. The claims provider and relying party must be members of an AD DS domain. You can configure the federation service proxy as a member of a workgroup or as a member of an extranet forest, and you can deploy it in a perimeter network.

Demonstration: Installing the AD FS server role

In this demonstration, you will see how to:

- Install AD FS.
- Add a DNS record for AD FS.
- Configure AD FS.

Demonstration Steps

Install AD FS

1. On **LON-SVR2**, use **Server Manager** to install the **Active Directory Federation Services** role on **LON-SVR2.Adatum.com**.

Add a DNS record for AD FS

1. On **LON-DC1**, use **DNS Manager** to add a new host record for AD FS in the **Adatum.com** forward lookup zone with the following settings:
 - Name: **adfs**
 - IP address: **172.16.0.12**

Configure AD FS

1. On **LON-SVR2**, 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**
 - **Create a database on this server using Windows Internal Database**

Verify the correctness of the statement by placing a mark in the column to the right.

Statement	Answer
In Windows Server 2016, the federation server proxy functionality is part of the Web Application Proxy role.	

Lesson 3

Implementing AD FS for a single organization

The simplest deployment scenario for AD FS is within a single organization. In this scenario, a single AD FS server can operate both as the claims provider and as the relying party. All users in this scenario are internal to the organization, as is the application that the users access.

This lesson provides details on the components that are required to configure AD FS in a single-organization deployment of AD FS. These components include configuring claims, claim rules, claims provider trusts, and relying party trusts.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe AD FS claims.
- Describe AD FS claim rules.
- Describe a claims provider trust.
- Describe a relying party trust.
- Explain how to configure a claims provider and relying party trusts.

AD FS claims

An *AD FS claim* is a statement that a trusted entity, such as a claims provider, makes about a particular subject, such as a user. AD FS claims provide the link between the claims provider and relying party roles in an AD FS deployment. The claims provider creates the claims, and the relying party consumes the claims. AD FS claims provide a standards-based and flexible way for claims provider organizations to provide specific information about users in their organizations. AD FS claims also provide a way for relying parties to define exactly what information they require to provide application access. The claim information provides the details required by applications to enable access to claims-aware applications.

- Claims provide information about users from the claims provider to the relying party
- AD FS:
 - Provides a default set of built-in claims
 - Enables the creation of custom claims
 - Requires each claim have a unique URI
- Claims can be:
 - Retrieved from an attribute store
 - Calculated based on retrieved values
 - Transformed into alternate values

Claim types

Each AD FS claim has a claim type, such as email address, UPN, or last name. Users can be issued claims based on any defined claim type. Therefore, a user might be issued a claim with a type of Last Name and a value of Weber, for example. AD FS provides many built-in claim types. You also can opt to create new ones based on organizational requirements.

A Uniform Resource Identifier (URI) uniquely identifies each AD FS claim type. This information is provided as part of the AD FS server metadata. For example, if the claims provider organization and the relying party organization decide to use a claim type of AccountNumber, both organizations must configure a claim type with this name. The claim type is published, and the claim type URI must be identical on both AD FS servers.

How claim values are populated

Claims issued by a claims provider contain the information that is required by the relying party to enable appropriate application access. One of the first steps in planning an AD FS deployment is to define exactly what information the applications must have about each user to provide that user access to the application. Once you define this information, the claims then are defined on the claims provider federation server. The AD FS server can obtain the information required to populate the claim in several ways:

- It can retrieve the claim from an attribute store. Frequently, the information required for the claim is already stored in an attribute store that is available to the federation server. For example, an organization might decide that the claim should include the user's UPN, email address, and specific group memberships. This information is stored in AD DS already, so the federation server can retrieve this information from AD DS when creating the claim. Because AD FS can use AD DS, AD LDS, SQL Server, a non-Microsoft LDAP directory, or a custom attribute store to populate claims, you can define almost any value within the claim.
- It can calculate the claim based on collected information. Claims provider federation servers also can calculate information based on data that is gathered from an attribute store. For example, you might want to provide information about a person's salary within a claim. This information likely is stored in a Human Resources database, but the actual value might be considered confidential. You can define a claim that categorizes salaries within an organization and then have the AD FS server calculate in which category a specific user belongs. In this way, the claim only includes the salary category information, not the actual user's salary value.
- AD FS can transform the claim from one value to another. In some cases, the information that is stored in an attribute store does not exactly match the information required by the application when making authorization information. For example, the application might define different user roles that do not directly match the attributes that are stored in any attribute store. However, the application role might correlate to the AD DS group membership. For example, users in the Sales group might correlate to one application role, while users in the Sales Management group might correlate to a different application role. To establish the correlation in AD FS, you can configure a claims transformation that takes the value provided by the claims provider and translates the value into to a claim that is useful to the application in the relying party.
- If you have deployed DAC, AD FS can transform a DAC device claim into an AD FS claim. This ensures that users can access an AD FS website only from trusted workstations that have been issued a valid device claim.

AD FS claim rules

Claim rules define how claims are sent and consumed by AD FS servers. Claim rules define the business logic that is applied to claims that are provided by claims providers and to claims that are accepted by the relying parties. You can use claim rules to:

- Define which incoming claims are accepted from one or more claims providers.
- Define which outbound claims are provided to one or more relying parties.

- Claim rules define how claims are sent and consumed by AD FS servers
- Claims provider rules are acceptance transform rules
- Relying party rules can be:
 - Issuance transform rules
 - Issuance authorization rules
 - Delegation authorization rules
- AD FS servers provide default claim rules, templates, and a syntax for creating custom claim rules

- Apply authorization rules to enable access to a specific relying party for one or more users or groups of users.

You can define two types of claim rules:

- Claim rules for a claims provider trust. A claims provider trust is the AD FS trust relationship that is configured between an AD FS server and a claims provider. You can configure claim rules to define how the claims provider processes and issues claims.
- Claim rules for a relying party trust. A relying party trust is the AD FS trust relationship that is configured between an AD FS server and a relying party. You can configure claim rules that define how the relying party accepts claims from the claims provider.

Claim rules configured on an AD FS claims provider all are considered acceptance transform rules. These rules determine what claim types are accepted from the claims provider and are then sent to a relying party trust. When configuring AD FS within a single organization, a default claims provider trust is configured with the local AD DS domain. This rule set defines the claims that are accepted from AD DS.

There are three types of claim rules for a relying party trust:

- Issuance transform rules. These rules define the claims that are sent to the relying party that was defined in the relying party trust.
- Issuance authorization rules. These rules define which users are permitted or denied access to the relying party defined in the relying party trust. This rule set can include rules that explicitly permit access to a relying party and rules that explicitly deny access to a relying party.
- Delegation authorization rules. These rules define the claims that specify which users can act on behalf of other users when accessing the relying party. This rule set can include rules that explicitly permit delegates for a relying party or rules that explicitly deny delegates to a relying party.



Note: A single claim rule can be associated only with a single federated trust relationship. This means that you cannot create a set of rules for one trust and then reuse those rules for other trusts that you configure on your federation server.

AD FS servers are preconfigured with a set of default rules and several default templates that you can use to create common claim rules. You can create custom claim rules by using the AD FS claim rule language.

Claims provider trust

A *claims provider trust* is configured on the relying party federation server. The claims provider trust identifies the claims provider and describes how the relying party consumes the claims that the claims provider issues. You must configure a claims provider trust for each claims provider. A claims provider trust for the local AD DS is configured by default. You must configure any additional claims providers.

By default, an AD FS server is configured with a claims provider trust named Active Directory. This trust defines the claim rules, which are all

acceptance transform rules that define how the AD FS server accepts AD DS credentials. For example, the default claim rules on the claims provider trust include rules that transmit user names, security identifiers

- Claims provider trusts:
 - Are configured on the relying party federation server
 - Identify the claims provider
 - Configure the claim rules for the claims provider
- In a single-organization scenario, a claims provider trust called Active Directory defines how AD DS user credentials are processed
- Claims provider trusts can be configured by:
 - Importing the federation metadata
 - Importing a configuration file
 - Configuring the trust manually

(SIDs), and group SIDs to the relying party. In a single-organization AD FS deployment where AD DS authenticates all users, the default claims provider trust might be the only required claims provider trust.

When you expand an AD FS deployment to include other organizations, you must create additional claims provider trusts for each federated organization that is an identity provider. When configuring a claims provider trust, you have three options:

- Import data about the claims provider through the federation metadata. If the AD FS federation server or federation server proxy is accessible through the network from your AD FS federation server, you can enter the host name or URL for the partner federation server. Your AD FS federation server connects to the partner server and downloads the federation metadata from the server. The federation metadata includes all the information that is required to configure the claims provider trust. As part of the federation metadata download, your federation server also downloads the SSL certificate that is used by the partner federation server.
- Import data about the claims provider from a file. Use this option if the partner federation server is not directly accessible from your federation server, but the partner organization has exported its configuration and provided you the information in a file. The configuration file must include configuration information for the partner organization and the SSL certificate that the partner federation server uses.
- Manually configure the claims provider trust. Use this option if you want to configure all of the settings for the claims provider trust. When you choose this option, you must provide the features that the claims provider supports and the URL that is used to access the claims provider AD FS servers. You also must add the SSL certificate that the partner organization uses.

Relying party trust

You define a *relying party trust* on the claims provider federation server. The relying party trust identifies the relying party and also defines the claim rules that define how the relying party accepts and processes claims from the claims provider.

In a single-organization scenario, the relying party trust defines how the AD FS server interacts with the applications deployed within the organization.

When you configure the relying party trust in a single organization, you provide the URL for the internal application. You can also configure settings such as the URL used by the web server, the issuance authorization rules for the application, and whether the application supports SAML 2.0 or requires AD FS 1.0 tokens.

- Relying party trusts:
 - Are configured on the claims provider federation server
 - Identify the relying party
 - Configure the claim rules for the relying party
- In a single-organization scenario, a relying party trust defines the connection to internal applications
- You can configure relying party trusts by:
 - Importing the federation metadata
 - Importing a configuration file
 - Manually configuring the trust

Configuring a relying party trust is similar to configuring a claims provider trust. When you expand the AD FS deployment to include other organizations, you must create additional relying party trusts for each federated organization. When configuring a relying party trust, you have three options:

- Import data about the relying party through the federation metadata. If the AD FS federation server or federation server proxy is accessible through the network from your AD FS federation server, you can enter the host name or URL for the partner federation server. Your AD FS federation server connects to the partner server and then downloads the federation metadata from the server. The federation metadata includes all the information that is required to configure the relying party trust. As part of the federation metadata download, your federation server also downloads the SSL certificate that the partner federation server uses.

- Import data about the relying party from a file. Use this option if the partner federation server is not accessible from your federation server directly. In this case, the partner organization can export its configuration information to a file and then provide it to you. The configuration file must include configuration information for the partner organization and the SSL certificate that the partner federation server uses.
- Manually configure the claims provider trust. Use this option if you want to configure all of the settings for the claims provider trust.

Demonstration: Configuring claims provider and relying party trusts

In this demonstration, you will see how to:

- Configure a claims provider trust.
- Configure a WIF application for AD FS.
- Configure a relying party trust.

Demonstration Steps

Configure a claims provider trust

1. On **LON-SVR2**, in **Server Manager**, open the **AD FS Management** tool.
2. Browse to **Claims Provider Trusts**, and then edit claim rules for **Active Directory**.
3. Add an acceptance transform rule with the following settings:
 - Claim rule template: **Send LDAP Attributes as Claims**
 - Claim rule name: **Outbound LDAP Attributes Rule**
 - Attribute store: **Active Directory**
 - Mapping of LDAP attributes:
 - E-Mail-Addresses: **E-Mail Address**
 - User-Principal-Name: **UPN**

Configure a WIF application for AD FS

1. On **LON-SVR1**, in **Server Manager**, open the **Windows Identity Foundation Federation Utility** tool.
2. Enter the following in the **Federation Utility Wizard**:
 - Application configuration location: **C:\inetpub\wwwroot\AdatumTestApp\web.config**
 - Application URI: **https://lon-svr1.adatum.com/AdatumTestApp/**
 - **Use an existing STS**
 - STS WS-Federation metadata document location:
https://adfs.adatum.com/federationmetadata/2007-06/federationmetadata.xml
 - **Disable certificate chain validation,**
 - **No encryption**

Configure a relying party trust

1. On **LON-SVR2**, in the AD FS console, add a **Relying Party Trust** with the following settings:
 - Import data about the relying party published online or on a local network
 - Federation Metadata address: **<https://lon-svr1.adatum.com/adatumtestapp/>**
 - Display name: **A. Datum Test App**
 - **Permit everyone**
2. Leave the **Edit Claims Issuance Policy for A. Datum Test App** window open for the next task. (This might be hidden behind Server Manager.)
3. On **LON-SVR2**, in the **Edit Claim Issuance Policy for A. Datum Test App** window, add a rule on the **Issuance Transform Rules** tab.
4. Complete the **Add Transform Claim Rule Wizard** with the following settings:
 - Claim rule template: **Pass Through or Filter an Incoming Claim**
 - Claim rule name: **Pass through Windows account name**
 - Incoming claim type: **Windows account name**
 - **Pass through all claim values**
5. Create three more rules to pass through the **E-Mail Address**, **UPN**, and **Name** claim types.

Question: What are claim rules? What can you use claim rules for?

Lab A: Implementing AD FS

Scenario

A. Datum Corporation plans to implement AD FS. In the initial deployment, the company plans to use AD FS to implement SSO for internal users who access an application on a web server. As one of the senior network administrators at A. Datum, it is your responsibility to implement this AD FS solution. As a proof of concept, you plan to deploy a sample claims-aware application, and configure AD FS to enable internal users to access the application.

Objectives

After completing this lab, you will be able to:

- Install and configure AD FS.
- Configure an internal application for AD FS.

Lab Setup

Estimated Time: 55 minutes

Virtual machines: **20743A-LON-DC1**, **20743A-LON-SVR1**, **20743A-LON-SVR2**, **20743A-LON-SVR3**, **20743A-LON-CL1**, and **20743A-LON-CL3**.

User name: **Adatum\Administrator**

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

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

1. On the host computer, start **Hyper-V Manager**
2. In Hyper-V Manager, click **20743A-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 using the following credentials:
 - o User name: **Administrator**
 - o Password: **Pa\$\$w0rd**
 - o Domain: **Adatum**
5. Repeat steps 2 through 4 for **20743A-LON-SVR1**, **20743A-LON-SVR2**, and **20743A-LON-SVR3**.
6. Repeat steps 2 and 3 for **20743A-LON-CL1** and **20743A-LON-CL3**.

Exercise 1: Installing and configuring AD FS

Scenario

To start the AD FS implementation, you must install AD FS on an A. Datum server. During the initial deployment, you will configure it as the first server in a farm with the option to expand the farm later. The certificate for AD FS has been installed on LON-SVR2.

The main tasks for this exercise are as follows:

1. Create a DNS record for AD FS.
2. Install AD FS.
3. Configure AD FS.
4. Verify AD FS functionality.

► Task 1: Create a DNS record for AD FS

1. On **LON-DC1**, from **Server Manager**, open **DNS Manager** and add a new host record for AD FS:
 - Forward lookup zone: **Adatum.com**
 - Name: **adfs**
 - IP address: **172.16.0.12**

► Task 2: Install AD FS

- On **LON-SVR2**, open **Server Manager**, and then add the **Active Directory Federation Services** role.

► Task 3: Configure AD FS

1. On **LON-SVR2**, in **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**
 - **Create a database on this server using Windows Internal Database**



Note: The adfs.adatum.com certificate was preconfigured for this task. In your own environment, you must obtain this certificate.

► Task 4: Verify AD FS functionality

1. On **LON-CL1**, sign in as **Adatum\Beth** with the password **Pa\$\$w0rd**.
2. Open **Internet Explorer**, and then access **https://adfs.adatum.com/federationmetadata/2007-06/federationmetadata.xml**
3. Verify that the file loads, and then close Internet Explorer.

Results: After completing this exercise, you will have installed and configured AD FS. You also will have verified that it is functioning by viewing the FederationMetaData.xml file contents.

Exercise 2: Configuring an internal application for AD FS

Scenario

The first scenario for implementing the proof-of-concept AD FS application is to ensure that internal users can use SSO to access the web application. You must plan to configure the AD FS server and a web application to enable this scenario. You also must verify that internal users can access the application.

The main tasks for this exercise are as follows:

1. Configure the Active Directory claims provider trust.
2. Configure the application to trust incoming claims.
3. Configure a relying party trust for the claims-aware application.
4. Configure claim rules for the relying party trust.
5. Test access to the claims-aware application.
6. Configure Internet Explorer to pass local credentials to the application automatically.
7. Prepare for the next lab.

► Task 1: Configure the Active Directory claims provider trust

1. On **LON-SVR2**, in **Server Manager**, open the **AD FS Management** tool.
2. Browse to the **Claims Provider Trusts**, and then edit claim rules for **Active Directory**.
3. Add an acceptance transform rule with the following settings:
 - Claim rule template: **Send LDAP attributes as claims**
 - Name: **Outbound LDAP Attributes Rule**
 - Attribute store: **Active Directory**
 - Mapping of LDAP attributes to outgoing claim types:
 - E-Mail-Addresses: **E-Mail Address**
 - User-Principal-Name: **UPN**
 - Display-Name: **Name**

► **Task 2: Configure the application to trust incoming claims**

- On **LON-SVR1**, open **Server Manager**, and then open the **Windows Identity Foundation Federation Utility** tool.
- Enter the following in the **Federation Utility Wizard**:
 - Application configuration location: **C:\inetpub\wwwroot\AdatumTestApp\web.config**
 - Application URI: **https://lon-svr1.adatum.com/AdatumTestApp/**
 - **Use an existing STS**
 - STS WS-Federation metadata document location: **https://adfs.adatum.com/federationmetadata/2007-06/federationmetadata.xml**
 - **Disable certificate chain validation**
 - **No encryption**

► **Task 3: Configure a relying party trust for the claims-aware application**

1. On **LON-SVR2**, in the **AD FS** console, add a **Relying Party Trust** with the following settings:
 - **Import data about the relying party published online or on a local network**
 - Federation Metadata address: **https://lon-svr1.adatum.com/adatumtestapp/**
 - Display name: **A. Datum Test App**
 - **Permit everyone**
2. Leave the **Edit Claims Issuance Policy for A. Datum Test App** window open for the next task. (This might be hidden behind Server Manager.)

► **Task 4: Configure claim rules for the relying party trust**

1. On **LON-SVR2**, in the **Edit Claim Issuance Policy for A. Datum Test App** window, add a rule on the **Issuance Transform Rules** tab.
2. Complete the **Add Transform Claim Rule Wizard** with the following settings:
 - Claim rule template: **Pass Through or Filter an Incoming Claim**
 - Claim rule name: **Pass through Windows account name**
 - Incoming claim type: **Windows account name**
 - **Pass through all claim values**
3. Create three more rules to pass through the **E-Mail Address**, **UPN**, and **Name** claim types.

► **Task 5: Test access to the claims-aware application**

1. On **LON-CL1**, use **Internet Explorer** to access **https://lon-svr1.adatum.com/AdatumTestApp/**.



Note: It is critical to use the trailing slash in the URL for step 1.

2. When you receive a prompt, sign in as **Adatum\Beth** with the password **Pa\$\$w0rd**.
3. Review the claim information that is displayed by the application, and then close **Internet Explorer**.

► **Task 6: Configure Internet Explorer to pass local credentials to the application automatically**

1. On **LON-CL1**, in **Internet Explorer**, open **Internet Options**.
2. On the **Security** tab, add the following sites to the **Local intranet** zone:
 - **https://adfs.adatum.com**
 - **https://lon-svr1.adatum.com**
3. Use **Internet Explorer** to access **https://lon-svr1.adatum.com/AdatumTestApp/**.



Note: It is critical to use the trailing slash in the URL for step 3.

4. Notice that you did not receive a prompt for credentials.
5. Review the claim information that is displayed by the application, and then close **Internet Explorer**.

► **Task 7: Prepare for the next lab**

- Leave all the virtual machines running at the end of this lab.

Results: After completing this exercise, you will have configured AD FS to support authentication for an application.

Question: Why is it important to configure adfs.adatum.com to use as a host name for the AD FS service?

Question: How can you test whether AD FS is functioning properly?

Lesson 4

Implementing Web Application Proxy

Many organizations need to extend the AD FS infrastructure beyond private networks and onto the Internet. To enhance security for AD FS and AD FS applications, use Web Application Proxy. It also is important to consider high availability for AD FS, because it is a critical service once it is in place.


Lesson Objectives

After completing this lesson, you will be able to:

- Describe the new features in Web Application Proxy in Windows Server 2016.
- Describe how to configure an application for Web Application Proxy.
- Describe Web Application Proxy and AD FS.
- Explain how to install and configure Web Application Proxy.


What is new in Web Application Proxy?

The Web Application Proxy role service is a reverse web proxy and provides access to internal organizational web applications for remote users that connect to your organization's network.

 **Note:** Web Application Proxy uses AD FS to preauthenticate Internet users; it acts as an AD FS proxy for publishing claims-aware applications.

AD FS provides users with SSO capability, which enables users to enter their credentials to access an organizational web application without receiving a prompt to enter their credentials again. With Web Application Proxy, you can publish both claims-aware applications that use AD FS preauthentication and web applications that use pass-through preauthentication.

Usually, you place the Web Application Proxy in your perimeter network between two firewall devices.

 **Note:** The AD FS server and applications that are published are located in the organizational network with domain controllers and other internal servers, and they are protected by the second firewall. This scenario helps provide secure access to organizational applications for users on the Internet. At the same time, this scenario helps protect the organization's IT infrastructure from security threats from the Internet.

Windows Server 2016 includes several improvements to the Web Application Proxy role, including:

- Preauthentication for HTTP Basic application publishing
- Wildcard domain publishing of applications
- HTTP to HTTPS redirection
- HTTP Publishing

Improvements in Web Application Proxy in Windows Server 2016

Windows Server 2016 includes a number of improvements to the Web Application Proxy role, including:

- Preauthentication for HTTP Basic application publishing. HTTP Basic is the authorization protocol that is used by many protocols, including Exchange ActiveSync, to connect devices, including smartphones, with Exchange Server mailboxes.



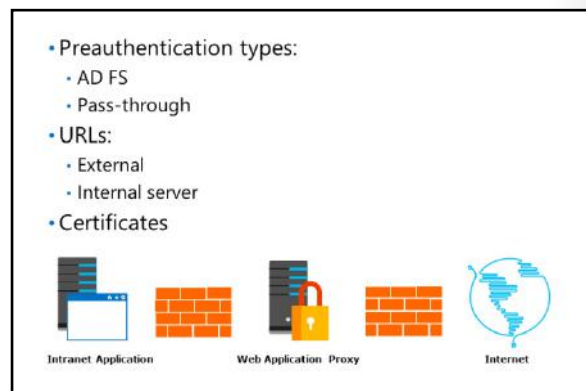
Note: Web Application Proxy interacts with AD FS using redirections, which is not supported on ActiveSync clients.

Web Application Proxy in Windows Server 2016 enables you to publish an app using HTTP Basic by enabling the HTTP app to receive a non-claims relying party trust for the application to the Federation Service.

- Wildcard domain publishing of applications simplifying the publishing of Microsoft Office SharePoint apps. To support scenarios such as SharePoint 2013, the external URL for the application can now include a wildcard that enables you to publish multiple applications from within a specific domain (for example, https://*.sp-apps.adatum.com).
- HTTP to HTTPS redirection. To ensure that your users can access your app, even if they neglect to type HTTPS in the URL, Web Application Proxy in Windows Server 2016 now supports HTTP to HTTPS redirection.
- HTTP Publishing. You can now publish HTTP applications using pass-through preauthentication.

Configuring an application

You should place the Web Application Proxy server in a perimeter network, where it protects web applications and AD FS when they are accessible from the Internet. To install Web Application Proxy, you must already have implemented AD FS in your organization. All configuration information for Web Application Proxy is stored in AD FS. To use Web Application Proxy as a reverse proxy for web applications, you must configure each application. For each application, you need to configure the type of preauthentication for the application and URLs.



Pass-through preauthentication

When you use pass-through preauthentication, no preauthentication is performed, and valid requests are passed to web-based applications on an internal network without performing authentication on a user. The application performs all authentication for an application only after a user is connected. You can use pass-through preauthentication for any web application.

A web application protected by preauthentication is protected from malformed packets that could cause a denial-of-service attack. However, the web application would not be protected from application-level threats where the application mishandles valid data. For example, an HTTPS request with valid HTTP commands would be passed through to the application even if the actions requested by the HTTP commands might cause the web application to fail.

AD FS preauthentication

You can configure Web Application Proxy to use AD FS preauthentication or pass-through authentication. When you use AD FS for preauthentication, AD FS authenticates a user request before passing it to an internal, web-based application. This ensures that only authorized users can send data to a web-based application. AD FS preauthentication provides a higher level of protection than pass-through authentication, because unauthenticated users cannot submit requests to the application.

Only a claims-aware application that uses AD FS for authentication can use AD FS preauthentication. You must configure the claims-aware application in AD FS as a relying party and select it from a list when Web Application Proxy is configured. Web Application Proxy is aware of the relying parties configured in AD FS because of the integration between AD FS and Web Application Proxy.

URLs

For each application that you publish, you must configure an external URL and internal server URL. External users use the external URL when they access the application. The Web Application Proxy server uses the internal server URL to access the application on behalf of external users.

If you use split DNS, it is common to have the same value for both the external URL and the internal server URL. Some applications experience errors when the external URL and the internal server URL are different. When the external URL and the backed server URL are different, only the host name in the URL can change. The path to the application must remain the same. For example, if the internal URL for an application is `https://server1.adatum.com/app1`, you cannot have an external URL of `https://extranet.adatum.com/application1`.

Certificates

When you define the external URL, you also must select a certificate that contains the host name in the external URL. This certificate must be installed on the local server. However, it does not need to match the certificate used on the backend server hosting the application. You can have one certificate for each host name used on the Web Application Proxy server or a single certificate with multiple names.

Web Application Proxy and AD FS proxy

Many organizations need to provide authentication for users and devices that are located on a network that is external to the organization. In most cases, allowing clients to access an AD FS server located on an internal network directly from the Internet is an unacceptable security risk. We strongly recommend an AD FS proxy to allow clients on the Internet to access AD FS.

An AD FS proxy is a reverse proxy, located in a perimeter network that is specifically for AD FS. Clients from the Internet communicate with the AD FS proxy in the perimeter network instead of directly with the AD FS server. The AD FS proxy mitigates the risks associated with Internet connectivity for AD FS.

- Web Application Proxy is an AD FS proxy
- The same certificate is used on the AD FS server and Web Application Proxy
- Split DNS allows the same name to resolve to different IP addresses



Authentication process

An internal AD FS server uses Windows authentication to prompt for authentication. This works well for internal computers that are joined to the domain and can pass workstation credentials automatically to AD FS to automate authentication. This prevents users from seeing a request for authentication credentials.

When computers that are not joined to the domain communicate with AD FS, the web browser presents the users with a logon prompt. This logon prompt asks for a user name and password but provides no context.

When you use an AD FS proxy, an authentication web page is provided for computers that are not joined to the domain. This provides better compatibility than browser-based Windows authentication for AD FS clients that use non-Microsoft operating systems. You also can customize the web page to provide more context for users, by adding a company logo, for example.

DNS resolution

To provide seamless movement between internal and external networks, Web Application Proxy uses the same host name when accessing AD FS internally and externally. On the internal network, the AD FS host name resolves to the IP address of the internal AD FS server. On the external network, the AD FS host name resolves to the IP address of the AD FS proxy. In both cases, the AD FS host name is different from the computers that host the AD FS roles.

Certificates

The certificate used on an internal AD FS server has a subject name that is the same as the host name for AD FS (for example, adfs.adatum.com). Because the same host name is used to access AD FS internally and externally through the AD FS proxy, you must configure the AD FS proxy with the same certificate as the AD FS server. If the certificate subject does not match the host name, AD FS authentication fails.



Note: To ensure that you have a certificate with the same subject name, export the certificate from the AD FS server and import it on the Web Application Proxy server. Remember to include the private key when you export the certificate.

Demonstration: Installing and configuring the Web Application Proxy

In this demonstration, you will see how to:

- Install Web Application Proxy.
- Export the certificate from the AD FS server.
- Import the certificate to the Web Application Proxy server.
- Configure Web Application Proxy.

Demonstration Steps

Install Web Application Proxy

1. On **LON-SVR3**, in **Server Manager**, add the **Remote Access** server role and the **Web Application Proxy** role service.

Export the adfs.adatum.com certificate from LON-SVR2

1. On **LON-SVR2**, open a **Microsoft Management** console, and then add the **Certificates** snap-in for the **Local Computer**.
2. From the **Personal** folder, export the **adfs.adatum.com** certificate:
 - o **Yes, export the private key**
 - o File format: **Personal Information Exchange – PKCS #12 (.PFX)**
 - o Password: **Pa\$\$w0rd**
 - o File name: **C:\adfs.pfx**

Import the adfs.adatum.com certificate on LON-SVR3

1. On **LON-SVR3**, open a **Microsoft Management** console, and then add the **Certificates** snap-in for the **Local Computer**.
2. From the **Personal** folder, import the **adfs.adatum.com** certificate.
 - o File name: **\\LON-SVR2\c\$\adfs.pfx**
 - o Password: **Pa\$\$w0rd**
 - o **Mark this key as exportable. This will allow you to back up or transport your keys at a later time**
 - o Certificate store: **Personal**

Configure Web Application Proxy

1. On **LON-SVR3**, in **Server Manager**, click the **Notifications** icon, and then click **Open the Web Application Proxy Wizard**.
2. In the **Web Application Proxy Wizard**, provide the following configuration settings:
 - o Federation service name: **adfs.adatum.com**
 - o User name: **Adatum\Administrator**
 - o Password: **Pa\$\$w0rd**
 - o Certificate to be used by the AD FS proxy: **adfs.adatum.com**

Check Your Knowledge

Question	
Which of the following statements about configuring Web Application Proxy is true? (Choose all that apply.)	
Select the correct answer.	
<input type="checkbox"/>	To install Web Application Proxy, you must have implemented AD FS in your organization.
<input type="checkbox"/>	To install Web Application Proxy, you need not have implemented AD FS in your organization.
<input type="checkbox"/>	For each application that you publish, you must configure an external URL and an internal server URL.

Question	
	When you define the external URL, you must also select a certificate that contains the host name in the internal URL.
	When you define the external URL, you must also select a certificate that contains the host name in the external URL.

Lab B: Implementing Web Application Proxy

Scenario

A. Datum plans to implement AD FS. You have successfully implemented AD FS to support an internal application. Now you must deploy and configure the Web Application Proxy to support remote clients.

Objectives

After completing this lab, you will be able to:

- Implement Web Application Proxy.

Lab Setup

Estimated Time: 20 minutes

Virtual machines: **20743A-LON-DC1**, **20743A-LON-SVR1**, **20743A-LON-SVR2**, **20743A-LON-SVR3**, **20743A-LON-CL1**, and **20743A-LON-CL3**.

User name: **Adatum\Administrator**

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

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

1. On the host computer, start **Hyper-V Manager**.
2. In Hyper-V Manager, click **20743A-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 using the following credentials:
 - User name: **Administrator**
 - Password: **Pa\$\$w0rd**
 - Domain: **Adatum**
5. Repeat steps 2 through 4 for **20743A-LON-SVR1**, **20743A-LON-SVR2**, and **20743A-LON-SVR3**.
6. Repeat steps 2 and 3 for **20743A-LON-CL1** and **20743A-LON-CL3**.



Note: These virtual machines should still be running from the previous lab. You must complete the previous lab before you start this lab.

Exercise 1: Implementing Web Application Proxy

Scenario

The final scenario for implementing the proof-of-concept AD FS application is to increase security for AD FS authentication by implementing an AD FS proxy for the AD FS and a reverse proxy for the application. You will implement Web Application Proxy to fulfill both of these roles.

The main tasks for this exercise are as follows:

1. Install Web Application Proxy.
2. Add the adfs.adatum.com certificate to **LON-SVR3**.
3. Add the **LON-SVR1.adatum.com** certificate to **LON-SVR3**.

4. Configure Web Application Proxy.
5. Configure the AD FS Proxy for the test application.
6. Test Web Application Proxy.
7. Prepare for the next module.

► **Task 1: Install Web Application Proxy**

- On **LON-SVR3**, open **Server Manager**, add the **Remote Access** server role and the **Web Application Proxy** role service.

► **Task 2: Add the adfs.adatum.com certificate to LON-SVR3**

1. On **LON-SVR2**, open a **Microsoft Management** console, and then add the **Certificates** snap-in for the **Local Computer**.
2. From the **Personal** folder, export the **adfs.adatum.com** certificate:
 - **Yes, export the private key**
 - File format: **Personal Information Exchange – PKCS #12 (.PFX)**
 - Password: **Pa\$\$w0rd**
 - File name: **C:\adfs.pfx**
3. On **LON-SVR3**, open a **Microsoft Management** console, and then add the **Certificates** snap-in for the **Local Computer**.
4. From the **Personal** folder, import the **adfs.adatum.com** certificate:
 - File name: **\\LON-SVR2\c\$\adfs.pfx**
 - Password: **Pa\$\$w0rd**
 - Certificate store: **Personal**

► **Task 3: Add the LON-SVR1.adatum.com certificate to LON-SVR3**

1. On **LON-SVR1**, open **Microsoft Management** console, and then add the **Certificates** snap-in for the **Local Computer**.
2. From the **Personal** folder, export the **lon-svr1.adatum.com** certificate:
 - **Yes, export the private key**
 - File format: **Personal Information Exchange – PKCS #12 (.PFX)**
 - Password: **Pa\$\$w0rd**
 - File name: **C:\lon-svr1.pfx**
3. On **LON-SVR3**, switch to **Console1**.
4. From the **Personal** folder, import the **lon-svr1.adatum.com** certificate:
 - File name: **\\LON-SVR1\c\$\lon-svr1.pfx**
 - Password: **Pa\$\$w0rd**
 - Certificate store: **Personal**

► Task 4: Configure Web Application Proxy


1. On **LON-SVR3**, in **Server Manager**, click the **Notifications** icon, and then click **Open the Web Application Proxy Wizard**.
2. In the **Web Application Proxy Configuration** Wizard, provide the following configuration settings:
 - Federation service name: **adfs.adatum.com**
 - User name: **Adatum\Administrator**
 - Password: **Pa\$\$w0rd**
 - Certificate to be used by the AD FS proxy: **adfs.adatum.com**
3. Leave the **Remote Access Management** console open for the next task.

► Task 5: Configure the AD FS Proxy for the test application


- On **LON-SVR3**, in the **Remote Access Management** console, publish a new application with the following settings:
 - Preauthentication: **Active Directory Federation Services (AD FS)**
 - Supported clients: **Web and MSOFBA**
 - Relying party: **A. Datum Test App**
 - Name: **A. Datum Test App Rule**
 - External URL: **https://lon-svr1.adatum.com/adatumtestapp/**
 - External certificate: **lon-svr1.adatum.com**
 - Backend server URL: **https://lon-svr1.adatum.com/adatumtestapp/**

► Task 6: Test Web Application Proxy

1. Switch to **LON-CL3**, and sign in as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. On **LON-CL3**, use Notepad to add the following lines to **C:\Windows\System32\Drivers\etc\hosts**:
 - **172.16.0.13 adfs.adatum.com**
 - **172.16.0.13 lon-svr1.adatum.com**

 **Note:** You edit the hosts to force **LON-CL3** to access the application through Web Application Proxy. In a production environment, you do this by using split DNS. The IPv4 address is for **LON-SVR3**, the Web Application Proxy.

3. Use Internet Explorer to access **https://lon-svr1.adatum.com/adatumtestapp/**.

 **Note:** If you receive two certificate errors, ignore them.

4. Sign in as **Adatum\Beth** with the password **Pa\$\$w0rd**.

► Task 7: Prepare for the next module

When you have finished the lab, revert the virtual machines to their initial state.

1. On the host computer, start **Hyper-V Manager**.
2. In the **Virtual Machines** list, right-click **20743A-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20743A-LON-SVR1**, **20743A-LON-SVR2**, **20743A-LON-SVR3**, **20743A-LON-CL3**, and **20743A-LON-CL1**.

Results: After completing this exercise, you will have configured Web Application Proxy to secure access to AdatumTestApp from the Internet.

Question: In the lab, you received a certificate error when connecting from LON-CL3 to the A. Datum Test App. Why did this error occur, and what can you do to avoid this?

Lesson 5

Implementing SSO with Microsoft Online Services

As organizations move services and applications to cloud-based services, it is increasingly important to provide a simple authentication and authorization experience to their users as they consume cloud-based services. Cloud-based services add another level of complexity to the IT environment, because they are located outside the direct administrative control of IT administrators, and they can run on many different platforms. This lesson identifies how you can use AD FS to support SSO with online services and provide high-level guidance on configuration steps for a typical scenario.

Lesson Objectives

After completing this lesson, you will be able to:

- Explain when to consider using AD FS with Microsoft online services.
- Explain how to configure SSO with Microsoft online services.

AD FS and SSO with online services

You can use AD FS to provide an SSO experience to users across various available cloud-based platforms. For example, once users authenticate with AD DS credentials, using those domain credentials, they then can access Microsoft online services, such as Microsoft Azure, Microsoft Intune, or Microsoft Office 365, if they use those domain credentials.



Note: AD FS also can provide SSO to other cloud providers. Because AD FS is based on open standards, it can interoperate with any compliant claims-based system.

A hybrid Microsoft Exchange deployment is an example of a cloud-based service that uses AD FS for authentication. In this type of deployment, an organization deploys some or all of its mailboxes in an Office 365 environment. However, the organization manages all of its user accounts in its on-premises AD DS environment. The deployment uses a directory synchronization tool to synchronize user account information from the on-premises environment to the Office 365 deployment.

When users try to sign in to their Office 365 mailbox, they must authenticate by using their internal AD DS credentials. If users try to sign in directly to the Office 365 environment, they are redirected back to the internal AD FS deployment to authenticate before they are given access.

The following steps describe what happens when a user tries to access his or her online mailbox by using a web browser:

1. The user opens a web browser and sends an HTTPS request to the Office 365 Outlook on the Web server.
2. The Outlook on the Web server receives the request and verifies whether the user is part of a hybrid Exchange Server deployment. If this is the case, the server redirects the client computer to the Azure AD Authentication System.
3. The client computer sends an HTTPS request to the Azure AD Authentication System.
4. The client computer is redirected again to the on-premises federation server. The redirection to the user's home domain is based on the UPN suffix of the user.
5. The client computer sends an HTTPS request to the on-premises federation server.

6. If the user is logged on to the domain already, the on-premises federation server can take the user's Kerberos ticket and request authentication from AD DS on the user's behalf by using Windows authentication. If the user logs on from outside of the network or from a computer that is not a member of the internal domain, the user receives a prompt for credentials.
7. The AD DS domain controller authenticates the user, and then sends the success message back to the federation server, along with other information about the user that the federation server can use to generate the user's claims.
8. The federation server creates the claim for the user based on the rules defined during the AD FS server setup. The claims data is placed in a digitally signed security token. Then the data is sent to the client computer, which posts it back to the Azure AD Authentication System.
9. The Azure AD Authentication System validates that the security token came from a trusted federation partner. This trust is configured when you configure the hybrid Exchange Server environment.
10. The Azure AD Authentication System creates and signs a new token that it sends to the client computer, which then sends the token back to the Outlook on the Web server.
11. The Outlook on the Web server receives the request and validates the signed tokens. The server issues the client a session cookie indicating that it has authenticated successfully. The user then is granted access to his or her Exchange Server mailbox.

Configuring SSO for integration with Microsoft online services

SSO, also referred to as identity federation, enables you to simplify your users' sign-in process when they access online services, such as Office 365 or Intune. By using SSO, users can use their internal AD DS credentials to access these online services. When you configure AD FS to provide SSO for Microsoft online services, you create a federated trust between your organization's on-premises directory and the federated domain you specify in your Azure AD tenant.

To configure SSO for integration with online services, you must:


1. Prepare for single sign-on
2. Set up your on-premises AD FS
3. Set up directory synchronization
4. Verify single sign-on




Note: For more information about Azure AD, see Lesson 4: "Integrating AD DS with Azure AD in Module 3: Implementing Directory Services," in this course.


To configure SSO for integration with online services, use the following high-level steps:


1. Prepare for SSO:
 - a. Deploy AD DS in your on-premises environment.
 - b. Install the AD FS role.
 - c. Prepare Active Directory. Depending on your domains, you might need to complete these tasks:
 - i. Verify that the UPNs are set and known by the users.
 - ii. Verify that the UPN domain suffix is under the domain that you choose to set up for SSO.

 **Note:** Remember that UPNs that you use for SSO must contain only letters, numbers, periods, dashes, and underscores.


- iii. Ensure that the domain you choose to federate is registered as a public domain with a domain registrar or within your own public DNS servers.

 **Note:** If your AD DS domain name is not a public Internet domain, you must set a UPN to have a domain suffix that can be registered publicly. In this situation, we recommend that you use something familiar to your users, such as their email domain.

 **Note:** To prepare your Active Directory environment for SSO, you can run the Microsoft Deployment Readiness Tool. This tool inspects your Active Directory environment and provides a report that includes information about whether you are ready to set up SSO. If not, it lists the changes you need to make to prepare for single sign-on.

 **Additional Reading:** To download the tool, go to Microsoft Office 365 Deployment Readiness Tool at: <http://aka.ms/Xzx0o8>

2. Set up your on-premises AD FS:
 - a. Deploy your AD FS server farm.
 - b. Configure extranet access:
 - i. Install the Web Application Proxy role.
 - ii. Configure the Web Application Proxy.
 - c. Establish a trust between AD FS and Azure AD:
 - i. Using Windows PowerShell and the Microsoft Azure Active Directory Module, add the required domains with the **New-MsolFederatedDomain** cmdlet.

 **Additional Reading:** For additional guidance on these steps, refer to: "Checklist: Use AD FS to implement and manage single sign-on" at: <http://aka.ms/Dxonts>

3. Set up directory synchronization:
 - o Download and install Azure AD Connect to enable synchronization of the domain in Microsoft Azure.
4. Verify SSO:
 - a. On a computer that is joined to the domain, sign in to your Microsoft cloud service using the same logon name that you use for your corporate credentials.
 - b. Click inside the password box. If single sign-on is set up, the password box will be shaded, and you will see the following message: "You are now required to sign in at <your company>."
 - c. Click the Sign in at <your company> link. If you are able to sign in, then single sign-on has been set up.

Question: How will your organization implement SSO with online services, such as Microsoft Office 365?

Module Review and Takeaways

Review Questions

Question: Your organization is planning to implement AD FS. In the short term, only internal clients will use AD FS to access internal applications. However, later you will provide access to web-based applications that are secured by AD FS to users at home. How many certificates should you obtain from a third-party CA?

Question: Your organization has implemented a single AD FS server and a single Web Application Proxy successfully. Initially, AD FS was used for only a single application, but now it is used for several business-critical applications. AD FS must be configured to be highly available.

During the installation of AD FS, you selected to use the Windows Internal Database. Can you use this database in a highly available configuration?

Module 5

Implementing network services

Contents:

Module Overview	5-1
Lesson 1: Overview of networking enhancements	5-2
Lesson 2: Implementing IPAM	5-20
Lesson 3: Managing IP address spaces with IPAM	5-30
Lab: Implementing network services	5-37
Module Review and Takeaways	5-45

Module Overview

Windows Server 2016 introduces changes to the features in networking server roles, such as Domain Name System (DNS) server and Dynamic Host Configuration Protocol (DHCP) server. New features include policies to specify how to handle queries in the DNS server role. Windows Server 2016 does not support Network Access Protection (NAP) for the DHCP server role. Additionally, the Network Policy Server (NPS) role in Windows Server 2016 no longer supports NAP. Furthermore, now you can use the IP Address Management (IPAM) feature, which was introduced in Windows Server 2012, to manage DNS zones and support multiple Active Directory Domain Services (AD DS) forests.

In this module, you will learn how to deploy and configure DNS enhancements in Windows Server 2016 and new features in IPAM and DHCP failover, which were introduced in Windows Server 2012. Additionally, you will learn about some enhancements introduced in Windows Server 2016.

Objectives

After completing this module, you will be able to:

- Describe the networking enhancements in Windows Server 2016.
- Implement IPAM.
- Manage IP address spaces with IPAM.

Lesson 1

Overview of networking enhancements

Datacenters are becoming increasingly software defined. In the “Overview of storage in Windows Server 2016” module, you learned about software-defined storage. In this lesson, you will learn about the new and improved software-defined networking technologies that Windows Server 2016 introduces. You can use these technologies, such as Microsoft Hyper-V virtual switch, Hyper-V Network Virtualization, and Windows Server Gateway, to create a fully realized, software-defined datacenter solution for your organization. Module 8, “Implementing software-defined networking,” goes into more detail regarding these technologies.

This lesson will introduce you to the new DNS features in Windows Server 2016, including DNS policies. Although this is not new in Windows Server 2016, you also will learn how to configure and deploy DHCP failover to achieve high availability in your DHCP infrastructure.

Lesson Objectives

After completing this lesson, you will be able to:

- Explain converged networking.
- Describe software-defined networking.
- Describe the components of a DNS solution.
- Describe the new DNS features in Windows Server 2016.
- Describe DNS policies.
- Configure DNS policies.
- Explain DHCP.
- Explain the changes in DHCP features in Windows Server 2016.
- Describe a DHCP failover.
- Configure failover for DHCP.

What is converged networking?

For many years, to ensure performance, organizations typically had several physical servers, each dedicated to a different function. Some organizations might configure two or more servers to ensure performance and availability of specific services. Then, after server virtualization became the standard, organizations typically configured the physical hosts running the virtual machines (VMs) with many network interface cards (NICs) to ensure network performance, and they used one NIC for each type of network traffic. In Hyper-V environments, organizations have used one NIC for management, one for storage, one for backup, one for migration, one for cluster communications, and at least one for the VMs. You typically team some of these NICs to ensure uptime and increase bandwidth. However, this requires physical hosts with at least six, and sometimes up to 10, NICs to ensure network performance. The traffic load on many of these NICs is not particularly high.

Converged networking:

- Uses fewer teamed NICs instead of dedicated NICs for each network traffic type
- Is used in Hyper-V solutions
- Includes the following components:
 - Network teaming
 - Virtual switch
 - Network QoS
 - Network isolation
 - Virtual NICs
- Has improved features in Windows Server 2016, including:
 - Converged RDMA
 - SET

Converged networking is a concept in which one NIC can carry multiple types of network traffic. Windows Server 2016 and Windows Server 2012 R2 support converged networking but with different features. Because there are fewer physical NICs to deliver the same network traffic, the total bandwidth of the NICs can quickly become a bottleneck. Therefore, Quality of Service (QoS) should be a part of the converged networking design. QoS can help ensure that important traffic comes through. When you configure QoS, you use QoS weight to adjust the importance of the traffic. You give each traffic type a weight from 1 to 100, and the higher the value, the more important the traffic.

A Hyper-V converged network consists of the following components:

- Network teaming. Teaming two or more NICs to improve bandwidth and availability.
- Virtual switch. The virtual switch in Hyper-V is a software-based network switch through which VMs communicate.
- Network QoS. QoS prioritizes network traffic.
- Network isolation. Using virtual local area network (VLAN) identifiers, you can direct traffic on the same physical NIC to different segments.
- Virtual NICs. A virtual NIC created on the physical host that you use instead of a physical NIC for special network traffic, such as management or clustering.

Remote Direct Memory Access (RDMA) over Ethernet is a technology that can transfer data more efficiently between two computers. RDMA-capable NICs can transfer data at full speed with low latency, while using very little central processing unit (CPU) power. Windows Server 2016 and Windows Server 2012 use RDMA to provide the Server Message Block (SMB) Direct functionality. Currently three different types of NICs exist that support RDMA:

- InfiniBand
- Internet Wide Area RDMA Protocol (iWARP)
- RDMA over Converged Ethernet (RoCE)

In Windows Server 2012, RDMA NICs do not work in a virtual switch, and you cannot team RDMA NICs. You can use RDMA NICs only for storage with Windows Server 2012.

New features of converged networking in Windows Server 2016

Windows Server 2016 improves on converged networking with the following features:

- Converged RDMA. You can team RDMA NIC, and use them for all types of network traffic.
- Switch Embedded Teaming (SET). You can group between one and eight physical NICs into software-based virtual NICs. SET supports RDMA and therefore offers better performance than Windows Server 2012. Furthermore, you can create converged networking with just two high-performance RDMA-capable NICs.



Additional Reading: For more information, refer to: "Remote Direct Memory Access (RDMA) and Switch Embedded Teaming (SET)" at: <http://aka.ms/Kjbew7>

Overview of software-defined networking

Software-defined networking is a method of configuring and managing physical and virtual network devices—including routers, switches, and gateways centrally in your datacenter. The Hyper-V virtual switch, Hyper-V Network Virtualization, and Windows Server Gateway are integral elements in your software-defined networking infrastructure. You will have better integration if the physical switches, routers, and other hardware devices are compatible with software-defined networking.

- Software-defined networking is to networks what server virtualization is to physical servers
- Windows Server 2016 features in software-defined networking include:
 - Network Controller
 - Hyper-V Network Virtualization
 - Hyper-V virtual switch
 - RRAS Multitenant Gateway
 - NIC Teaming

Software-defined networking is possible because the network planes—management, control, and data—are not bound to the network devices themselves. Therefore, other entities, such as datacenter-management software, including Microsoft System Center Virtual Machine Manager, can configure rules and policies for both physical and virtual network devices.

By using software-defined networking, you can automate management of your datacenter network to dynamically adapt to your applications' requirements. Software-defined networking enables you to:

- Virtualize your network, that so you can abstract your applications and workloads from the physical network. Software-defined networking works with your applications in a nondisruptive manner, similar to server virtualization.
- Centrally define policies to control physical and virtual networks, including the flow of traffic between them.
- Implement network policies consistently, even when you deploy new applications or move applications across virtual or physical networks.

Windows Server 2016 includes several technologies to support software-defined networking, such as:

- **Network Controller.** Windows Server 2016 introduces Network Controller, which provides a centralized point of automation to manage and configure your datacenter's virtual and physical networks.
- **Hyper-V Network Virtualization.** Network Virtualization continues to provide abstraction from the physical network and isolation on a shared physical network.
- **Hyper-V virtual switch.** The virtual switch is a software-based, layer-2, Ethernet-only network switch that is available on Hyper-V servers. The switch includes extensible capabilities to connect VMs to both virtual networks and a physical network.
- **RRAS Multitenant Gateway.** The Routing and Remote Access Service (RRAS) gateway helps you move VM workloads between datacenters and cloud providers, making it easy to extend your datacenter to Microsoft Azure. The gateway is multitenant aware, and hosting providers can use it to provide their customers with a hybrid infrastructure.
- **NIC Teaming.** NIC Teaming allows you to place multiple NICs in a team for bandwidth aggregation and high-availability scenarios.



Additional Reading: For more information, refer to: Module 8, "Implementing Software Defined Networking."


The Packet Direct feature

Packet Direct is a new Windows Server 2016 feature that increases network throughput and delivers low latency on your network connections. To use Packet Direct, you need to use Packet Direct-compatible NICs and drivers. It uses a new I/O model that extends the network driver interface specification (NDIS).

Packet Direct is useful in both physical and virtual scenarios. It uses the current NDIS miniport driver with a Packet Direct extension but implements its own Packet Direct client that does not use the NDIS protocol driver. The Packet Direct client has direct access to the read and transmit queues on the NIC, so network throughput increases as compared to throughput when you use NDIS. In Hyper-V, you can enable Packet Direct on the virtual switch.

Components of a DNS solution

DNS is a name-resolution service that resolves names to IP addresses. The DNS service is a logically partitioned, hierarchically distributed database that enables many different servers to host a worldwide database of DNS names. In Windows Server 2016, DNS is a server role that provides a solution to ensure that client computers can find resources on a domain, local area network (LAN), and the Internet. It also facilitates user and computer authentication in a domain. The components of a DNS solution include internal DNS servers, DNS servers on the Internet, and DNS resolvers or clients.



DNS servers

A DNS server responds to name and IP address-resolution queries from the DNS client service, also known as the *DNS resolver* on other computers. DNS servers can host one or more zones of a particular domain, and zones contain different resource records. DNS servers also can cache lookups to save time for common queries and can store service locator records in the zones that enable clients to find domain controllers in AD DS. Domain controllers add or register their service locator records to the DNS server's zone to which the domain controller belongs. This enables clients to find all domain controllers for the domain to which they belong. In a domain-based corporate network, you need to secure and protect these DNS servers and their resource records. The best practice is to implement Active Directory-integrated zones, thereby combining the DNS server roles and the AD DS role on your domain controllers. This helps enhance security, and it facilitates zone transfers and delegation.

DNS servers on the Internet

DNS servers on the Internet host public zone information, root server information, and other common top-level domains (TLDs), such as .com, .net, and .edu. Other organizations that have their own domain names, such as companies, government agencies, and nonprofit organizations, also have their own DNS servers to which you can send iterative queries through the root and TLD servers. There are millions of these DNS servers, and each one can host resource records of web services that your DNS servers will use to resolve names to IP addresses.

DNS resolvers

A DNS resolver is a service that runs on a client computer, generating and sending iterative or recursive queries to a DNS server. A DNS resolver can be any computer that is performing a DNS lookup that requires interaction with a DNS server. DNS servers also can issue DNS requests to other DNS servers.

When a DNS server responds to a name-resolution request, the DNS resolver caches that information so that it can access it again, if necessary. It stores this information locally, rather than going back to the DNS server each time. However, each record is marked with a Time-to-Live (TTL) time stamp that automatically flushes the record out of the cache when the TTL expires.

New DNS features in Windows Server 2016

DNS is an integral part of all information technology (IT) infrastructures, because organizations use it for host-name resolution. Windows Server 2012 made it easier to configure Domain Name System Security Extensions (DNSSEC). The DNS server functionality continues to provide improved functionality in Windows Server 2016, and there are several improvements to security.

- DNS policies
- Response rate limiting
- DNS-based authentication of named entities
- Unknown record support
- IPv6 root hints
- Enhanced Windows PowerShell support

New DNS features in Windows Server 2016

Windows Server 2016 introduces new and improved features in the DNS server role, including:

- DNS policies. DNS policies specify how a DNS server responds to queries. You can configure policies to allow responses particular to a client IP address, the time of the day, and other parameters. An upcoming topic in this lesson covers DNS policies.
- Response rate limiting. When you enable response rate limiting, you can prevent malicious users from starting denial-of-service attacks using your DNS servers.
- DNS-based authentication of named entities. You can use Transport Layer Security (TLS) authentication records to provide information on the certification authority (CA) that your domain names use for certificates. This prevents man-in-the-middle attacks, in which malicious users corrupt a DNS cache and point it to their website, and might even provide a legitimate certificate issued from a different CA. Request for comments (RFC) 6394 and 6698 describe DNS-based authentication of named entities.
- Unknown record support. DNS in previous versions of Windows Server operating systems required that all records had to be a supported type. The DNS server role in Windows Server 2016, by using the unknown record functionality, supports records that are not supported explicitly.
- IPv6 root hints. Windows Server 2016 now has both the native IP version 4 (IPv4) and IP version 6 (IPv6) root hints to support internet name resolution. The root hints are in the **C:\Windows\system32\dns\cache.dns** file.
- Windows PowerShell support. New Windows PowerShell cmdlets are available in Windows Server 2016, and they facilitate administration of new DNS features.

New DNS features in Windows Server 2012 R2

Windows Server 2012 R2 introduces improved features in the DNS server role, including:

- Enhanced zone-level statistics. New statistics information is available for resource records, zone transfers, and dynamic updates. You can retrieve statistics information by using the **Get-DnsServerStatistics** cmdlet in Windows PowerShell.

- Enhanced DNSSEC support. Several DNSSEC technologies feature improvements, such as key management and enhanced support for signed zones saved in files.
- Enhanced Windows PowerShell support. Windows PowerShell includes new cmdlets for configuring DNSSEC.



Additional Reading: For more information, refer to: "What's New in DNS Server in Windows Server 2012 R2" at: <http://aka.ms/q8slci>

DNSSEC

Intercepting and tampering with an organization's DNS query response is a common attack method that malicious users utilize. If an attacker can alter the response from a DNS server, or send a spoofed response to point client computers to their own servers, then they can gain access to sensitive information. This is a *man-in-the-middle attack*. Any service that relies on DNS for an initial connection, such as e-commerce web servers and email servers, is vulnerable to these attacks. DNSSEC helps protect clients that are making DNS queries from accepting false DNS responses. DNSSEC continues to be available, and is easier to implement, in Windows Server 2016.

New resource records

You achieve validation of DNS responses by associating a private/public key pair, which you as the administrator generate, with a DNS zone, and then defining additional DNS resource records to sign and publish keys. Resource records distribute the public key, while the private key remains on the server. DNS clients also need to support public key infrastructure (PKI) to query for validation of these signed DNS records. When the client requests validation, DNSSEC adds data to the response that enables the client to authenticate the response. Windows Server 2012 defined the new resource records in the following table.

Resource record	Purpose
DNSKEY	This record is the public key for the zone. The DNSKEY decrypts responses that an authoritative DNS server generates. These keys require periodic replacement, or <i>key rollover</i> . Windows Server 2012 and later operating systems support automated key rollovers. DNSKEYs are divided into zone-signing keys (ZSKs) and key-signing keys (KSKs). You can use ZSKs to sign all record sets in a zone, but they typically have short validity periods. However, KSKs sign ZSKs, and typically have longer validity periods than ZSKs.
Delegation signer (DS)	A DS is a delegation record that contains the public key's hash, which is a child zone's DNSKEY. The parent zone's private key signs the record. If a child zone of a signed parent also is signed, the DS records from the child must be manually added to the parent. Its DNSSEC records must be signed again, so that you create a chain of trust.
Resource record digital signature (RRSIG)	This record is the digital signature for a DNS record, and it asserts that a DNS record is authentic.
Next Secure (NSEC)	This record type is returned when DNS is queried for hosts that do not exist in the zone. The non-existence of the host is proved because the NSEC response includes the last and first valid host names within the range being queried.
NSEC3	This hashed version of the NSEC record helps prevent alphabet attacks by enumerating the zone and is a successor to NSEC. NSEC3 returns a hashed value of the last and first valid hosts, so it is a more secure method of handling queries if no host records exist.

Trust anchors

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



Note: All domain controllers that are hosting DNS must run Windows Server 2012 or later before you can distribute trust anchors to them.

NRPT

The name resolution policy table (NRPT) contains rules that control the DNS client behavior for sending DNS queries and processing the responses from those queries. For example, a DNSSEC rule prompts the client computer to check if the response for a particular DNS domain suffix is valid. We recommend Group Policy for configuring the NRPT. If there is no NRPT present, the client computer does not validate responses.

Factors to consider before implementing DNSSEC

Consider the following factors before you implement DNSSEC:

- The zone replication scope or type cannot be changed while a zone is signed.
- The encryption algorithm affects the zone signing time, which in turn affects the overall server reboot time.
- When DNSSEC is deployed, DNS response messages are larger than when DNSSEC is not deployed.
- When DNSSEC is deployed, zone files are larger than when DNSSEC is not deployed.
- Queries for DNSKEY records cause DNS traffic to increase.
- The client computer must spend more time authenticating responses.
- There is an additional level of administration necessary to maintain DNSSEC.



Note: Only Windows Server 2008 R2, Windows 7, and newer operating systems support DNSSEC. During the name-resolution process, the DNS client sends information to the DNS server, regardless of whether the DNS client supports DNSSEC.

DNS policies

Windows Server 2016 introduces DNS policies, which you can use to control how a DNS server handles queries based on several parameters. For example, you can create a DNS policy to respond to a query asking for a web server's IP address to respond with a different IP address based on the time of day. You can use DNS policies in many scenarios, including when you want to:

- Provide application load balancing. You can redirect DNS clients to each endpoint for a given application, based on configurable percentages of traffic.
- Manage traffic. You can redirect DNS clients to the closest datacenter.
- Use split-brain DNS. Typically, split-brain DNS requires two zones that are on two different servers. However, when you use DNS policies, you can split DNS records into different zone scopes, and DNS clients receive a response based on whether they are internal or external clients.
- Provide filtering. You can block or allow DNS queries from a list of malicious or approved IP addresses or fully qualified domain names (FQDNs).
- Utilize redirection based on time. You can redirect DNS clients to different datacenters based on the time of the day.

- Use DNS policies to:
 - Redirect DNS clients to endpoints based on configurable percentages of traffic
 - Redirect DNS clients to the closest datacenter
 - Configure split-brain DNS
 - Block or allow DNS queries from a list of malicious or approved IP addresses or FQDNs
 - Redirect DNS clients to different datacenters based on the time of the day
- You configure DNS policies by using Windows PowerShell

Currently, you can only configure DNS policies using Windows PowerShell. DNS policies require Windows Server 2016, so if you decide to configure DNS policies for a zone, all DNS servers responsible for the zone need to run Windows Server 2016. There is no requirement for the DNS clients, because the DNS servers internally handle all aspects of DNS policies. You can create different types of DNS policies as listed in the following table.

Type	Level	Description
Query resolution policy	Server or zone	You can use this type of policy to specify how DNS servers handle incoming DNS resolution queries.
Recursion policy	Server	You can configure DNS policies to allow the DNS server to perform recursion for specific DNS clients.
Zone transfer policy	Server or zone	You can allow or deny zone transfers from your DNS server based on different criteria, such as client subnet, DNS server interface IP address, time of day, and use of IPv4 or IPv6.

The following scenarios explain some uses for DNS policies.

Block DNS clients from a subnet from making name resolution

In this scenario, you want to block computers from a certain subnet from making name resolution queries against the DNS servers. To block a subnet from using the DNS server for name resolution, type the following Windows PowerShell cmdlets at a Windows PowerShell prompt, and press Enter after each:

```
Add-DnsServerClientSubnet -Name "MaliciousSubnet01" -IPv4Subnet 192.168.33.0/24 -PassThru
Add-DnsServerQueryResolutionPolicy -Name "BlackholePolicyMalicious01" -Action IGNORE -
ClientSubnet "EQ,MaliciousSubnet01" -PassThru
```

Apply load balancing

In another scenario, you might place one web farm in each of your two datacenters. However, because your web farm in datacenter 1 has twice the capacity of the web farm in datacenter 2, you want the big farm to handle two-thirds of the load and the small farm in datacenter 2 to handle one-third of the load.

To achieve this, you have to configure zone scopes, which are unique instances of a zone. One DNS zone can have several zone scopes. Each zone scope contains its own set of DNS records. You can create the same records in multiple scopes, each with different IP addresses.

You can use the **Add-DnsServerZoneScope** Windows PowerShell cmdlet to configure zone scopes. To do so, type the following Windows PowerShell cmdlets at a Windows PowerShell prompt, and press Enter after each:

```
Add-DnsServerZoneScope -ZoneName "contoso.com" -Name "Datacenter1ZoneScope"
Add-DnsServerZoneScope -ZoneName "contoso.com" -Name "Datacenter2ZoneScope"
```

The next step is to add the resource records to the zone scopes. You use the **Add-DnsServerResourceRecord** cmdlet with the **ZoneScope** parameter. Type the following Windows PowerShell cmdlets at a Windows PowerShell prompt, and press Enter after each:

```
Add-DnsServerResourceRecord -ZoneName "contoso.com" -A -Name "intranet" -IPv4Address
"192.168.5.55" -ZoneScope "Datacenter1ZoneScope"
Add-DnsServerResourceRecord -ZoneName "contoso.com" -A -Name "intranet" -IPv4Address
"10.10.5.55" -ZoneScope "Datacenter2ZoneScope"
```

You also have to specify the subnets from where your DNS clients originate. To do this, use the **Add-DnsServerClientSubnet** cmdlet. Type the following Windows PowerShell cmdlets at a Windows PowerShell prompt, and press Enter after each:

```
Add-DnsServerClientSubnet -Name "Location1Subnet" -IPv4Subnet "192.168.5.0/24" -PassThru
Add-DnsServerClientSubnet -Name "Location2Subnet" -IPv4Subnet "10.10.5.0/24" -PassThru
```

Finally, you have to create a policy that states that the web farm in datacenter 1 should receive twice the amount of traffic as the web farm in datacenter 2. You use the **Add-DnsServerQueryResolutionPolicy** cmdlet to create the policy, by typing the following Windows PowerShell cmdlet at a Windows PowerShell prompt, and then pressing Enter:

```
Add-DnsServerQueryResolutionPolicy -Name "IntranetPolicy" -Action ALLOW -ZoneScope
"Datacenter1ZoneScope,2;Datacenter2ZoneScope,1" -ZoneName "contoso.com"
```



Additional Reading: For more information, refer to "DNS Policies Overview" at: <http://aka.ms/lm8s95>

Demonstration: Configuring DNS policies

In this demonstration, you will see how to configure and verify DNS policies.

Demonstration Steps

1. On **TREY-DC1**, start the **DNS** console.
2. Create a new Conditional Forwarder for the **Adatum.com** domain, and then configure the master server to **172.16.0.10**.
3. Open a **Windows PowerShell** window.
4. In the **Windows PowerShell** window, type the following three commands, pressing Enter after each command:

```
Clear-DnsClientCache  
Clear-DnsServerCache  
Resolve-DnsName LON-DC1.Adatum.com
```

5. Verify that the name resolves to an IP address.
6. Switch to **LON-DC1**.
7. On **LON-DC1**, start a **Windows PowerShell** window.
8. In the **Windows PowerShell** window, type the following two commands, pressing Enter after each command:

```
Add-DnsServerClientSubnet -Name "TreyResearchSubnet" -IPv4Subnet 172.16.10.0/24 -  
PassThru  
Add-DnsServerQueryResolutionPolicy -Name "BlackholePolicyTreyResearch" -Action IGNORE  
-ClientSubnet "EQ,TreyResearchSubnet" -PassThru
```

9. Switch to **TREY-DC1**.
10. On **TREY-DC1**, in the **Windows PowerShell** window, type the following three commands, pressing Enter after each command:

```
Clear-DnsClientCache  
Clear-DnsServerCache  
Resolve-DnsName LON-DC1.Adatum.com
```

11. Verify that the last command returns an error because the DNS policy no longer allows **TREY-DC1** to perform name resolution on **LON-DC1**.

Overview of DHCP

DHCP is a server role that you can install on Windows Server 2016. By using the DHCP server role, you can ensure that all clients have appropriate IP addresses and network-configuration information. This can help eliminate errors that you might make during configuration. A DHCP client is any device that takes a DHCP address and that can request and retrieve network settings from a DHCP server service. DHCP clients might be computers, mobile devices, printers, or switches. DHCP also can provide IP address information to network boot clients.

- DHCP components consist of:
 - The DHCP server service
 - DHCP scopes
 - DHCP options
 - The DHCP database
 - The DHCP console
- When you use DHCP:
 - Clients request IP configuration through a broadcast
 - IP addresses are leased to clients for a configurable period and are renewed regularly
 - DHCP servers must be authorized in AD DS

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

DHCP consists of the components listed in the following table.

Component	Description
DHCP server service	After you install the DHCP server role, you can implement the DHCP server as a service. This service can distribute IP addresses and other network configuration information to clients that request it.
DHCP scopes	<p>The DHCP administrator configures the range of IP addresses and related information that is allotted to the server for distribution to requesting clients. You can associate each scope with a single IP subnet only, and a scope must consist of:</p> <ul style="list-style-type: none"> • A name and description • A range of addresses that the server distributes • A subnet mask <p>Additionally, a scope also can define:</p> <ul style="list-style-type: none"> • IP addresses that should be excluded from distribution • The duration of the IP address lease • DHCP options <p>You can configure a single DHCP server with multiple scopes, but you must connect the server directly to each subnet that it serves or have a supporting and configured DHCP relay agent in place. Scopes also are the primary method for a server to manage and distribute any related configuration parameters, such as DHCP options, to network clients.</p>
DHCP options	<p>When you assign an IP address to a client, you also can simultaneously assign many other network-configuration parameters. The most common DHCP options include:</p> <ul style="list-style-type: none"> • Default gateway IP address • DNS server IP address

Component	Description
	<ul style="list-style-type: none"> DNS domain suffix Windows Internet Name Service (WINS) server IP address <p>You can apply the options at different levels, as follows:</p> <ul style="list-style-type: none"> Globally to all scopes Specifically, to particular scopes To specific clients, based on a class ID value To clients that have specific IP address reservations configured
DHCP database	The DHCP database contains configuration data about the DHCP server, and it stores information about the IP addresses that have been distributed. By default, the DHCP database files are stored in the %systemroot%\System32\Dhcp folder. The DHCP database is a Microsoft Jet database.
DHCP console	The DHCP console is the main administrative tool for managing all aspects of the DHCP server. This management console is installed automatically on any server on which you install the DHCP server role. However, you also can install it on a remote server or Windows 10 client by using the Remote Server Administration Tools (RSATs) and by connecting to the DHCP server for remote management.

How clients acquire IP addresses

When you configure a Windows client operating system to use the DHCP service, when it starts, the client will use an Address Resolution Protocol (ARP) broadcast in its subnet to request IP configuration from any DHCP server that might receive the request. However, because DHCP uses broadcasts to initiate communications, DHCP servers can communicate only within their IP subnets. This means that there must be a DHCP server on each IP subnet or you must configure a router to forward bootstrap protocol (BOOTP) traffic to the DHCP relay agent that is configured on the remote subnet. Either the DHCP relay service or BOOTP forwarding can relay DHCP broadcast packets as directed messages across a router and into other IP subnets. The relay agent acquires an IP address configuration for the remote subnet's requesting client and then forwards that configuration to the client.

DHCP leases

DHCP allocates IP addresses on a dynamic basis, or as a *lease*. You can configure a lease's duration. The default lease time for wired clients is eight days, but mobile or handheld devices, such as tablets, usually have a shorter lease duration. Typically, if you have a higher turnover of devices or users, the lease time should be shorter. Conversely, if there is more device and user permanency, it can be longer. You can configure the lease settings in the **DHCP** console, under the server name and the IPv4 or IPv6 node, by clicking **Scope**, and then clicking **Properties**.

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

DHCP server authorization

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

Windows PowerShell

You can use Windows PowerShell cmdlets to provide command-line support for managing DHCP. To use the DHCP cmdlets, you must load the **DhcpServer** module. In addition to providing command-line support, you can use Windows PowerShell cmdlets if you want to script your DHCP management. The following table includes a subset of the Windows Server 2016 Windows PowerShell cmdlets for managing DHCP.

Cmdlet	Additional information
Add-DhcpServerInDC	Use to add the specified computer that runs the DHCP server service as an authorized DHCP server in AD DS.
Add-DhcpServerv4Class	Use to add an IPv4 vendor or user class to the DHCP server service.
Add-DhcpServerv4ExclusionRange	Use to add an IP address-exclusion range to an IPv4 scope.
Add-DhcpServerv4Failover	Use to add a new IPv4 failover relationship on the DHCP server service.
Add-DhcpServerv4FailoverScope	Use to add one or more scopes to an existing failover relationship.
Add-DhcpServerv4Filter	Use to add a media access control (MAC) address filter to the DHCP server service. You can use the filter on an Allow or Deny list.
Add-DhcpServerv4Lease	Use to add a new IPv4 address lease in the DHCP server service for testing purposes.
Add-DhcpServerv4OptionDefinition	Use to add a new DHCP version 4 (DHCPv4) option definition to the DHCP server service.
Add-DhcpServerv4Policy	Use to add a new IPv4 policy to a DHCP server or a DHCP scope.
Add-DhcpServerv4PolicyIPRange	Use to add an IP range to an existing scope policy.
Add-DhcpServerv4Reservation	Use to reserve the specified IPv4 address in the specified DHCP scope for a specified client.
Add-DhcpServerv4Scope	Use to add an IPv4 scope on the DHCP server service.

You can see all cmdlets by running the following command on a Windows Server 2016 server that has the DHCP server role installed:

```
Get-Command -Module DhcpServer
```

Changes in DHCP features in Windows Server 2016

The DHCP server role in Windows Server 2016 has no new features or improvements; however, there is one big change that might affect organizations that have implemented NAP. In Windows Server 2016, the DHCP server role no longer supports NAP. DHCP clients that send a statement of health together with their DHCP request will receive a normal DHCP address lease from a DHCP server that is running Windows Server 2016.



Additional Reading: For more information, refer to: "What's New in DHCP in Windows Server 2016" at: <http://aka.ms/lyimjo>

- DHCP does not support NAP
- Features new to Windows Server 2012 and Windows Server 2012 R2 include:
 - DHCP failover
 - DHCP policies
 - DHCP name protection
 - Disable PTR record registration
 - Improved Windows PowerShell support



Note: The rest of this topic discusses important DHCP features that are not new in Windows Server 2016.

DHCP failover

DHCP failover is a feature in Windows Server 2016 that allows you to configure high availability for the DHCP service. In Windows Server 2008 R2 and earlier versions, DHCP failover was not possible because DHCP servers were independent and unaware of one another. Configuring two separate DHCP servers to distribute IP addresses within the same IP address pool could lead to duplicate address assignment if the administrator did not configure the split-scope overlapping ranges correctly.

The DHCP failover feature enables an alternative DHCP server to distribute IP addresses and associated option configuration to the same subnet or scope. The two DHCP servers replicate lease information between them through an established partner relationship, which provides one server with the configuration to know from where to take over if the other DHCP server fails. If one of the DHCP servers fails, the other DHCP server provides service for client computers. In Windows Server 2016, you can configure one alternative DHCP server for failover, but only for IPv4 scopes.

DHCP policy-based assignment

Policy-based assignment enables the DHCP server to evaluate DHCP requests based on policies that an administrator defines. You apply policies at the server or scope level, and a policy contains a set of conditions based on fields in the client request, such as the:

- Vendor class
- User class
- MAC address
- Client Identifier
- Relay agent information

The DHCP server can assign different DHCP options and addresses based on the policy criteria that the client request matches. For example, you could add a vendor class that matches a particular type of printer and have DHCP addresses from a specific range assigned within the scope when a printer that matches those criteria requests a DHCP address.

DHCP name protection

Windows Server 2016 supports DHCP name protection. When the DHCP server registers names in DNS for computers that run Windows operating systems, DHCP must protect the DNS records from being overwritten by non-Microsoft operating systems that have the same name. For example, a UNIX-based system with the name **Client1** potentially could overwrite the DNS address that the DHCP server assigned and registered on behalf of a Windows-based operating system that also has the name **Client1**. DHCP name protection addresses this issue.

DHCP Name Protection uses a DHCP Information (DHCID) resource record to track which computer originally requested the name. The DHCP server provides this record, and it is stored in DNS. When the DHCP server receives a request to update a host record that is currently associated with a different computer, the DHCP server can verify the DHCID in DNS to check whether the requester is the original owner of the name. If it is not the same computer, DHCP will not update the record in DNS.

To resolve this issue, either the current host name owner must release the IP address, or the requester must use a different host name. You can implement name protection for both IPv4 and IPv6 by configuring it on the properties page on the **DNS** tab at the IP address level or the scope level.

New and improved DHCP features in Windows Server 2012 R2

Windows Server 2012 R2 introduced new and improved features in the DHCP server role, including:

- DNS registration enhancements. Administrators can configure different DHCP policies based on the DHCP client's FQDNs. The DHCP server can register DHCP clients with a different DNS suffix and override the DNS suffix previously configured on the DHCP client.
- DNS PTR record registration options. The DHCP server can register only host (A) resource records of DHCP clients with the DNS server but not the pointer (PTR) record. This scenario addresses organizations that do not have a reverse lookup zone created, so that only the A record is registered in DNS. Administrators can disable PTR record registration for all DHCP clients or for specific DHCP clients. You can select clients for which to disable PTR record registration based on different criteria, such as subnet, the DHCP clients' locations, or specific DHCP client attributes.
- New and improved Windows PowerShell cmdlets for DHCP. Administrators can use new Windows PowerShell cmdlets to perform different management tasks, such as creating DHCP security groups or creating super scopes.



Reference Links: For more information, refer to: "Windows PowerShell for DHCP server" at: <http://aka.ms/Pj3sjc>

What is DHCP failover?

DHCP manages IP address distribution in TCP/IP networks of all sizes. When this service fails, clients lose connectivity to the network and all of its resources. DHCP failover is a feature of Windows Server 2016 that addresses this issue.

DHCP failover

DHCP clients renew their leases on their IP addresses at regular, configurable intervals. If the DHCP service fails and the leases time out, the clients no longer have IP addresses. For Windows Server 2008 R2 and earlier versions, DHCP failover

DHCP failover:

- Enables two DHCP servers to provide IP addresses and optional configurations to the same subnets or scopes
- Requires failover relationships to have unique names
- Supports the hot standby and load sharing modes

When you use DHCP failover:

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

was not possible because DHCP servers were independent and unaware of each other. Therefore, if you configured two separate DHCP servers to distribute the same pool of addresses, you could have duplicate addresses. Additionally, to provide redundant DHCP services, you had to configure clustering and perform a significant amount of manual configuration and monitoring.

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



Note: You can configure only two DHCP servers for failover in one failover relationship and only for IPv4 scopes and subnets.

Configuring DHCP failover

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



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

The following table details the two modes for which you can configure failover.

Mode	Characteristics
Hot standby	<p>One server is the primary server, and the other is the secondary server. The primary server actively assigns IP configurations for the scope or subnet. The secondary DHCP server assumes this role only if the primary server becomes unavailable. A DHCP server can simultaneously act as the primary for one scope or subnet, and the secondary for another.</p> <p>Administrators must configure a percentage of the scope addresses assigned to the standby server. If the primary server is down, the standby server receives these addresses during the maximum client lead-time interval. The default value is 5 percent of the scope, meaning 5 percent of the available addresses are reserved for the secondary server. The secondary server takes control of the entire IP range after the maximum client lead-time interval passes. When the primary server is down, addresses from the secondary server use a lease time that is equal to the maximum client lead-time, which is one hour by default.</p> <p>Hot standby mode is best for deployments in which a disaster-recovery site is at a different location. This ensures that the DHCP server will not service clients unless there is a main server outage.</p>
Load sharing	<p>This is the default mode. In this mode, both servers supply IP configuration to clients simultaneously. The configuration decides which server will respond to IP configuration requests, depending on how the administrator configures the load distribution ratio. The default ratio is 50:50.</p>

Maximum client lead time

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

Auto state switchover interval

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

Message authentication

Windows Server 2016 can authenticate the failover message traffic between replication partners. The administrator can establish a shared secret, much like a password, in the **Configuration Failover Wizard** for DHCP failover. This validates that the failover message comes from the failover partner.

Firewall considerations

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

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

Demonstration: Configuring DHCP failover

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

Demonstration Steps

Configure a DHCP failover relationship

1. On **LON-SVR1**, open the **DHCP** console. Verify that the server is authorized but that no scopes are configured.
2. In the **DHCP** console, add **LON-DC1** as a server.
3. In the **lon-dc1** node, in **IPv4**, launch **Configure Failover Wizard**.
4. Configure failover replication with the following settings:
 - o Partner server: **172.16.0.11**
 - o Relationship Name: **Adatum**
 - o Maximum Client Lead Time: **15 minutes**
 - o Mode: **Load balance**
 - o Load Balance Percentage: **50%**
 - o State Switchover Interval: **60 minutes**
 - o Message authentication shared secret: **Pa\$\$w0rd**
5. Complete the **Configure Failover Wizard**.
6. In the **lon-svr1.adatum.com** node, refresh the **IPv4** node. Notice that the **Adatum** scope is configured and active.

Check Your Knowledge

Question	
Which of the following options are new or different in Windows Server 2016? (Select all that apply.)	
Select the correct answer.	
<input type="checkbox"/>	DNS policies
<input type="checkbox"/>	DHCP failover
<input type="checkbox"/>	IPv6 root hints in DNS
<input type="checkbox"/>	Hyper-V virtualized networking
<input type="checkbox"/>	No DHCP server support for NAP

Verify the correctness of the statement by placing a mark in the column to the right.

Statement	Answer
You need more physical NICs when you configure converged networking on a Hyper-V host.	

Lesson 2

Implementing IPAM

With the development of IPv6 and the increase in devices that require IP addresses, networks have become more complex and difficult to manage. Maintaining an updated list of static IP addresses issued to devices has often been a manual task, which can lead to errors. Now, organizations can manage IP addresses by using the IPAM tool provided by Windows Server 2016.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe IPAM.
- Describe the IPAM architecture.
- Describe scenarios in which to use IPAM.
- Describe the requirements for IPAM implementations.
- Implement IPAM.
- Explain how to manage IPAM.

What is IPAM?

IP address management is a difficult task in large networks, because tracking IP address usage is largely a manual operation. Windows Server 2016 features the IPAM tool, which is a framework for discovering, auditing, monitoring utilization of, and managing the IP address space in a network. IPAM enables the administration and monitoring of DHCP and DNS, and it provides a comprehensive view of where you use IP addresses. IPAM collects information from domain controllers and NPSs and then stores that information in a database. In Windows Server 2016, you can store that database in a Windows Internal Database (WID) or external database.

- IPAM includes administration and monitoring of:
 - IP addresses
 - DHCP services
 - DNS services
- IPAM benefits for network administrators include:
 - Planning and allocation functionality for IPv4 and IPv6 address spaces
 - Utilization statistics and trend monitoring for IP address spaces
 - Static IP inventory management, lifetime management, and DHCP and DNS record creation and deletion
 - Service and zone monitoring of DNS services

The following table provides details about how you can use IPAM when administering IP addresses.

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

Characteristics of IPAM

The characteristics of IPAM include that:

- A single IPAM server can support up to 150 DHCP servers and 500 DNS servers.
- A single IPAM server can support up to 6,000 DHCP scopes and 150 DNS zones.
- IPAM stores three years of forensics data (IP address leases, host MAC addresses, and user sign in and sign out information) for 100,000 users in a WID when you use Windows Server 2012. In Windows Server 2016 and Windows Server 2012 R2, you can choose between a WID and Microsoft SQL Server. There is no database purge policy provided, and the administrator must purge the data manually, as necessary.
- IPAM on Windows Server 2012 supports WID only. You can use an external database only when you implement IPAM on Windows Server 2016 or Windows Server 2012 R2.
- Only IPv4 has IP address utilization trends.
- Only IPv4 has IP address reclamation support.
- IPAM does not check for IP address consistency with routers and switches.

Benefits of IPAM

The benefits of IPAM include that it provides:

- Planning and allocation for IPv4 and IPv6 address spaces.
- Utilization statistics and trend monitoring for IP address spaces.
- Static IP inventory management, lifetime management, and DHCP and DNS record creation and deletion.
- Service and zone monitoring of DNS services.
- Tracking for IP address leases and logon events.
- Role-based access control (RBAC).
- Remote administration support through Remote Server Administration Tools.
- Reporting functionality in the IPAM management console.



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

Windows Server 2016 enhancements to IPAM

Windows Server 2016 improves and adds functionality for IPAM. The IPAM framework now includes the following functionality:

- Enhanced IP address management. IPAM now supports /31, /32, and /128 subnets for use in point-to-point and certain switch configurations. New Windows PowerShell cmdlets are available to return available subnets and IP ranges.
- Enhanced DNS service management. IPAM improves DNS service management, including that you now can:
 - Collect DNS zones and resource records (other than those pertaining to DNSSEC) from DNS servers that run Windows Server 2008 or newer.
 - Configure properties and operations on all types of resource records.

- Configure properties and operations on all types of DNS zones.
- Perform tasks on secondary and stub zones. You now can perform tasks such as **Transfer from Master** or **Transfer new copy of zone from Master**.
- Utilize RBAC for supported DNS configurations.
- Collect conditional forwarders, and then configure them.
- Integrated DNS, DHCP, and IP address management. During DNS resource record collection, IPAM will create IP addresses from PTR records collected, if you map the reverse lookup zone to an IP address range.
- Multiple AD DS forest support. IPAM now supports discovering and managing DHCP and DNS servers in a different AD DS forest than the one to which the IPAM server belongs, as long as there is a two-way forest trust between the two AD DS forests.
- Windows PowerShell support for RBAC. You can use Windows PowerShell cmdlets to configure RBAC.

Windows Server 2012 R2 enhancements to IPAM

Windows Server 2012 R2 improved and added functionality for IPAM. The IPAM framework was expanded with the following:

- RBAC. You can use RBAC for IPAM to customize roles, access scopes, and access policies for IPAM administrators.
- Virtual address space management. You can use IPAM to manage IP addresses in a Microsoft-based network. You can manage both physical and virtual addresses. Integration between IPAM and Virtual Machine Manager allows end-to-end address space management. You can view virtual address space in the **IPAM** console's **VIRTUALIZED ADDRESS SPACE** node: this was new in Windows Server 2012 R2.
- Enhanced DHCP server management. DHCP management was improved in Windows Server 2012 R2 and includes new DHCP scope and DHCP server operations. Additionally, views were added for DHCP failover, DHCP policies, DHCP superscopes, DHCP filters, and DHCP reservations.
- External database support. You can configure IPAM to use a WID. Support for using SQL Server was added in Windows Server 2012 R2.
- Upgrade and migration support. You can upgrade the IPAM database from Windows Server 2012 to Windows Server 2012 R2.
- Enhanced Windows PowerShell support. IPAM includes more than 50 Windows PowerShell commands.

IPAM architecture

IPAM consists of four main modules, which the following table details.

IPAM modules	IPAM topologies	IPAM components
<ul style="list-style-type: none"> • IPAM discovery • IP address space management • Multi-server management and monitoring • Operational auditing and IP address tracking 	<ul style="list-style-type: none"> • Centralized • Distributed • Hybrid 	<ul style="list-style-type: none"> • IPAM server • IPAM client

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

In Windows Server 2016, the IPAM server can manage multiple AD DS forests if there is a two-way trust between the forests. IPAM is deployed in one of three topologies:

- **Centralized.** This model consists of a single IPAM server that manages all infrastructure in the environment. This model is used when centralized administration and reporting is required or for smaller organizations.
- **Distributed.** This model involves IPAM servers deployed for each site in the organization. This model is useful when an organization assigns infrastructure management to individual business groups, for their specific area, or when the number of IPAM clients is too large for one IPAM server.
- **Hybrid.** The hybrid model consists of a centralized IPAM server for management and reporting and individual IPAM servers at each site for increased redundancy, load handling, or administrative role separation.



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

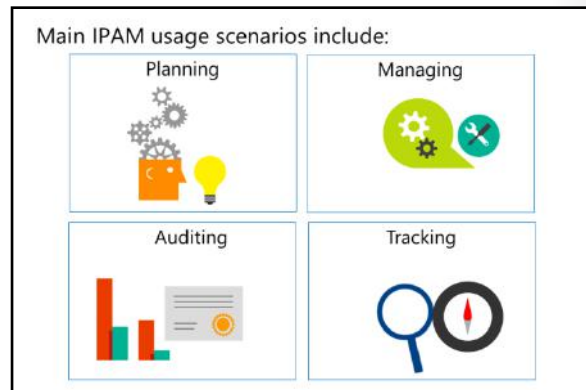
IPAM has two main components:

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

Scenarios in which to use IPAM

Organizations will choose to deploy and use IPAM based on their business requirements. There are several usage scenarios for organizations that deploy IPAM, including for:

- Planning. Organizations might need to reevaluate and plan reconfiguration of their IP address space. Administrators can use IPAM to plan IP address space and optimize capacity for IP addresses, DHCP, and DNS servers.
- Managing. Administrators can use IPAM's single and centralized console to manage and administer IP addresses, DNS resource records and zones, DHCP, and DNS servers. Centralized management using IPAM is more efficient than administering networks by using multiple consoles on different remote servers.
- Tracking. Administrators can use IPAM for tracking IP address utilization. In previous versions of the Windows Server operating system before IPAM, administrators had to use manual documentation, Microsoft Excel worksheets, or other non-Microsoft software to document IP address usage and assignment.
- Auditing. You can use IPAM for addressing compliance and change-management procedures that an organization's information security officers or compliance managers require. IPAM provides reporting for security audits or forensic investigation, based on regulations such as SOX.



Requirements for implementing IPAM

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

- The IPAM server must be a domain member, but it cannot be a domain controller.
- The IPAM server should be a single-purpose server. Do not install other network roles, such as DHCP or DNS, on the same server. IPAM then cannot detect services on other servers.
- You must have IPv6 enabled on the IPAM server, if you want to manage an IPv6 address

Prerequisites to ensure a successful IPAM implementation:

- IPAM server must belong to the domain
- IPAM server cannot be a domain controller
- Enable IPv6
- Sign in with a domain account
- Ensure correct IPAM local security group
- Enable logging of account logon events

space.

- You must sign in to the IPAM server with a domain account and not with a local account.
- You must be a member of the correct IPAM local security group on the IPAM server.
- You must enable logging of account logon events on domain controllers and NPS servers for IPAM's IP address tracking and auditing feature.

If you manage Windows Server 2008 and Windows Server 2008 R2 with IPAM, the Windows Server 2008 or Windows Server 2008 R2 servers require that you:

- Install Service Pack 2 (SP2) on Windows Server 2008.
- Install the full version of the Microsoft .NET Framework 4.0.
- Install Windows Management Framework 3.0 (KB2506146).
- Install Windows Management Framework Core (KB968930) if you are using Windows Server 2008 SP2.
- Enable Windows Remote Management.
- Verify that service principal names (SPNs) are written to AD DS.

After you have decided which IPAM topology to use, you can deploy IPAM servers by performing the following steps:

1. Install the IPAM Server feature by using Server Manager or with the following Windows PowerShell cmdlet:

```
Install-WindowsFeature IPAM -IncludeManagementTools
```

2. Provision IPAM servers. After the feature installation, you must provision each IPAM server to create the permissions, file shares, and settings on managed servers. You can do this either manually, or by deploying Group Policy Objects (GPOs). There are several advantages to using GPOs instead of manual provisioning, including that:
 - GPO-applied settings are less prone to human configuration errors.
 - GPO settings are applied automatically to servers when you assign them a status of **managed**.
 - You can remove GPO settings easily by disabling or deleting the GPO link.
3. Configure and run server discovery. You must configure the scope of discovery for servers that you are going to manage. Discovery scope is determined by selecting the domain or domains on which the IPAM server runs the discovery. You also can add a server manually in the **IPAM management** console by specifying the FQDN of the server that you want to manage.
4. Choose and manage discovered servers. After discovery is complete, and you have added any servers manually that were not discovered, you must choose the servers that you want to manage. To do this, you must edit the server properties in the **IPAM** console and change **Manageability Status** to **Managed**. Refresh Group Policy on the managed servers, and after the management permission for a server is set successfully, you will see a status indicator in the IPAM server inventory displaying **IPAM Access Unblocked**.

Demonstration: Implementing IPAM

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

Demonstration Steps

Install IPAM

1. On **LON-SVR2**, start **Server Manager**.
2. In the **Server Manager** console, add the **IP Address Management (IPAM) Server** feature and other required features.

Configure IPAM

1. In the **Server Manager** navigation pane, click **IPAM**.
2. In the **IPAM Overview** content pane, verify that IPAM is connected to **LON-SVR2.ADATUM.COM** as **Adatum\Administrator**.
3. Click **Provision the IPAM server**.
4. In the **Provision IPAM Wizard**, select the default values, and then configure the following setting in the wizard:
 - o GPO name prefix: **IPAM**
5. In the **IPAM Overview** content pane, click **Configure server discovery**.
6. Click **Get forests**. In the **Configure Server Discovery** dialog box, click **OK**, and then click **Cancel**.
7. In the **IPAM Overview** content pane, click **Configure server discovery**.
8. Add the **Adatum.com** domain.
9. In the **IPAM Overview** content pane, click **Start server discovery**. Discovery might take 5 to 10 minutes to run. The yellow status bar indicates when discovery is complete.
10. In the **IPAM Overview** content pane, click **Select or add servers to manage and verify IPAM access**. Notice that the **IPAM Access Status** is blocked for **LON-DC1**. Scroll down to the **Details** view, and then notice the status report.
11. Start a **Windows PowerShell** window.
12. At the Windows PowerShell prompt, type the following command, and then press Enter:

```
Invoke-IPamGpoProvisioning -Domain Adatum.com -GpoPrefixName IPAM -IpamServerFqdn LON-SVR2.adatum.com -DelegatedGpoUser Administrator
```
13. When you receive a prompt to confirm the action, type **Y**, and then press Enter. The command will take a few minutes to complete.
14. Switch to **LON-DC1**.
15. On **LON-DC1**, open the **Group Policy Management** console, and verify that three GPOs that begin with **IPAM_** are created and linked to the **Adatum.com** domain. For the **IPAM_DNS GPO**, verify that the security filtering section is empty.
16. Switch to **LON-SVR2**.
17. On **LON-SVR2**, in **Server Manager**, right-click the line with **LON-DC1**, and then click **Edit Server**.
18. Set the **Manageability status** field to **Managed**, and then click **OK**.
19. Switch to **LON-DC1**.

20. On **LON-DC1**, in the **Group Policy Management** console, verify that the **IPAM_DNS GPO** security filtering section now contains **LON-DC1\$**.
21. Open a command prompt. At the command prompt, type **Gpupdate /force**, and then press Enter.
22. Wait for the Gpupdate process to complete, and then close the command prompt.
23. Switch to **LON-SVR2**.
24. On **LON-SVR2**, in the **Server Manager** window, right-click the line with **LON-DC1**, and then click **Refresh Server Access Status**. Discovery might take 5 to 10 minutes to run. The yellow status bar indicates when discovery is complete.
25. After discovery is complete, refresh IPv4 by clicking the **Refresh** icon. It might take up to 5 minutes for the status to change. Verify that the **IPAM Access Status** for **LON-DC1** is now **Unblocked**.
26. In the **IPAM Overview** content pane, click **Retrieve data from managed servers**. This action will take a few minutes to complete.

IPAM management and monitoring

One of the design criteria for IPAM is to perform day-to-day DHCP and DNS operations in the **IPAM** console. In Windows Server 2012 and Windows Server 2012 R2, DHCP was the focus. In Windows Server 2016, the DNS management features several enhancements. It is now possible to manage both DNS zones and resource records. In the **IPAM** console, you can use the **MONITOR AND MANAGE** section and the **DNS and DHCP Servers**, **DHCP Scopes**, **DNS Zone Monitoring**, and **Server Groups** views to view and monitor health and configuration for all DNS and DHCP servers that IPAM manages. IPAM uses scheduled tasks to collect data from managed servers periodically. You also can retrieve data on demand by using the **Retrieve All Server Data** option. Address space management supports IPv4 public and private addresses, and IPv6 global and unicast addresses.

You can use the IPAM console to:

- Configure many DHCP properties and values
- Configure DNS zones and resource records
- Monitor DHCP and DNS server health

DHCP server management

From the **IPAM** console, you can manage DHCP servers and also:

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

Starting with Windows Server 2012 R2, you also can:

- Create a new MAC address filter for DHCP.
- Configure, import, activate, and deactivate DHCP policies.
- Perform DHCP failover replication for a DHCP server.
- Set an access scope for a DHCP scope or DHCP server.

DNS server management

You can start the **DNS management** console for any managed DNS server from the IPAM server's central console. After you start the **DNS management** console, you can retrieve server data from the selected set of servers. The **DNS Zone Monitoring** view displays all the forward lookup and reverse lookup zones on all the DNS servers that IPAM is currently managing. For the forward lookup zones, IPAM also displays all the servers that are hosting the zone, the aggregate health of the zone across all these servers, and the zone properties. Furthermore, beginning with Windows Server 2016, you can perform several **DNS Server management** actions in the **IPAM** console, including:

- Collect DNS zones and resource records from DNS servers that run Windows Server 2008 or newer. This does not include DNSSEC resource records.
- Configure all types of resource records by performing actions such as creating, modifying, and deleting. This does not include DNSSEC resource records.
- Configure all types of DNS zones, including primary, secondary, and stub zones. This includes creating, modifying, and deleting zones.
- Perform tasks on secondary and stub zones: both forward lookup zones and reverse lookup zones. You can perform tasks such as **Transfer from Master** or **Transfer new copy of zone from Master**.
- RBAC for the DNS configuration.
- Collect and configure conditional forwarders. This includes creating, modifying, and deleting conditional forwarders.

Monitoring DHCP and DNS

By using IPAM, you can monitor DHCP and DNS servers from any physical location of the enterprise. One of IPAM's primary benefits is that you can use it to manage multiple DHCP servers simultaneously or DHCP scopes spread across one or more DHCP servers.

You can use the IPAM monitoring view to have a single-console view of the status and health of selected sets of DNS and DHCP servers. IPAM's monitoring view displays the basic health of servers and recent configuration events that occurred on these servers and you can use it to organize the managed servers into logical server groups.

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

Check Your Knowledge

Question	
Which of the following statements are true? (Choose all that apply.)	
Select the correct answer.	
<input type="checkbox"/>	You can install IPAM on an AD DS domain controller.
<input type="checkbox"/>	You can install IPAM on a domain-joined server.
<input type="checkbox"/>	IPAM can manage servers that are running Windows Server 2012 and newer only.
<input type="checkbox"/>	IPAM can manage servers that are running Windows Server 2008 and newer only.
<input type="checkbox"/>	You must install SQL Server locally on the IPAM server.

Lesson 3

Managing IP address spaces with IPAM

There are multiple phases through which you proceed when you use IPAM to manage IP addresses. You can have IPAM automatically manage IP addresses that DHCP servers issue, or you can create IP address ranges manually that you want to manage. In this lesson, you will learn how to manage all aspects of IPAM: from configuring automatic management to manually adding and updating address information. Finally, you will learn how to monitor your IP address usage and create reports from IPAM.


Lesson Objectives

After completing this lesson, you will be able to:

- Use IPAM to manage IP addressing.
- Explain how to add address spaces to IPAM.
- Explain how to administer IPAM.
- Explain how to implement IPAM reporting and monitoring.

Using IPAM to manage IP addressing

Administrators can use IP address space management in IPAM to manage, track, audit, and report on an organization's IPv4 and IPv6 address spaces. The **IP ADDRESS SPACE** section in the **IPAM** console provides administrators with IP address-utilization statistics and historical trend data that they can use to make informed planning decisions for dynamic, static, and virtual address spaces. Periodically, IPAM performs tasks that discover the address space and utilization data automatically, according to the configuration on the DHCP servers that you manage in IPAM. You also can import IP address information from comma separated value (.csv) files.



IPAM also enables administrators to detect overlapping IP address ranges that are defined on different DHCP servers, find free IP addresses within a range, create DHCP reservations, and create DNS records. IPAM provides a number of ways to filter the view of the IP address space. You can customize how you view and manage the IP address space by using the following views:

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

IP address blocks

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

IP address subnets

After IP address blocks, IP address subnets are the next hierarchical level of address space entities. IPAM summarizes utilization statistics and trends at the IP address subnet level for the IP address ranges that the IP address subnet contains. Additionally, you can create subnets as either physical or virtual. If subnets are virtual, you can assign them to either a provider or a customer virtual network.

IP address ranges

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

IP addresses

IP addresses are the addresses that make up the IP address range. IPAM enables end-to-end lifecycle management of IPv4 and IPv6 addresses, including record synchronization with DHCP and DNS servers. IPAM automatically maps an address to the appropriate range based on the range's start and end address. An IP address is uniquely identifiable by the value of mandatory **Managed By Service** and **Service Instance** options that help IPAM manage and maintain duplicate IP addresses from the same console. You can add or import IP addresses from within the **IPAM** console. In Windows Server 2016, IPAM can create IP addresses from the discovery of PTR records from reverse lookup zones in discovered DNS servers.

IP address inventory

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

IP address range groups

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

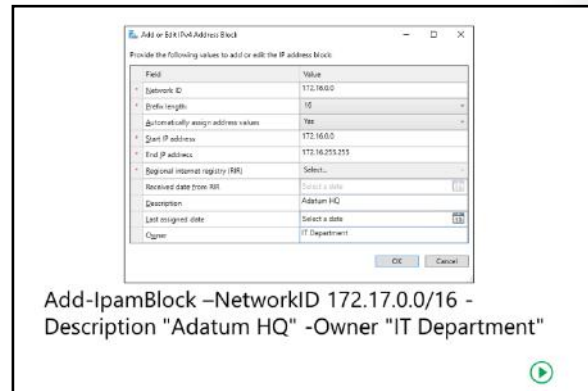


Note: The term *prefix length* is equivalent to using the term *subnet mask* when you define an address range. Windows PowerShell uses prefix length and refers to the routing prefix that Classless Inter-Domain Routing (CIDR) notation uses. For example: 192.168.2.0/24, that is, the 192.168.2.0 network with a prefix length of 24, is equivalent to 192.168.2.0/255.255.255.0, that is, 192.168.2.0 with a network mask of 255.255.255.0.

Adding address spaces to IPAM

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

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



- **IP address block.** When you add an **IP address block**, and you supply the **Network ID** and **Prefix length**, IPAM calculates the **Start IP address** and **End IP address** automatically. Additionally, if you enter a non-private IP address range, you must specify the RIR where the addresses are registered and the registration date range. Optionally, you can add a brief description and an owner.

To add an IP address block, type the Windows PowerShell cmdlet, **Add-IpamBlock**, at a Windows PowerShell prompt, and then press Enter:

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

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

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

There are several optional settings when you add an **IP Address subnet**. You can specify one or more VLANs to be associated with the subnet, virtualized subnet or not, or custom fields such as **AD Site** or **VMM IP Pool Name**. As with the other IP address types, you can add a brief description and an owner.

You can use the following Windows PowerShell cmdlet **Add-IpamSubnet** to add an **IP address subnet**. You can specify if the network type is **NonVirtualized**, **Provider**, or **Customer IP Subnet**. If you add a **Customer IP Subnet**, you must specify the address space to which you add the **Customer IP Subnet**. Type the following cmdlet at a Windows PowerShell prompt, and then press Enter:

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

- **IP address range.** You can use an **IP address range** to subdivide an IP address subnet. When you create an **IP address range**, you must specify the **Network ID**, including the **Prefix length** or **Subnet mask**. Additionally, if an IP address subnet does not exist already that contains the addresses in the IP address range that you create, you can select to have an IP address subnet created

automatically. The other required fields, including **Managed by Service**, **Service Instance**, and **Assignment Type**, will use default values unless otherwise specified. As with the other IP address types, a large variety of custom fields is available to describe the IP address range.

You can use the Windows PowerShell cmdlet **Add-IpamRange** to add an IP Address range. When you use **Add-IpamRange**, you can specify if the network type is **NonVirtualized**, **Provider**, or **Customer IP range**. If you add a **Customer IP range**, you must specify the address space to which you add the **Customer IP range**. Type the following cmdlet at a Windows PowerShell prompt, and then press Enter:

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

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

You can use the Windows PowerShell cmdlet **Add-IpamAddress** to add an IP Address. When you use **Add-IpamAddress**, you also must specify the IP address. Type the following cmdlet at a Windows PowerShell prompt, and then press Enter:

```
Add-IpamAddress -IpAddress <X.X.X.X>
```

Administering IPAM

The day-to-day operations of administering your IPAM solution might include updating address spaces and maintaining your IP address inventory. You can use a text file to import IP address information into IPAM. When you import information from a file, you must include the required fields for the address type, just as you do when you add addresses through the console. The file type that you use is a .csv file with the field names in the first row.

You can import information into custom fields. However, they must exist in IPAM before you import the data, and you must include the defined field name in the first line. The fields are not required to be in any particular order. However, the data must be in the same order as the fields.

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

- You can enclose field names and data in quotation marks.
- Field names and data can contain spaces.

Daily operations in IPAM include:

- Using .csv files to import individual IP addresses, IP address blocks, subnets, and ranges
- Finding available IP addresses
- Reclaiming used IP addresses
- Creating, modifying, and editing IP addresses
- Creating and deleting:
 - DHCP reservations
 - DNS Host records
 - DNS PTR records

- Field names and data are not case-sensitive.
- Data must be valid for the field to which it belongs.

For example, you can use the following entries in a text file to import two addresses into the IPAM database that manages a DHCP server named **DHCP1.adatum.com**:

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

When you import IP address blocks, subnets, and ranges from a file, you combine the network ID and network prefix length in a single field named **Network**. For example, to import an IP Address block of 65.52.0.0/14 assigned by the American Registry for Internet Numbers (ARIN) regional authority, use the following entries in a text file:

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

If a required field is missing or if you try to import the wrong data type for a field, IPAM creates an error report in the signed-in user's **Documents** folder. The mandatory fields for importing data are as follows:

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

Importing and updating IP address ranges

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

Exporting IP information from DNS and DHCP

Because of the potential impact of importing information from your existing DHCP and DNS servers automatically, IPAM does not import all the available IP information automatically. To import existing DNS and DHCP IP address information into IPAM, you must first export the information into text files.



Additional Reading: For more information, refer to: "Manage IPAM" at:
<http://aka.ms/sjmpco>

Finding and allocating available IP addresses

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

When you complete the searches, you can allocate one of the addresses that you found. By default, IPAM highlights the last address it finds, and you can allocate that address or select a different address. Choose each section to configure the IP address as needed. IPAM automatically fills some of the fields in the **Basic Configurations** section with the selected IP address and the default values for the mandatory fields.

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

Reclaiming IP addresses

When you no longer need IP addresses that you added manually, you can *reclaim* them, which makes them available for use with other devices. Additionally, the reclaim operation cleans DHCP reservations and DNS records on managed DNS and DHCP servers. There are two ways to reclaim IP addresses:

- To reclaim IP addresses in a range, in the **IP ADDRESS SPACE**, change the current view to **IP Address Ranges**. The **Reclaim IP Addresses** task is available if you right-click the desired IP address range. If you choose this operation, it opens the **Reclaim IP Addresses** dialog box.
- The **Reclaim IP Addresses** dialog box displays all the utilized IP addresses for the range, the **IP Address State**, and additional information, such as the **Device Name** and **Device Type**. When you have determined the IP addresses that you want to reclaim, select **IP addresses**, and then click the **Reclaim** button. By default, this operation removes the DNS resource records and DHCP reservations.

Editing IP addresses

You can use the **Edit IP Address** dialog box to add information to an IP address or change the currently configured information. You can modify all aspects of the IP address information.

Creating records

There are three options that you can use to create records for an IP address, including:

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

Deleting records

You can use the following four options to delete IP addresses and related information:

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

Implementing IPAM reporting and monitoring

You can use the IPAM address space management feature to efficiently view, monitor, and manage the IP address space on the network. The Event Catalog allows you to see configuration changes to IPAM and on DHCP servers that you are managing by using IPAM, and you also can track your DHCP leases.

Utilization monitoring

IPAM maintains utilization data for IP address ranges and blocks and for IP range groups. You can configure thresholds for utilization percentages of the IP address space and then use those thresholds to determine under-utilization and over-utilization.

You can perform utilization trend building and reporting for IPv4 address ranges, blocks, and range groups. Use the utilization trend window to view trends over specific periods, such as daily, weekly, monthly, or annually. You also can view trends over custom date ranges. Utilization data from managed DHCP scopes is autodiscovered, and you can view this data.

The event catalog

The IPAM event catalog provides a centralized repository, so that you can audit all configuration changes that you perform on DHCP servers that you manage from this IPAM management console. The IPAM configuration events console gathers all of the configuration events. These configuration event catalogs allow you to view, query, and generate reports of the consolidated configuration changes, along with details specific to each record.

IP address tracking

You can use IPAM to track IP addresses with associated DHCP lease events on managed DHCP servers and on user and computer authentication events on AD DS domain controllers that you manage in IPAM. You can track IP addresses by:

- IP address
- Client ID (MAC address)
- Host name
- User name

To find relevant events, in the **IPAM** navigation pane, click **EVENT CATALOG**, and then select how you want to track IP addresses. If you want to track IP addresses by host name, click **By Host Name**, and then in the **By Host Name** text box, type the host name that you want to track. Type a start and end date in the relevant text boxes, and then click **Search**. You now will see a list of events for the host name that you entered, for the period that you specified. You can export the events to a .csv file for further investigation or reporting.

The **IPAM** console does not contain extensive reporting tools. You can track and audit events and configuration changes. You can use the utilization trend of IP addresses in IP ranges to predict when to expand **DHCP** scopes. You must implement additional tools, if you want more reporting options. This can include the reporting infrastructure of Microsoft SQL Server Reporting Services.

- IPAM reporting includes:
 - Monitoring IP address space utilization
 - IP address tracking
- IPAM monitoring includes:
 - Monitoring DNS and DHCP health
 - Using the event catalog to view a centralized repository for all configuration changes
- Use Microsoft SQL Server Reporting Services for extensive reporting

Lab: Implementing network services

Scenario

A. Datum Corporation has deployed several new branch offices and significantly increased the number of users in the organization. A. Datum also has expanded the number of partner organizations and customers that are accessing A. Datum websites and applications. This expansion has resulted in increasing complexity of the A. Datum network infrastructure, which means that the organization must be much more aware of network-level security.

You are responsible for implementing several new features in the Windows Server 2016 environment, including new DNS and DHCP features, and then implementing IPAM to simplify the process for managing the IP infrastructure.

Objectives

After completing this lab, you will be able to:

- Configure DNS policies.
- Configure DHCP failover.
- Configure IPAM.

Lab Setup

Estimated Time: 40 minutes

Virtual machines: **20743A-LON-DC1**, **20743A-LON-SVR1**, **20743A-LON-SVR2**, **20743A-LON-CL1**, and **20743A-TREY-DC1**

User name: **Adatum\Administrator**

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

User name: **TreyResearch\Administrator**

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

For this lab, you will use the available VM environment. Before you begin the lab, you must complete the following procedure:

1. On the host computer, start **Hyper-V Manager**.
2. In Hyper-V Manager, click **20743A-LON-DC1**, and then in the **Actions** pane, click **Start**.
3. In the **Actions** pane, click **Connect**. Wait until the VM starts.
4. Sign in by using the following credentials:
 - User name: **Adatum\Administrator**
 - Password: **Pa\$\$w0rd**
5. Repeat steps 2 through 4 for **20743A-LON-SVR1**, **20743A-LON-SVR2**, and **20743A-LON-CL1**.
6. Repeat steps 2 through 4 for **20743A-TREY-DC1**, but sign in as **TreyResearch\Administrator** with the password **Pa\$\$w0rd**.

Exercise 1: Configuring DNS policies

Scenario

You want to make DNS zone management easier. You want to configure split-brain DNS by using the new DNS policies in Windows Server 2016 so that partners can connect to a different web server. Trey Research has agreed to test the new DNS setup. If you configure the DNS policies correctly, clients in the TreyResearch domain should receive different name-resolution answers than clients that are internal to the Adatum domain.

The main tasks for this exercise are as follows:

1. Check DNS name resolution before configuring DNS policies.
2. Configure DNS policies.
3. Check DNS name resolution after configuring DNS policies.

► Task 1: Check DNS name resolution before configuring DNS policies

1. Switch to **LON-CL1**.
2. On **LON-CL1**, flush the DNS client cache, and then perform name resolution on **www.adatum.com**. Verify that the name resolves to an IP address of **172.16.0.10**.
3. Switch to **TREY-DC1**.
4. On **TREY-DC1**, start a **Windows PowerShell** window, and in the **Windows PowerShell** window, type the following command, and then press Enter:

```
Add-DnsServerConditionalForwarderZone -Name "adatum.com" -MasterServers 172.16.0.10
```

5. In the **Windows PowerShell** window, type the following three commands, pressing Enter after each command:

```
Clear-DNSServerCache
Clear-DNSClientCache
Resolve-DNSName adatum.com
```

6. When you receive a prompt, type **Y**, and then press Enter. Verify that the last command returns an IP address of **172.16.0.10**. If the last command returns an error, retry the last command.

► Task 2: Configure DNS policies

1. Switch to **LON-DC1**.
2. On **LON-DC1**, start a **Windows PowerShell** window, and then in the **Windows PowerShell** window, type the following two commands, pressing Enter after each command:

```
Add-DnsServerZoneScope -ZoneName "adatum.com" -Name "partners"
Add-DnsServerClientSubnet -Name "TreyPartner" -IPv4Subnet 172.16.10.0/24
```

3. In the **Windows PowerShell** window, type the following command, and then press Enter:

```
Add-DnsServerResourceRecord -ZoneName "adatum.com" -A -Name "www" -IPv4Address
"131.107.0.200" -ZoneScope "partners"
```

4. In the **Windows PowerShell** window, type the following command, and then press Enter:

```
Add-DnsServerQueryResolutionPolicy -Name "SplitBrainZonePolicy" -Action ALLOW -
ClientSubnet "eq,TreyPartner" -ZoneScope "partners,1" -ZoneName adatum.com
```

► **Task 3: Check DNS name resolution after configuring DNS policies**

1. Switch to **LON-CL1**.
2. On **LON-CL1**, flush the DNS client cache, and then perform name resolution on **www.adatum.com**. Verify that the name still resolves to an IP address of **172.16.0.10**.
3. Switch to **TREY-DC1**.
4. On **TREY-DC1**, in the **Windows PowerShell** window, type the following three commands, pressing Enter after each command:

```
Clear-DNSServerCache
Clear-DNSClientCache
Resolve-DNSName www.adatum.com
```

5. Verify that the last command returns an IP address of **131.107.0.200**.

Results: After completing this exercise, you should have successfully configured DNS policies for split-brain DNS and then tested that the policies work.

Exercise 2: Configuring DHCP failover

Scenario

The A. Datum network expansion, and the increased availability and security requirements, necessitate configuration of DHCP failover.

The main tasks for this exercise are as follows:

1. Install DHCP on an additional server.
2. Configure DHCP scopes.
3. Configure DHCP failover.
4. Test DHCP failover.

► **Task 1: Install DHCP on an additional server**

1. Switch to **LON-SVR1**.
2. On **LON-SVR1**, click **Start**, and then click **Server Manager**.
3. Use the **Server Manager** console to install the **DHCP server** role with dependent roles and features.
4. In the menu bar, click the **Notifications** flag, and then click **Complete DHCP Configuration**.
5. Complete the **DHCP Post-Install configuration** wizard, accepting the default values.
6. Use the **Server Manager** console to start the **DHCP** console.
7. In the **DHCP** console, verify that no scopes exist.

► **Task 2: Configure DHCP scopes**

1. Switch to **LON-DC1**.
2. Open **Active Directory Users and Computers**, and then add the **LON-DC1** and **LON-SVR1** computer accounts to the **DnsUpdateProxy** group. The group is located in the **Users** container.
3. Open **Windows PowerShell**.

4. In the **Windows PowerShell** window, type the following two commands, pressing Enter after each command:

```
Add-DhcpServerv4Scope -Name 'Building 2 Scope' -StartRange 192.168.2.11 -EndRange 192.168.2.100 -SubnetMask 255.255.255.0 -Description 'For new building at Adatum HQ'
Set-DhcpServerv4OptionValue -ScopeID 192.168.2.0 -OptionId 6 -Value 172.16.0.9 -Force
```

5. Start the **DHCP** console, and then verify that the scope is listed.

► Task 3: Configure DHCP failover

1. Switch to **LON-SVR1**.
2. On **LON-SVR1**, open the **DHCP** console, and add **LON-DC1** as a server.
3. In the **lon-dc1** node, in **IPv4**, launch the **Configure Failover Wizard**.
4. Configure failover replication with the following settings:
 - Partner server: **172.16.0.11**
 - Relationship Name: **Adatum**
 - Maximum Client Lead Time: **15 minutes**
 - Mode: **Load balance**
 - Load Balance Percentage: **50%**
 - State Switchover Interval: **60 minutes**
 - Message authentication shared secret: **Pa\$\$w0rd**
5. Complete the **Configure Failover Wizard**.
6. In the **lon-svr1.adatum.com** node, refresh the **IPv4** node, and then note that the two scopes are **configured** and **active**.

► Task 4: Test DHCP failover

1. Switch to **LON-CL1**.
2. On **LON-CL1**, modify the network settings to receive an IP address from a DHCP server.
3. Record your IP address and DHCP server IP address.
4. Switch to **LON-SVR1**.
5. On **LON-SVR1**, in the **DHCP** console, stop the DHCP server service for the server that provided the DHCP lease to **LON-CL1**.
6. Switch to **LON-CL1**.
7. On **LON-CL1**, renew your IP address, and then record your IP address and DHCP server IP address. The DHCP server IP address should have changed, but the IP address of **LON.CL1** should be the same.
8. Switch to **LON-SVR1**.

9. On **LON-SVR1**, in the **DHCP** console, start the DHCP server service that you stopped previously for the DHCP server.

Results: After completing this exercise, you should have:

- Installed an additional DHCP server.
- Created a new DHCP scope.
- Configured DHCP failover.
- Verified DHCP failover.

Exercise 3: Configuring IPAM

Scenario

Because of A. Datum's expansion, you find it difficult to manage your network and IP address usage. You must implement IPAM to simplify how you manage your IP infrastructure. You need to prepare for a new building that will open soon and for the installation of a new domain controller. You want to try this using the **IPAM** console instead of the tools that you typically use.

The main tasks for this exercise are as follows:

1. Install the IPAM feature.
2. Configure IPAM provisioning.
3. Configure IPAM server discovery.
4. Configure managed servers.
5. Configure and verify a new DHCP scope with IPAM.
6. Import IP addresses into IPAM.
7. Manage DHCP scopes in IPAM.
8. Manage DNS zones in IPAM.
9. Implement IP address tracking.
10. Prepare for the next module.

► Task 1: Install the IPAM feature

1. Switch to **LON-SVR2**.
2. On **LON-SVR2**, open **Server Manager**, and then install the **IPAM Server** feature by using the **Add Roles and Features Wizard**.

► Task 2: Configure IPAM provisioning

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

► Task 3: Configure IPAM server discovery

1. On **LON-SVR2**, in the **IPAM Overview** pane, configure server discovery for the Adatum domain.
2. In the **IPAM Overview** pane, start the server discovery process. Discovery might take 5 to 10 minutes to run. The yellow status bar will indicate when discovery is complete.

► Task 4: Configure managed servers

1. On **LON-SVR2**, in the **IPAM Overview** pane, add the servers that you need to manage. Verify that IPAM access is blocked currently for both **LON-DC1** and **LON-SVR1**.
2. Use Windows PowerShell to grant the IPAM server permission to manage by running the following command:

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

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

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

1. On **LON-SVR2**, use IPAM to create a new DHCP scope on **LON-SVR1** with the following parameters:
 - Scope Name: **Building 3 Scope**
 - Scope start address: **10.0.0.50**
 - Scope end address: **10.0.0.100**
 - Subnet mask: **255.255.255.0**
 - Default gateway (Option 003): **10.0.0.1**
 - DNS Servers (Option 006): **172.16.0.9**
2. Use IPAM to configure failover for the Building 3 Scope on **LON-DC1** with the following parameters:
 - Configuration option: **Use an existing relationship**
 - Relationship name: **Adatum**
3. Switch to **LON-DC1**.
4. On **LON-DC1**, verify the scope in the **DHCP** console.

► Task 6: Import IP addresses into IPAM

1. Switch to **LON-SVR2**.
2. In the **IPAM** navigation pane, under **IP ADDRESS SPACE**, click **IP Address Inventory**.
3. Import into IP addresses the file **\\LON-SVR1\Labfiles\Mod05\IP-addresses.csv**.

4. In the **Import IP Addresses** dialog box, verify that 30 records were successfully imported.
5. In content pane, verify that the list contains 30 records.

► **Task 7: Manage DHCP scopes in IPAM**

1. On **LON-SVR2**, in the **IPAM** navigation pane, under **MONITOR AND MANAGE**, navigate to **DHCP Scopes**.
2. Search for DHCP scopes that contain the **[006] DNS Servers** option with a value of **172.16.0.9**.
3. Select **Building 2 Scope** and **Building 3 Scope** that in the **Server Name** column have the value **lon-svr1.adatum.com**.
4. Edit the scopes, perform a find and replace for option **006 DNS Servers**, and then replace the value **172.16.0.9** with the value **172.16.0.10**.

► **Task 8: Manage DNS zones in IPAM**

1. On **LON-SVR2**, in the **IPAM** navigation pane, under **MONITOR AND MANAGE**, click **DNS Zones**.
2. Start the **Retrieve Server Data** task.
3. After the refresh completes, click the **Server Manager** console **refresh** button. Verify that the zone status for the **Adatum.com** zone is now green.
4. Edit the **Adatum.com** zone, and then add an NS record for **LON-DC2.Adatum.com**.
5. Using the **IPAM** console, delete the **www** resource record.
6. Using the **IPAM** console, add an A record to the **Adatum.com DNS zone** for **lon-dc2.adatum.com** with an IP address of **172.16.0.20**.
7. Switch to **LON-DC1**.
8. On **LON-DC1**, start the **DNS Manager** console, and then verify that for the **Adatum.com zone** there is both an NS record for **lon-dc2.adatum.com**, an A record for **lon-dc2** and that the A record for **www** is missing.

► **Task 9: Implement IP address tracking**

1. Switch to **LON-SVR2**.
2. On **LON-SVR2**, in the **Server Manager** console, under **EVENT CATALOG**, navigate to **By Host Name**.
3. Retrieve **Event Catalog Data** to get the latest events from the **DHCP** servers. The yellow status bar indicates when discovery is complete.
4. Search for events with the following criteria:
 - By Host Name: **LON-CL1**.
 - First date text box: Yesterday's date in the format m/d/yyyy, for example **12/31/2016**.
 - Second date text box: Tomorrow's date in the same format.
5. Verify that at least one event is shown below the search.

► Task 10: Prepare for the next module

1. On the host computer, start the **Hyper-V Manager**.
2. In the **Virtual Machines** list, right-click **20743A-LON-DC1**, and then click **Revert**.
3. In the **Revert VM** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20743A-LON-SVR1**, **20743A-LON-SVR2**, **20743A-LON-CL1**, and **20743A-TREY-DC1**.

Results: After completing this exercise, you should have:

- Installed and configured the IPAM feature.
- Configured IPAM provisioning.
- Configured IPAM server discovery.
- Configured managed servers.
- Configured and verified a new DHCP scope with IPAM.
- Imported IP addresses into IPAM.
- Managed DHCP scopes in IPAM.
- Managed DNS scopes in IPAM.
- Implemented IP address tracking.

Question: Will you be implementing DNS policies in your DNS infrastructure? Discuss your answers with the rest of the students.

Question: What is the difference between a centralized and a distributed IPAM topology?

Module Review and Takeaways

Review Questions

Question: How many DHCP servers can you configure as part of a DHCP failover relationship?

Question: Describe a limitation of IPAM.

Real-world Issues and Scenarios

Scenario: What are some scenarios in which you would use DNS policies?

Answer: You can use DNS policies to configure split-brain DNS, DNS load balancing, and DNS responses based on criteria such as time of day, client IP address, DNS server IP address, and query type.

Scenario: What are some methods that you can use to guard against DHCP failures?

Answer: You can use DHCP failover protection, a DHCP split-scope solution, or you can cluster the DHCP servers.

Tools

The following table includes the tools that are needed for this module:

Tool	Use	Where to find it
The DNS management console	Configure all aspects of DNS.	In Server Manager, in the Tools drop-down list box.
The DHCP console	Configure all aspects of DHCP.	In Server Manager, in the Tools drop-down list box.
The IPAM console	Configure IP address management.	In Server Manager.

MCT USE ONLY. STUDENT USE PROHIBITED

Module 6

Implementing Hyper-V

Contents:

Module Overview	6-1
Lesson 1: Configure the Hyper-V role in Windows Server 2016	6-2
Lesson 2: Configuring Hyper-V storage	6-11
Lesson 3: Configuring Hyper-V networking	6-21
Lesson 4: Configuring Hyper-V virtual machines	6-28
Lab: Implementing server virtualization with Hyper-V	6-38
Module Review and Takeaways	6-45

Module Overview

A decade ago, server virtualization was rarely deployed on organizational networks. Today, it is a core networking technology. Server administrators must be able to distinguish which server workloads might run effectively on virtual machines and which need to remain in a traditional, physical deployment. This module introduces you to the new features of the Hyper-V role in Windows Server 2016, the components of the role, and the best practices for deploying the role.



Note: Many of the features that this module describes are also available in Windows Server 2012 R2. Some features also are available in Windows Server 2012. This module explicitly states the features that are new to Windows Server 2016.

Objectives

After completing this module, you will be able to:

- Configure the Hyper-V role in Windows Server 2016.
- Configure Hyper-V storage.
- Configure Hyper-V networking.
- Configure Hyper-V virtual machines.

Lesson 1

Configure the Hyper-V role in Windows Server 2016

The Hyper-V role has undergone substantial changes in Windows Server 2016. New features, such as shielded virtual machines, production checkpoints, hot add and hot remove for network adapters and memory, and Windows PowerShell Direct, provide administrators new options for managing and deploying virtualization. In this lesson, you will learn about the new features in Windows Server 2016 Hyper-V. You will also learn about Hyper-V integration services and the factors that you need to consider when configuring Hyper-V hosts.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the new features in Windows Server 2012 R2 Hyper-V.
- Describe the new features in Windows Server 2016 Hyper-V.
- Describe the improvements in Hyper-V Manager.
- Describe the prerequisites and requirements for installing Hyper-V.
- Install and configure the Hyper-V role.
- Describe the best practices for configuring Hyper-V hosts.
- Describe nested virtualization.
- Describe migration to Microsoft Azure virtual machines.

New features in Windows Server 2012 R2 Hyper-V

The Hyper-V role in Windows Server 2012 R2 includes a large number of improvements and new features that are not available in Windows Server 2012. The following table lists some of these new features.

New or improved?	Feature
New in Windows Server 2012 R2	<ul style="list-style-type: none"> • Shared virtual hard disk • Automatic virtual machine activation • Enhanced session mode • Storage quality of service • Virtual machine generation
Improved in Windows Server 2012 R2	<ul style="list-style-type: none"> • Resize virtual hard disk • Live migration • Protected network and storage • Guest service integration services • Export • Hyper-V Replica • Linux support • Management

Feature	Description
Shared virtual hard disk	Use this feature to cluster virtual machines by using shared virtual hard disk (.vhdx format) files.
Automatic virtual machine activation	Use this feature to activate virtual machines automatically on computers that are running the Datacenter edition of Windows Server 2012 R2 or newer Windows Server operating systems.

Feature	Description
Enhanced session mode	Use this feature to provide support for redirecting an increased number of local resources.
Storage quality of service	Use this feature to specify the minimum and maximum I/O loads in terms of I/O operations per second (IOPS) on a per-virtual-hard-disk basis.
Virtual machine generation	Use this feature to provide support for Generation 1 and Generation 2 virtual machines.

The following table lists some of the improved features in Windows Server 2012 R2 Hyper-V.

Feature	Improvement
Resize virtual hard disk	Use this feature to resize virtual hard disks while a virtual machine is running.
Live migration	This feature provides improved performance, including compression of virtual machine RAM and cross-version live migration between Windows Server 2012 Hyper-V and Windows Server 2012 R2 Hyper-V.
Protected network and storage	This feature provides virtual network adapter protection and virtual machine storage protection.
Guest service integration services	This feature provides the ability to copy files to a virtual machine without using a network connection or having to shut down the virtual machine.
Export	Use this feature to export a virtual machine or virtual machine checkpoint while the virtual machine is running.
Hyper-V Replica	This feature supports extended replication and configurable replication frequency.
Linux support	This feature provides support for Linux virtual machine backup and dynamic memory.
Management	This feature provides support for managing Hyper-V on Windows Server 2012 R2 from computers that run Windows 8 or Windows Server 2012.



Additional Reading: For more information, refer to: "What's New in Hyper-V in Windows Server 2012 R2" at: <http://aka.ms/y3gvwz>

New features in Windows Server 2016 Hyper-V

The Hyper-V role in Windows Server 2016 includes a large number of improvements and new features that are not available in Windows Server 2012 and Windows Server 2012 R2. The following tables list some of these features and the improvements to previous features.

The following table lists some of the new features in Windows Server 2016 Hyper-V.


- Host Resource Protection
- Nested Virtualization
- PowerShell Direct
- Shielded Virtual Machines
- Virtual TPM/Bitlocker
- Hot add/remove memory and network adapters
- Checkpoints
- Linux
- Hyper-V Manager
- Storage Quality of Service


Feature	Description
Windows Server Hyper-V Containers and Docker	This feature provides a more isolated container than the normal Windows Server container by using Hyper-V.
Host Resource Protection	This feature helps protect virtual machines against malicious activities by detecting increased resource utilization of an excessive amount that could affect other virtual machines by consuming all resources available on that host.
Nested virtualization	Use this feature to create a Hyper-V host from virtual machine guests.
Windows PowerShell Direct	Use this feature to run commands in the Windows PowerShell command-line interface on virtual machines, without any special configuration or remoting.
Shielded virtual machines	Use this feature to encrypt virtual machines and force the virtual machine to run on Host Guardian Service clients only.
Virtual trusted platform modules (TPMs) and BitLocker Drive Encryption on virtual machines	Use this feature to enable TPMs and BitLocker on virtual machines.

The following table lists some of the improved features in Windows Server 2016 Hyper-V.

Feature	Improvement
Host add/remove memory and network adapters	Static memory and network adapters can now be added while a virtual machine is online.
Checkpoints	Checkpoints now support production checkpoints.
Linux	Some of the latest versions of Linux can now use Secure Boot with Hyper-V. These include CentOS 7.0 and later versions, Red Hat Enterprise Linux 7.0 and later versions, SUSE Linux Enterprise Server 12 and later versions, and Ubuntu 14.04 and later versions.
Hyper-V Manager	This feature supports alternate credentials.

Feature	Improvement
Storage quality of service	This feature has expanded to include Scale-Out File Servers and centrally managed Quality of Service (QoS) policies.

 **Additional Reading:** For more information, refer to: "What's new in Hyper-V on Windows Server 2016 Technical Preview" at: <http://aka.ms/t8ye3k>

 **Note:** Subsequent topics in this module discuss these new features in more detail.

Hyper-V Manager improvements

When you remotely administer a Hyper-V environment, you will often use the Hyper-V Manager snap-in to the Microsoft Management Console. Hyper-V Manager in Windows Server 2016 features the following improvements:

- Management protocol. Hyper-V Manager now supports connections over the Web Services Management Protocol (WS-Management Protocol). This allows Hyper-V Manager to communicate by using the Kerberos protocol, NT LAN Manager (NTLM), or Credential Security Support Provider. When using Credential Security Support Provider, you remove the need for Active Directory Domain Services (AD DS) delegation. This makes it easier to enable remote administration because WS-Management Protocol communicates over port 80, the default open port.
- Alternate credentials support. Communicating over WS-Management Protocol opens up the ability to add different credentials in Hyper-V Manager and to save the credentials for ease of management. Alternative credentials, however, only work with Windows 10 and Windows Server 2016 hosts. Older servers running Hyper-V do not support the WS-Management Protocol for Hyper-V Manager communication.
- Previous version support. When running Hyper-V Manager on Windows Server 2016 and Windows 10, you will still be able to manage your previous hosts that are running Windows 2012, Windows 2012 R2, Windows 8, or Windows 8.1.

- Management protocol: WS-Management protocol, allowing Kerberos, NTLM, or Credential Security Support Provider authentication
- Alternate credentials support: connect to hosts with various authentication requirements
- Previous version support: continue managing older hosts

Prerequisites and requirements for installing Hyper-V on Windows Server 2016

To install Hyper-V, the host computer must meet the following requirements:

- It should have an x64 processor with hardware-assisted virtualization.
- Both hardware-assisted virtualization and DEP must be enabled in the BIOS or UEFI.
- The central processing unit (CPU) must support second-level address translation (SLAT) for the processor and VM Monitor Mode extensions. Unlike in Windows Server 2012 R2 where SLAT was optional and beneficial, you must install Hyper-V on Windows Server 2016.
- It needs a minimum of 4 gigabytes (GB) of RAM.

- Server hardware must support:
 - Hardware-assisted virtualization
 - Data execution prevention
 - SLAT and VM Monitor Mode extensions
 - 4 GB of RAM
- Hardware must be adequate to support the needs of virtual machines with respect to:
 - Memory
 - Disk I/O
 - Processing capability
 - Network throughput, typically NIC Teaming

When deciding on the server hardware on which you plan to install the Hyper-V role, ensure that the server meets the following requirements:

- The server must have enough memory to support the memory requirements of all the virtual machines that must run concurrently. The server also must have enough memory to run the host Windows Server 2016 operating system. When assessing memory requirements, reserve 1 GB of RAM for the host operating system.
- Storage subsystem performance must meet the I/O needs of the guest virtual machines. You might need to place different virtual machines on separate physical disks to deploy a high-performance redundant array of independent disks (RAID), solid-state drives (SSDs), a hybrid SSD, or a combination of all three.
- The CPU capacity of the host server must meet the requirements of the guest virtual machines.
- The host server's network adapters must be able to support the network throughput requirements of the guest virtual machines. This might require installing multiple network adapters and using network interface card (NIC) Teaming for virtual machines that have high network-use requirements.

Demonstration: Installing and configuring the Hyper-V role

It is necessary to start a traditionally deployed server to run this demonstration or to configure a nested virtualization virtual machine host.

Demonstration Steps

1. Sign in to **LON-HOST1**.
2. Open **Hyper-V Manager**.
3. In the **Hyper-V Settings** dialog box, review the following settings:
 - Virtual Hard Disks
 - Virtual Machines
 - Physical GPUs
 - NUMA Spanning

Best practices for configuring Hyper-V hosts

You should consider the following best practices when provisioning Windows Server 2016 to function as a Hyper-V host:

- Provision the host with adequate hardware.
- Deploy virtual machines on separate disks or Cluster Shared Volumes (CSVs) if using shared storage.
- Do not collocate other server roles.
- Manage Hyper-V remotely.
- Run Hyper-V by using the Server Core installation or Nano Server.
- Run the Best Practices Analyzer and resource metering.
- Use Generation 2 virtual machines if they are supported by the guest operating system.

- Provision the Hyper-V host with adequate hardware resources
- Deploy virtual machines on separate disks or CSVs if using shared storage
- Do not collocate other Windows Server roles on the Hyper-V host
- Manage Hyper-V remotely
- Run Hyper-V by using the Server Core or Nano configuration
- Use resource metering and the Best Practices Analyzer
- Use Generation 2 virtual machines where possible

Provision the host with adequate hardware

Perhaps the most important best practice is to ensure that you have provisioned the Hyper-V host with adequate hardware. You should ensure that there is appropriate processing capacity, an appropriate amount of RAM, and fast and redundant storage. You should ensure that you have provisioned the Hyper-V host with multiple network adapters that you configure as a team. Inadequately provisioning the Hyper-V host with hardware affects the performance of all virtual machines that are hosted on the server.

Deploy virtual machines on separate disks

You should use separate disks to host virtual machine files rather than storing virtual machine files on the same disk as the host operating system files. Doing this minimizes contention and ensures that read/write operations occurring on virtual machine files do not conflict with read/write operations occurring at the host operating system level. It also minimizes the chance that virtual machine hard disks will grow to consume all available space on an operating system volume.

The impacts on performance are reduced if you deploy to a disk that uses striping, such as a RAID 1+0 array. If you are using shared storage, you can provision multiple virtual machines on the same logical unit number (LUN) if you use CSVs. However, choosing between separate LUNs for each virtual machine or a shared LUN depends heavily on virtual machine workload and host configuration.

Do not collocate other server roles

You should ensure that Hyper-V is the only server role that is deployed on the server. Do not collocate the Hyper-V role with other roles, such as the domain controller or file server role. Each role that you deploy on a server requires resources, and when deploying Hyper-V, you want to ensure that virtual machines have access to as much of a host server's resources as possible. If locating these roles on the same hardware is necessary, deploy these roles as virtual machines rather than installing them on the physical host.

Manage Hyper-V remotely

When you sign in locally to a server, your session consumes server resources. If you configure remote management of a Hyper-V server instead of performing administrative tasks by signing in locally, you can ensure that all possible resources on the Hyper-V host are available to the hosted virtual machines. You also should restrict access to the Hyper-V server so that only administrators who are responsible for virtual machine management can make connections. A configuration error on a Hyper-V host can cause downtime for all hosted virtual machines.

Run Hyper-V by using the Server Core or Nano configuration

You should run Hyper-V by using the Server Core or Nano configuration. Doing so provides the following benefits:

- Running Windows Server 2016 in the Server Core or Nano configuration minimizes hardware-resource utilization for the host operating system. As a result, the hosted virtual machines have more hardware resources.
- Server Core and Nano configurations require fewer software updates, which in turn require fewer restarts. When you restart a Hyper-V host and it is unavailable, all virtual machines that the server hosts also become unavailable. To prevent this issue, we recommend configurations that use clustered environments. Because a Hyper-V host can host many critical servers as virtual machines, ensure that you minimize downtime.

Run the Best Practices Analyzer and use resource metering

You can use the Best Practices Analyzer to determine any specific configuration issues that you should address. You can use resource metering, introduced in Windows Server 2012 Hyper-V, to monitor how hosted virtual machines use server resources and determine if specific virtual machines are using a disproportionate amount of a host server's resources. If the performance characteristics of one virtual machine are lowering the performance of other virtual machines that are hosted on the same server, consider migrating that virtual machine to another Hyper-V host.

Use Generation 2 virtual machines if supported by the guest operating system

Generation 2 virtual machines have slightly faster start times than Generation 1 virtual machines. Generation 2 virtual machines use a simplified hardware model and allow advanced features such as:

- Pre-Boot Execution Environment (PXE) boot from a standard network adapter.
- Small computer system interface (SCSI) controller boot.
- Secure boot.



Additional Reading: For more information, refer to: "Tip: 6 Best Practices for Physical Servers Hosting Hyper-V Roles" at: <http://aka.ms/aquwzd>

Nested virtualization

Windows Server 2016 has introduced support for *nested virtualization*. Nested virtualization converts a Hyper-V guest virtual machine to a Hyper-V host so that it can host other guest virtual machines. Doing this can be extremely useful on development and test servers in order to build out Hyper-V host testing scenarios that previously would have required more physical hardware just to run tests.

To enable nested virtualization, you need at least 4 GB of RAM and Windows Server 2016 or Windows 10 as the host operating system.

Additionally, the virtual machine that is running Hyper-V must be the same build as the host.

- Converts a Hyper-V guest virtual machine to a Hyper-V host
- Extremely useful with development and test servers
- Requirements:
 - At least 4 GB of static RAM
 - Windows Server 2016 or Windows 10 as the host operating system
 - The virtual machine that is running Hyper-V must be the same build as the host
- Some features are disabled or will fail after you enable nested virtualization

To enable nested virtualization, run the following script on the host Hyper-V server. "DemoVM" would be the virtual machine on which you want to enable nested virtualization.

```
Invoke-WebRequest https://raw.githubusercontent.com/Microsoft/Virtualization-Documentation/master/hyperv-tools/Nested/Enable-NestedVm.ps1 -OutFile ~/Enable-NestedVm.ps1
~/Enable-NestedVm.ps1 -VmName "DemoVM"
```



Note: This command is accurate for Windows Server 2016 Technical Preview 5. The command might change in future releases.

After enabling nested virtualization, you can install Hyper-V on a virtual machine in the same way that you would for a Hyper-V host. The following features are disabled or will fail after you enable nested virtualization:

- Virtual-based security
- Device Guard
- Dynamic Memory
- Hot add Static Memory
- Checkpoints
- Live migration

Migration to Azure virtual machines

Many organizations are considering moving all or part of their server infrastructure to cloud-based platforms such as Azure. Organizations want to take advantage of features in Azure such as decreased datacenter management, quick scalability, and load flexibility. One of Azure's primary uses can be to build a temporary test or development environment quickly without needing to purchase new hardware. With Windows Server 2016 Hyper-V, you can push virtual machines to Azure.

- Use Azure virtual machines to build a test and development environment
- Two methods for moving virtual machines from Hyper-V to Azure:
 - Virtual hard disk manual moves are free but require a virtual hard disk on which you have run Sysprep
 - Azure Site Recovery:
 - Supported but not designed for single virtual machine moves

You can migrate virtual machines from Hyper-V to Azure without using non-Microsoft tools in two ways.

- Upload any Hyper-V virtual hard disk to Azure:
 - This method requires that you use the System Preparation Tool (Sysprep) and generalize a virtual hard disk. Move the generalized virtual hard disk to an Azure storage account container, and then create an image using the generalized virtual hard disk in cloud storage. Because you need to generalize the virtual hard disk first, this is not a practical solution if you want to migrate a virtual machine that you are currently running in your private environment to Azure.
 - This process is designed mostly for when you want to deploy multiple virtual machines with the same configuration. For instance, you might want to deploy multiple Web servers and be able to add and remove servers based on demand.

- Azure Site Recovery:
 - This is a full disaster-recovery solution for replicating virtual machines from a private cloud to the public cloud. Because of the infrastructure requirements, you would normally not use this option to move just one virtual machine to Azure. This is currently the only supported solution for moving virtual machines to Azure.

When deciding whether to move virtual machines to Azure, consider the following:

- Azure does not allow console access until the virtual machine is online.
- Most Azure virtual machines only have one network adapter and one external IP address.
- Azure currently only supports Generation 1 servers.
- Azure will convert any .vhdx file to a fixed-size .vhd file format when you move it to Azure.



Note: Azure does not support Generation 2 virtual machines natively. However, when using Azure Site Recovery on a Generation 2 virtual machine, Azure will convert the virtual machine to a Generation 1 virtual machine for use in Azure and then convert back to a Generation 2 virtual machine when returning to on-premises hardware.

Question: In a scenario where you are limited to only one physical server, how can you build a test or development environment with multiple hosts and virtual machines on those hosts?

Check Your Knowledge

Question	
When you are configuring a Hyper-V host, what are some best practices guidelines to follow? Select all that apply.	
Select the correct answer.	
<input type="checkbox"/>	Use Generation 2 virtual machines if supported by the guest operating system.
<input type="checkbox"/>	Manage Hyper-V locally.
<input type="checkbox"/>	Run Hyper-V by using the Server Core or Nano configuration.
<input type="checkbox"/>	Do not collocate other server roles.
<input type="checkbox"/>	Provision the host with adequate hardware.

Lesson 2

Configuring Hyper-V storage

Hyper-V provides many different virtual machine storage options. If you know which option is appropriate for a given situation, you can ensure that a virtual machine performs well. If you do not understand the different virtual machine storage options, you might end up deploying virtual hard disks that consume unnecessary space or place an unnecessary performance burden on the host Hyper-V server. This lesson describes different virtual hard disk types, different virtual hard disk formats, and the benefits and limitations of using virtual machine checkpoints.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe virtual hard disk properties in Windows Server 2016 Hyper-V.
- Describe the types of disks and select a virtual hard disk type.
- Describe shared virtual hard disks.
- Convert between virtual hard disk types and resize virtual hard disks.
- Describe the requirements for providing Fibre Channel support in virtual machines.
- Determine where to deploy virtual hard disks.
- Describe the requirements for storing Hyper-V data on Server Message Block (SMB) 3.0 file shares.
- Manage virtual hard disks in Hyper-V.

Virtual hard disk file formats

A virtual hard disk is a special file format that represents a traditional hard disk drive. You can configure a virtual hard disk with partitions and an operating system. Additionally, you can use virtual hard disks with virtual machines, and you can mount virtual hard disks by using Windows Server 2008, Windows 7, or newer operating systems.

Windows Server 2016, Windows Server 2012 R2, and Windows Server 2012 support *starting to virtual hard disks*. You can use this to configure the computer to start into a Windows Server 2016, Windows Server 2012 R2, or Windows Server 2012 operating system, or certain editions of the Windows 10, Windows 8.1, or Windows 8 operating systems that are deployed on a virtual hard disk. You can create a virtual hard disk by using:

- The Hyper-V Manager console.
- The Disk Management console.
- The **Diskpart** command-line tool.
- The **New-VHD** Windows PowerShell cmdlet.

- Windows Server 2012 introduces the new .vhdx format for virtual hard disks, which:
 - Supports bigger disks
 - Is less likely to lose data during unexpected outages
 - Supports better alignment when deployed to a large-sector disk
 - Allows larger block size for dynamic and differencing disks, which provides better performance
- Windows Server 2016 introduces the .vhds format for virtual hard disks
- Windows Server 2012 R2 and Windows Server 2016 support storage QoS for virtual hard disks



Note: Some editions of the Windows 7 and Windows Server 2008 R2 operating systems also support starting to virtual hard disks.

Comparing .vhdx, .vhd, and .vhds

Virtual hard disks use the .vhd extension. Windows Server 2012 introduced the new .vhdx format for virtual hard disks. Compared to the .vhd format that was used in Hyper-V on Windows Server 2008 and Windows Server 2008 R2, the .vhdx format offers the following benefits:

- .vhdx files can be as large as 64 TB; .vhd files are limited to 2 TB.
- The .vhdx file structure minimizes the chance that a disk will become corrupted if the host server suffers an unexpected power outage.
- The .vhdx format supports better alignment when deployed to large-sector disks.
- .vhdx allows larger block size for dynamically expanding and differencing disks, which provides better performance for these workloads.

Windows Server 2016 introduces the .vhds format, which is specific to shared virtual hard disks. If you have upgraded a Windows Server 2008 server or a Windows Server 2008 R2 Hyper-V server to Windows Server 2012, Windows Server 2012 R2, or Windows Server 2016, you can convert an existing .vhd file to .vhdx format by using the Edit Disk tool. It also is possible to convert from .vhdx format to .vhd.



Additional Reading: For more information, refer to: "Hyper-V Virtual Hard Disk Format Overview" at: <http://aka.ms/gsfjsjo>

Storage QoS for Hyper-V and Scale-Out Fileservers

You can use Windows Server 2016 and Windows Server 2012 R2 to configure QoS parameters for virtual hard disks individually for each virtual machine.

In Windows Server 2016 while using a Scale-Out Fileserver to store virtual machines' virtual hard disks you can utilize Storage QoS to monitor and set QoS policies. Instead of going to each virtual machine and applying a QoS to each virtual disk, you can apply a policy to many virtual disks simultaneously by using Windows PowerShell.

Dedicated and aggregated policies

The two different policies can be set for different situations or required levels of Input/Output Operations Per Second.

Dedicated policies operate in the same way that QoS is set in Windows Server 2012 R2. This allows you to set a minimum and maximum number of IOPS per virtual hard disk. However, you can create one policy and set it on multiple virtual machines and the virtual machines' virtual disks at the same time.

By using aggregated policies, you can set a shared minimum and maximum Input/Output Operations Per Second for groups of virtual hard disks.

Types of virtual hard disks

When you configure a virtual hard disk, you can choose one of the following disk types:

- Fixed-size
- Dynamically expanding
- Pass-through
- Differencing

Type of disk	Description
Fixed	All of the hard disk space is allocated during the creation process
Dynamic	The disk itself only uses the amount of space that needs to be allocated, and it grows as necessary
Pass-through	Virtual machines access a physical disk drive rather than use a virtual hard disk
Differencing	The amount of hard disk space that virtual hard disks consume is reduced at the cost of disk performance

Fixed-size virtual hard disks

When you create a fixed-size virtual hard disk, all of the hard disk space is allocated during the creation process. This has the advantage of minimizing fragmentation, which improves virtual hard disk performance when the disks are hosted on traditional storage devices. However, one disadvantage is that a fixed-size .vhd or .vhdx disk requires that all of the space that the virtual hard disk potentially can use is allocated on the host partition. In many situations, you will not know precisely how much disk space a virtual machine needs. If you use fixed-size hard disks, you might end up allocating more space to storage than is required.



Note: Disk fragmentation is a less important issue when virtual hard disks are hosted on RAID volumes or on SSDs. Since its introduction in Windows Server 2008, Hyper-V improvements also minimize the performance differences between dynamically expanding and fixed-size virtual hard disks.

Dynamically expanding virtual hard disks

When you create a dynamically expanding virtual hard disk, you specify a maximum size for the file. The disk itself only uses the amount of space that needs to be allocated, and it grows as necessary. For example, if you create a new virtual machine and specify a dynamically expanding disk, only a small amount of disk space is allocated to the new disk, as follows:

- Approximately 260 KB for a .vhd virtual hard disk
- Approximately 4,096 KB for a .vhdx virtual hard disk

As storage is allocated, such as when you deploy the operating system, the dynamically expanding hard disk grows. If you delete files from a dynamically expanding virtual hard disk, the virtual hard disk file does not shrink. You can only shrink a dynamically expanding virtual hard disk file by performing a shrink operation.

Creating a dynamically expanding virtual hard disk is similar to creating a fixed-size disk. In the **New Virtual Hard Disk Wizard**, on the **Choose Disk Type** page, select **Dynamically expanding size** instead of **Fixed**.

You can create a new dynamically expanding hard disk by using the **New-VHD** Windows PowerShell cmdlet with the *-Dynamic* parameter.

Pass-through disks

Virtual machines use pass-through disks to access a physical disk drive rather than using a virtual hard disk. You can use pass-through disks to connect a virtual machine directly to an Internet SCSI (iSCSI) LUN or a directly attached disk. When you use pass-through disks, the virtual machine must have exclusive access to the target disk. To use the pass-through disks, you must use the host's Disk Management console to take the disk offline. After the disk is offline, you can connect it to one of the virtual machine's disk controllers.

You can attach a pass-through disk by performing the following steps:

1. Ensure that the target hard disk is offline.
2. Use Hyper-V Manager to edit an existing virtual machine's properties.
3. Click an integrated drive electronics (IDE) or SCSI controller, click **Add**, and then click **Hard Drive**.
4. In the **Hard Drive** dialog box, select **Physical Hard Disk**. In the drop-down list box, select the disk that you want to use as the pass-through disk.

Differencing disks

Differencing disks record the changes that are made to a parent disk. You can use differencing disks to reduce the amount of hard disk space that virtual hard disks consume, but that comes at the cost of disk performance. Differencing disks work well with SSD where limited space is available on the drive, and the performance of the disk compensates for the performance drawbacks of using a differencing disk.

Differencing disks have the following properties:

- You can link multiple differencing disks to a single parent disk.
- When you modify the parent disk, all linked differencing disks fail.

You can reconnect a differencing disk to the parent by using the Inspect Disk tool, which is available in the **Actions** pane of the Hyper-V Manager console. You also can use the Inspect Disk tool to locate a differencing disk's parent disk.

To create a differencing disk, follow these steps:

1. Open **Hyper-V Manager**.
2. In the **Actions** pane, click **New**, and then click **Hard Disk**.
3. On the **Before You Begin** page of the **New Virtual Hard Disk** wizard, click **Next**.
4. On the **Choose Disk Format** page, select **VHD**, and then click **Next**.
5. On the **Choose Disk Type** page, select **Differencing**, and then click **Next**.
6. On the **Specify Name and Location** page, provide the location of the parent hard disk, and then click **Finish**.

You can create a differencing hard disk by using the **New-VHD** Windows PowerShell cmdlet. For example, to create a new differencing disk named `c:\diff-disk.vhd` that uses the virtual hard disk `c:\parent.vhd`, run the following Windows PowerShell command:

```
New-VHD c:\diff-disk.vhd -ParentPath C:\parent.vhd
```




Note: You do not have to shut down a virtual machine if you connect the pass-through disk to a virtual machine's SCSI controller. However, if you want to connect to a virtual machine's IDE controller, it is necessary to shut down the virtual machine.

Shared virtual hard disks

In Windows Server 2016, you can share a single virtual hard disk file between multiple virtual machines in a shared storage configuration. For example, the volume that hosts the database files in a Hyper-V guest failover cluster is configured to host a highly available Microsoft SQL Server deployment. Prior to Windows Server 2012 R2, you needed to use storage technologies such as iSCSI or Fibre Channel to provision shared storage for a Hyper-V guest failover cluster.

Shared virtual hard disks have the following properties:

- In Windows Server 2016, a virtual hard disk uses the .vhdx format. You cannot use virtual hard disks with the .vhd format.
- Generation 1 and Generation 2 virtual machines support Hyper-V guest failover clusters by using shared virtual hard disks.
- Virtual machines running Windows Server 2012 and newer operating systems have native support for using shared virtual hard disks as shared storage.
- Virtual machines that run Windows Server 2012 support using shared virtual hard disks as shared storage if Windows Server 2012 R2 or Windows Server 2016 integration services are installed.
- Shared virtual hard disks must be stored on CSVs or a file server with SMB 3.0 file-based storage.



Virtual machine storage resiliency

Windows Server 2016 introduces virtual machine storage resiliency. Before Windows Server 2016, if a transient failure occurred to the storage communication for a virtual machine, it would go into an off state as soon as it could no longer write or read from storage.

Now in Windows Server 2016 when a virtual hard disk (VHD) (.vhd, .vhdx, or .vhds) is stored on a CSV or SMB share and when a storage communication failure is detected, the virtual machine goes into a critical pause state. You can think of this as a critical "saved" state. This causes the entire compute and memory state to freeze and stops all I/O communication. When the storage communication begins to allow reads and writes, the virtual machine resumes operations and returns to a normal state.

When running a virtual machine with a shared VHD and guest cluster, the new storage resilience feature is designed to remove the attached disk if one virtual machine cannot see the shared disk. This allows the failover cluster to detect the storage failure and take action according to the failover cluster configuration.

- Critical pause state
- Shared VHD considerations
- Configuration options

A virtual machine stays in critical paused state for a default amount of time before going into an off state. This can be changed using the following Windows PowerShell command:

```
Set-VM -AutomaticCriticalErrorActionTimeout <value in minutes>
```

The new storage resiliency is enabled by default. To disable it, run the following command:

```
Set-VM -AutomaticCriticalErrorAction <None | Pause>
```

Converting and resizing disks

Periodically, it is necessary to perform maintenance operations on virtual hard disks. You can perform the following maintenance operations on virtual hard disks:

- Convert a disk from fixed-size to dynamically expanding.
- Convert a disk from dynamically expanding to fixed-size.
- Convert a virtual hard disk in .vhd format to .vhdx format.
- Convert a virtual hard disk in .vhdx format to .vhd format.

- You can perform the following maintenance operations on virtual hard disks:
 - Convert a disk from fixed to dynamic
 - Convert a disk from dynamic to fixed
 - Convert a virtual hard disk in .vhd format to .vhdx format
 - Convert a virtual hard disk in .vhdx format to .vhd format
- With Windows Server 2012 R2 and Windows Server 2016 Hyper-V, you can resize virtual hard disks while the virtual machine is running
- With Windows Server 2016, you can resize shared virtual hard disks while the virtual machine is running

You can shrink a dynamically expanding virtual hard disk that is not taking up all the space that is allocated to it. For example, a dynamically expanding virtual hard disk might be 60 GB on the host volume, but it only uses 20 GB of that space. You shrink a virtual hard disk by choosing the **Compact** option in the **Edit Virtual Hard Disk Wizard**.

You cannot shrink fixed-size virtual hard disks. You must convert a fixed-size virtual hard disk to dynamically expanding before you can compact the disk. You can use the **resize-partition** and **resize-vhd** Windows PowerShell cmdlets to compact a dynamically expanding virtual hard disk.

You also can use the **Edit Virtual Hard Disk Wizard** to expand a disk. You can expand dynamically expanding and fixed-size virtual hard disks.

Because Azure currently only supports importing virtual hard disks that are in .vhd format, when you want to import a current virtual hard disk to Azure, you need to be able to convert virtual hard disks. This requirement means that you need to convert these disks before importing them to Azure. When you import a dynamically expanding hard disk to Azure, the virtual hard disk will automatically convert to a fixed-size hard disk that is 127 GB in size.

Resizing virtual hard disks

In Windows Server 2016, you can resize virtual hard disks that are used by a running virtual machine, only if the following prerequisites are met:

- You can only resize a virtual hard disk if the virtual hard disk is in .vhdx format and it is connected to a virtual SCSI controller.
- You can only resize a shared virtual hard disk if the virtual machine is running Windows Server 2016.
- You cannot resize virtual hard disks that are connected to a virtual IDE controller.

- You cannot shrink a virtual hard disk beyond the size of the current volumes that are hosted on the virtual hard disk. Before attempting to shrink a virtual hard disk, use Disk Manager in the guest virtual machine operating system to reduce the size of the volumes that are hosted on the virtual hard disk.



Note: The prerequisites listed above are not required when resizing a virtual hard disk that a running virtual machine is not using.

Fibre Channel support in Hyper-V

Hyper-V virtual Fibre Channel is a virtual hardware component that you can add to a virtual machine and that enables the virtual machine 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 adapter (HBA).
- The Fibre Channel HBA must have a driver that supports virtual Fibre Channel.
- The virtual machine must support virtual machine extensions.

The Fibre Channel adapter:

- Allows a virtual machine to connect to a Fibre Channel SAN directly
- Requires that the Hyper-V host has a Fibre Channel HBA
- Requires that the Fibre Channel HBA driver supports virtual Fibre Channel

Virtual Fibre Channel adapters support port virtualization by exposing HBA ports in the guest operating system. Doing so allows the virtual machine to access the SAN by using a standard World Wide Name that is associated with the virtual machine.

You can deploy up to four virtual Fibre Channel adapters on each virtual machine.



Additional Reading: For more information, refer to: "Hyper-V Virtual Fibre Channel Overview" at: <http://aka.ms/gpv90h>

Location considerations for virtual hard disks

A key consideration when provisioning virtual machines is to ensure correct placement of virtual hard disks. Virtual hard disk performance can affect virtual machine performance dramatically. Servers that are otherwise well-provisioned with RAM and processor capacity can still experience poor performance if the storage system is overwhelmed.

Consider the following factors when you plan the location of virtual hard disk files:

- High-performance connection to storage. You can locate virtual hard disk files on local or remote storage. When you locate them on remote storage, you need to ensure that there is adequate bandwidth and minimal latency between the host and the remote storage. Both slow network

When planning the location of virtual hard disks, ensure that the virtual hard disk files are:

- Stored on disks that can be accessed quickly from the Hyper-V host
- Stored on a volume that is configured for redundancy
- Stored on high-performance storage
- Placed on volumes with adequate space if they are configured for growth

connections to storage, and 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. Even though hard disk failures are common, the virtual machine and the Hyper-V host need to remain in operation after a disk failure. Replacement of failed disks also must not affect the operation of the Hyper-V host or virtual machines.
- **High-performance storage.** The storage device on which you store virtual hard disk files should have excellent I/O characteristics. Many enterprises use hybrid SSD drives in RAID 1+0 arrays to achieve maximum performance and redundancy. Multiple virtual machines that are running simultaneously on the same storage can place a tremendous I/O burden on a disk subsystem. Therefore, you need to ensure that you choose high-performance storage. If you do not, virtual machine performance suffers.
- **Adequate growth space.** If you have configured virtual hard disks to grow automatically, ensure that the files have adequate space to grow into. In addition, carefully monitor growth so that you are not surprised when a virtual hard disk fills the volume that you allocated to host it.

Storing virtual machines on SMB 3.0 file shares

Hyper-V supports storing virtual machine data such as virtual machine configuration files, checkpoints, and virtual hard disk files on SMB 3.0 file shares. The file share must support SMB 3.0. This limits placement of virtual hard disks to file shares that are hosted on Windows Server 2012 or newer file servers. Older versions of Windows Server do not support SMB 3.0.



Note: The recommended bandwidth for network connectivity to the file share is 1 Gigabit per second (Gbps) or more.

- SMB 3.0 is available in Windows Server 2012 but not in earlier Windows Server versions
- Hyper-V can store the following on SMB 3.0 file shares:
 - Configuration files
 - Virtual hard disk files, in .vhd or .vhdx format
 - Checkpoint files


SMB 3.0 file shares provide an alternative to storing virtual machine files on iSCSI or Fibre Channel SAN devices. When creating a virtual machine in Hyper-V on Windows Server 2012 or newer, you can specify a network share when choosing the virtual machine location and the virtual hard disk location. You also can attach disks stored on SMB 3.0 file shares. You can use .vhd, .vhdx, and .vhds files with SMB 3.0 file shares.

Windows Server 2016 now uses Storage QoS to manage QoS policies for Hyper-V and Scale-Out File Servers. This allows deployment of QoS policies for SMB 3.0 storage.



Additional Reading: For more information, refer to: "Server Message Block Overview" at: <http://aka.ms/obyww0>

Since the release of Windows Server 2012 R2, SMB 3.0 has improved to allow shared storage for guest clustering to be stored on an SMB 3.0 file server.


 **Note:** The module, “Implementing Failover Clustering with Windows Server 2016 Hyper-V,” covers in more detail the concepts of using SMB 3.0 to host virtual machine files for Hyper-V Live Migration.

Demonstration: Managing virtual hard disks in Hyper-V

In this demonstration, you will see how to create a differencing disk based on an existing disk by using both Hyper-V Manager and Windows PowerShell.

Demonstration Steps

1. Use File Explorer to create the following folders on the physical host drive:
 - o **E:\Program Files\Microsoft Learning\20743\Drives\LON-GUEST1**
 - o **E:\Program Files\Microsoft Learning\20743\Drives\LON-GUEST2**

 **Note:** The drive letter might depend on the number of drives on the physical host machine.

2. On **LON-HOST1**, run Windows PowerShell Integrated Scripting Environment (ISE) as an administrator, go to **E:\Program Files\Microsoft Learning\20743\Drives**, and then run the **LON-HOST1_VM-Pre-Import-20743A.ps1** file.
3. In Hyper-V Manager, create a virtual hard disk with the following properties:
 - o Disk Format: **VHD**
 - o Disk Type: **Differencing**
 - o Name: **LON-GUEST1.vhd**
 - o Location: **E:\Program Files\Microsoft Learning\20743\Drives\LON-GUEST1**
 - o Parent Location: **E:\Program Files\Microsoft Learning\20743\Drives\20743A-BASE.vhd**
4. Open **Windows PowerShell**, import the Hyper-V module, and then run the following command:

```
New-VHD "E:\Program Files\Microsoft Learning\20743\Drives\LON-GUEST2\LON-GUEST2.vhd"  
-ParentPath "E:\Program Files\Microsoft Learning\20743\Drives\20743A-BASE.vhd"
```
5. Inspect the **E:\Program Files\Microsoft Learning\20743\Drives\LON-GUEST2\LON-GUEST2.vhd** disk.
6. Verify that **LON-GUEST2.vhd** is configured as a differencing virtual hard disk with **E:\Program Files\Microsoft Learning\20743\Drives\20743A-BASE.vhd** as a parent.

Question: You want to create a guest cluster for which you can manage and resize the shared storage through Hyper-V while the servers are online. What requirement must you consider?

Check Your Knowledge

Question	
When you configure a Virtual Hard Disk, which of the following options are available? Select all that apply.	
Select the correct answer.	
<input type="checkbox"/>	Pass-through
<input type="checkbox"/>	Dynamic
<input type="checkbox"/>	Differencing
<input type="checkbox"/>	Fixed

Lesson 3

Configuring Hyper-V networking

Hyper-V provides several different options for allowing network communication between virtual machines. You can use Hyper-V to configure virtual machines that communicate with an external network in a manner similar to physical hosts that you deploy traditionally. You also can use Hyper-V to configure virtual machines that are able to communicate only with a limited number of other virtual machines that are hosted on the same Hyper-V host. This lesson describes the various options that are available for Hyper-V virtual networks, which you can use to best meet your organization's needs.

Lesson Objectives

After completing this lesson, you will be able to:


- Describe the new Hyper-V networking features in Windows Server 2012.
- Describe the new Hyper-V networking features in Windows Server 2012 R2.
- Describe the new Hyper-V networking features in Windows Server 2016.
- Describe types of Hyper-V networks.
- Create Hyper-V network types.
- Describe Hyper-V virtual networking.
- Describe best practices for configuring virtual networks.

New Hyper-V networking features in Windows Server 2012

Several new features in Windows Server 2012 Hyper-V networking improve network performance and the flexibility of virtual machines in private and public cloud environments.

The new features that were introduced in Windows Server 2012 Hyper-V networking include:

- Network virtualization. This feature allows IP addresses to be virtualized in hosting environments so that virtual machines that migrate to the host can keep their original IP addresses, rather than being allocated IP addresses on the Hyper-V server's network.
- Bandwidth management. You can use this feature to specify a minimum and maximum bandwidth that Hyper-V allocates to the adapter. Hyper-V reserves the minimum bandwidth allocation for the network adapter even when other virtual network adapters on virtual machines that are hosted on the Hyper-V host are functioning at capacity.
- Dynamic Host Configuration Protocol (DHCP) guard. This feature drops DHCP messages from virtual machines that are functioning as unauthorized DHCP servers. This might be necessary in scenarios where you are managing a Hyper-V server that hosts virtual machines for others but in which you do not have direct control over the virtual machines' configuration.



- Router guard. This feature drops router advertisement and redirection messages from virtual machines that are configured as unauthorized routers. This might be necessary in scenarios where you do not have direct control over the configuration of virtual machines.
- Port mirroring. You can use this feature to copy incoming and outgoing packets from a network adapter to another virtual machine that you have configured for monitoring.
- NIC Teaming. You can use this feature to add a virtual network adapter to an existing team on the host Hyper-V server.
- Virtual Machine Queue (VMQ). This feature requires the host computer to have a network adapter that supports the feature. VMQ uses hardware packet filtering to deliver network traffic directly to a guest. This improves performance because the packet does not need to be copied from the host operating system to the virtual machine. Only Hyper-V–specific network adapters support this feature.
- Single-root I/O virtualization (SR-IOV). This feature requires that specific hardware and special drivers be installed on the guest operating system. SR-IOV enables multiple virtual machines to share the same Peripheral Component Interconnect Express physical hardware resources. If sufficient resources are not available, network connectivity fails to the virtual switch. Only Hyper-V–specific network adapters support this feature.
- IP security (IPsec) task offloading. This feature requires that the guest operating system and network adapter are supported. This feature allows a host's network adapter to perform calculation-intensive security-association tasks. If sufficient hardware resources are not available, the guest operating system performs these tasks. You can configure a maximum number of offloaded security associations from one to 4,096. Only Hyper-V–specific network adapters support this feature.

New Hyper-V networking features in Windows Server 2012 R2

The enhancements to virtual switches in Windows Server 2012 R2 in comparison to the virtual switches in Windows Server 2012 Hyper-V are:

- Extended port access control lists (ACLs)
- Dynamic load balancing
- Coexistence with non-Microsoft forwarding extensions
- Receive Side Scaling supported on the virtual machine network path
- Network tracing enhancements

The virtual switch improvements in Windows Server 2012 R2 include:

- Extended port access control lists
- Dynamic load balancing
- Coexistence with non-Microsoft forwarding extensions
- Receive Side Scaling supported on the virtual machine network path
- Network tracing improvements

Extended port ACLs

You can use extended port ACLs in a Hyper-V virtual switch to enforce security policies and firewall protection at the switch level for virtual machines. The difference between ACLs in Windows Server 2012 and Windows Server 2012 R2 Hyper-V include:

- Administrators can now include socket port numbers when developing ACLs.
- Hyper-V switches support unidirectional stateful rules with a timeout parameter.

Dynamic load balancing of network traffic

When you map a virtual network to a network adapter team on a Windows Server 2012 R2 Hyper-V host, network traffic will be continuously load balanced across network adapters, with traffic streams moved as necessary to maintain this balance. In Windows Server 2012 Hyper-V, traffic streams remained with the network adapter in the team that they were assigned to initially, and they would not be dynamically moved to other network adapters in the team.

Coexistence with non-Microsoft forwarding extensions


The Hyper-V Network Virtualization module forwards network traffic that is encapsulated through Network Virtualization using Generic Routing Encapsulation (NVGRE). Non-Microsoft switch extensions are supported in coexistence scenarios with Hyper-V virtual switches. When a non-Microsoft extension is present, any non-NVGRE network traffic is forwarded by using the non-Microsoft forwarding extensions.

Receive Side Scaling on the virtual machine network path

Windows Server 2012 R2 supports virtual Receive Side Scaling on the virtual machine network path. This allows virtual machines to support greater network traffic loads. Virtual Receive Side Scaling accomplishes this by spreading the processing load across multiple processor cores on both the Hyper-V host and the virtual machine. A virtual machine can take advantage of virtual Receive Side Scaling improvements only if the processor on the Hyper-V host supports Receive Side Scaling and the virtual machine is configured to use multiple processor cores.

Network tracing improvements

You use **Netsh Trace** commands to trace packets. The improvements in Windows Server 2012 R2 enable you to view port and switch information as you trace network traffic through Hyper-V virtual switches.

 **Additional Reading:** For more information, refer to: "What's New in Hyper-V Virtual Switch in Windows Server 2012 R2" at: <http://aka.ms/gd08pi>

New Hyper-V networking features in Windows Server 2016

In Windows Server 2016, Microsoft continues to improve the Software Defined Networking infrastructure. Software Defined Networks are a primary building block of software-defined datacenters. You can manage many of these features easily through Virtual Machine Manager. However, you can also use Windows PowerShell commands to configure implementations of these features. The default settings are adequate for most small-scale environments. Some of the new or improved features are:

- Network function virtualization
- Network Controller
- Switch Embedded Teaming
- Remote Direct Memory Access
- Virtual machine multi queues
- Converged NICs

- Network function virtualization. In most datacenters, hardware appliances handle some network functions such as software load balancing and network address translation, datacenter firewalls, and Remote Access Service gateways. However, with Software Defined Networking, more appliances are becoming virtualized. All three functions are available in Windows Server 2016.



Note: To use containers and build out Hyper-V virtualized networks more efficiently, it is important that you have the ability to utilize network address translation with Windows Server 2016 as a built-in feature of a virtual switch. You can create a virtual switch in a virtual machine container host by running the following command:

```
New-VMSwitch -Name "Virtual Switch Name" -SwitchType NAT
```

- **Network Controller.** By using Network Controller, you can have a central location to monitor, manage, troubleshoot, and configure both your physical and virtual environments.
- **Switch Embedded Teaming (SET).** SET is a new NIC Teaming option that you can use for Hyper-V networks. SET has some integrated functionality with Hyper-V that provides faster performance and better fault tolerance than traditional teams.
- **Remote Direct Memory Access (RDMA) with Hyper-V.** RDMA services can now use Hyper-V switches. You can enable this feature with or without SET.
- **Virtual machine multi queues (VMMQ).** This feature allocates multiple hardware queues for each virtual machine, thereby improving throughput as compared to Windows Server 2012 R2.
- **Converged network adapter.** A converged network adapter supports using a single network adapter or a team of network adapters to handle multiple forms of traffic, management, RDMA, and virtual machine traffic. This reduces the number of specialized adapters that each host needs.

Types of Hyper-V networks

Virtual switches are virtual devices that you can manage through the Virtual Switch Manager, which enables you to create three types of virtual switches—external, internal, and private. Virtual switches control how network traffic flows between virtual machines that are hosted on a Hyper-V server, in addition to how network traffic flows between virtual machines and the rest of the organizational network.

Hyper-V on Windows Server 2012 and Windows Server 2016 supports three types of virtual switches, which the following table details.




Type	Description
External	You use this type of switch to map a network to a specific network adapter or network adapter team. Windows Server 2012 supports mapping an external network to a wireless network adapter if you have installed the Wireless LAN service on the host Hyper-V server and if the Hyper-V server has a compatible network adapter.
Internal	You use internal virtual switches to communicate between the virtual machines on a Hyper-V host and to communicate between the virtual machines and the Hyper-V host itself.
Private	You use private switches only to communicate between virtual machines on a Hyper-V host. You cannot use private switches to communicate between the virtual machines and the Hyper-V host.

When configuring a virtual network, you also can configure a virtual LAN (VLAN) ID to associate with the network. You can use this configuration to extend existing VLANs on an external network to VLANs within the Hyper-V host's network switch. You can use VLANs to partition network traffic. VLANs function as separate logical networks. Traffic can pass only from one VLAN to another if it passes through a router.

You can configure the following extensions for each virtual switch type:

- Microsoft Network Driver Interface Specification (NDIS) Capture. This extension allows the capture of data that travels across a virtual switch.
- Microsoft Windows Filtering Platform. This extension allows filtering of data that travels across a virtual switch.

 **Additional Reading:** For more information, refer to: "Hyper-V Virtual Switch Overview" at: <http://aka.ms/jqu2uq>

Demonstration: Creating Hyper-V network types

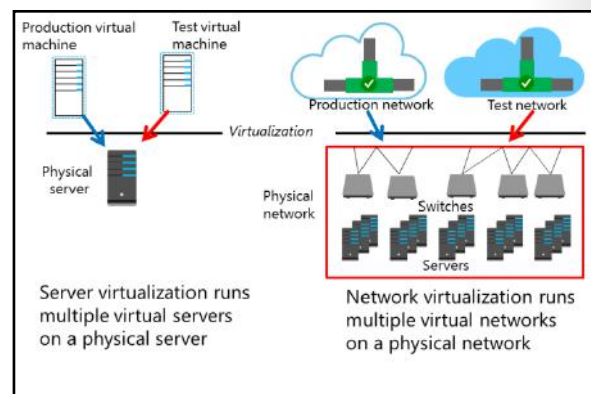
In this demonstration, you will see how to create two types of virtual switches.

Demonstration Steps

1. In Hyper-V Manager, use the Virtual Switch Manager to create a new external virtual network switch with the following properties:
 - o Name: **Corporate Network**
 - o Connection Type: **External Network:** Mapped to the host computer's physical network adapter. Varies depending on the host computer.
2. In Hyper-V Manager, use the Virtual Switch Manager to create a new virtual switch with the following properties:
 - o Name: **Private Network**
 - o Connection type: **Private network**

Hyper-V virtual networking

You can use Software Defined Networks to isolate virtual machines from different organizations even if they are connected to the same physical network. You can also use network virtualization to separate development and production virtual machines. For example, you might be providing infrastructure as a service (IaaS) to competing businesses. You can use network virtualization to go beyond assigning these virtual machines to separate VLANs, as a way of isolating network traffic. Network virtualization is a technology that you would deploy primarily in scenarios where you use Hyper-V to host virtual machines for non-Microsoft organizations. Network virtualization has the advantage that you can configure all network isolation on the Hyper-V host. With VLANs, it also is necessary to configure switches with the appropriate VLAN IDs.



When you configure network virtualization, each guest virtual machine has two IP addresses, which work as follows:

- **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 though a Hyper-V server that is connected to a separate public IP network might host the virtual server. Using the **ipconfig** command on the virtual machine shows the customer IP address.
- **Provider IP address.** The physical network 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 the virtual machine.

You can use network virtualization to host multiple virtual machines that use the same customer IP address, such as 192.168.15.101, on the same Hyper-V host. When you do this, the virtual machines are assigned different IP addresses by the hosting provider; although, this address will not be apparent from within the virtual machine.

You manage network virtualization by using Windows PowerShell cmdlets or System Center Virtual Machine Manager (SCVMM). All network virtualization cmdlets are in the NetWNV module for Windows PowerShell. 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.



Additional Reading: For more information, refer to: "Hyper-V Network Virtualization Overview" at: <http://aka.ms/vfku5o>

Best practices for configuring virtual networks

Best practices for configuring virtual networks typically focus on ensuring that you provision virtual machines with adequate bandwidth. You do not want the performance of all virtual machines to be affected if a bandwidth-intensive operation, such as a large file copy or website traffic spike, occurs on one virtual machine on the same host.

The following general best practices apply to configuring virtual networks:

- **Considerations for NIC Teaming.** You should deploy multiple network adapters to a Hyper-V host and then configure those adapters as part of a team. Doing so ensures that network connectivity is retained if individual network adapters fail. Configure multiple teams with network adapters that connect to different switches to ensure that connectivity remains if a hardware switch fails.
- **Considerations for 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 that is 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 a VMQ
- Use network virtualization when you have to ensure that virtual machines keep their original IP addresses after migrating to a new host

- Considerations for VMQ. You should provision a Hyper-V host with an adapter that supports VMQ. VMQ uses hardware packet filtering to deliver network traffic directly to a virtual machine. Doing so improves performance because you do 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.
- Considerations for network virtualization. Network virtualization is complicated to configure, but it has an advantage over VLAN—it is not necessary to configure VLANs on all the switches that are connected to the Hyper-V host. You can perform all necessary configurations when you need to isolate servers on a Hyper-V host without needing to involve the network team. If you are hosting large numbers of virtual machines and need to isolate them, use network virtualization rather than VLANs.

Question: When configuring a Hyper-V host with multiple network adapters on the same network, what should you do to provide redundancy and performance?

Lesson 4

Configuring Hyper-V virtual machines

When planning a server virtualization strategy, you need to know what you can and cannot accomplish when you are using Windows Server 2016 as a virtual machine host. In this lesson, you will learn how to configure virtual machine settings to represent different hardware configurations.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe virtual machine settings.
- Describe virtual machine configuration versions.
- Describe Generation 2 virtual machines.
- Create a virtual machine.
- Explain how memory works in Hyper-V.
- Describe checkpoints and explain how to manage checkpoints in Hyper-V.
- Create checkpoints.
- Explain how to import, export, and move virtual machines in Hyper-V.
- Describe the best practices for configuring virtual machines.
- Use Windows PowerShell Direct to restart a virtual machine.

Virtual machine generation versions

Windows Server 2012 R2 introduced a new type of virtual machine called a *Generation 2 virtual machine*. With this new name, all virtual machines that were created on platforms such as Windows Server 2012 and Windows Server 2008 R2 Hyper-V are termed *Generation 1 virtual machines*. Generation 2 virtual machines use a different hardware model and do not support many of the older devices that Generation 1 virtual machines supported, such as COM ports and the emulated floppy disk drive.

Generation 2 virtual machines provide the following functionality:

- Secure boot
- Boot from a virtual hard disk that is connected to a virtual SCSI controller
- Boot from a virtual DVD that is connected to a virtual SCSI controller
- PXE boot by using a standard Hyper-V network adapter
- UEFI firmware support

You determine the generation of a virtual machine during virtual machine creation. After a virtual machine is created, you cannot migrate it from Generation 1 to Generation 2, or from Generation 2 to Generation 1.

With Windows Server 2016, it is a best practice to use Generation 2 virtual machines if the guest is a supported operating system.

Generation 2 virtual machines support the following functionality:

- Secure boot
- Boot from a virtual hard disk connected to a virtual SCSI controller
- Boot from a virtual DVD connected to a virtual SCSI controller

- PXE boot by using a standard Hyper-V (not legacy) network adapter
- Unified Extensible Firmware Interface (UEFI) firmware support

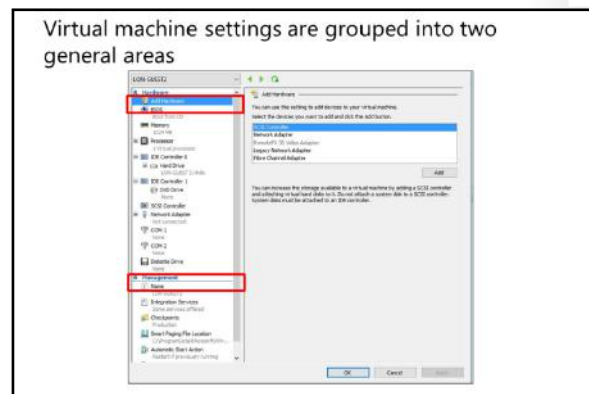
You can use Generation 2 virtual machines only with the following guest operating systems:

- Windows Server 2012 and later
- x64 editions of Windows 8 and later
- CentOS and Red Hat Enterprise Linux 7.1x and later
- Ubuntu 14.04 and later
- Oracle Linux 7.x series
- SUSE Enterprise Server 12

Windows Server 2016 Hyper-V supports secure boot for Ubuntu, CentOS, Red Hat Enterprise Linux, and SUSE Enterprise Server operating systems that run on Generation 2 virtual machines.

Overview of virtual machine settings

The two groups of virtual machine settings are hardware and management. In Windows Server 2016, the configuration files that store the hardware and management information are separated into two formats, .vmcx and .vmrs. The .vmcx format is used for configuring virtual machines, and the .vmrs format is used for runtime data. This helps decrease the chance of data corruption during a storage failure.



Hardware

Virtual machines use simulated hardware. Hyper-V uses this virtual hardware to mediate access to actual hardware. Depending on the scenario, you might not need to use all available simulated hardware. For example, you can map a virtual network adapter to a virtual network that, in turn, maps to an actual network interface, without having to use a COM port or a disk drive.

Generation 1 virtual machines have the following hardware by default:

- Firmware. Virtual hardware simulates a computer's BIOS. You can configure the virtual machine to switch Num Lock on or off. You also can choose the startup order for a virtual machine's virtual hardware. You can start a virtual machine from a DVD drive, an IDE device, a legacy network adapter, or a floppy disk.
- Memory. You can allocate memory resources to a virtual machine. An individual virtual machine can allocate as much as 1 TB of memory.
- Processor. You can allocate processor resources to a virtual machine. You can allocate up to 64 virtual processors to a single virtual machine.

- IDE controller. A virtual machine can support two IDE controllers only and, by default, allocates two IDE controllers to a virtual machine. These are IDE controller 0 and IDE controller 1. Each IDE controller can support two devices. You can connect virtual disks or virtual DVD drives to an IDE controller. If starting from a hard disk drive or DVD-ROM, the boot device must be connected to an IDE controller. IDE controllers are the only way to connect virtual hard disks and DVD-ROMS to virtual machines that use operating systems that do not support integration services.
- SCSI controller. You can use SCSI controllers only on virtual machines that you deploy with operating systems that support integration services. SCSI controllers allow you to support up to 256 disks, using four controllers with a maximum of 64 connected disks each. You can add and remove virtual SCSI disks while a virtual machine is running.
- Hyper-V–specific network adapter. Hyper-V–specific network adapters represent virtualized network adapters. You can only use Hyper-V–specific network adapters with supported virtual machine guest operating systems that support integration services.
- COM port. A COM port enables connections to a simulated serial port on a virtual machine.
- Diskette drive. You can map a .vfd floppy disk image to a virtual disk drive.

Generation 2 virtual machines have the following hardware by default:

- Firmware. UEFI, this allows all the features of the BIOS in Generation 1 virtual machines. However, it allows secure boot as well.
- Memory. Same as Generation 1 virtual machine.
- Processor. Same as Generation 1 virtual machine.
- SCSI Controller. Generation 2 virtual machines can boot to a virtual disk on the SCSI Controller.
- Network Adapter. Same as Generation 1 virtual machine.

You can add the following hardware to a virtual machine by editing the virtual machine's properties and clicking **Add Hardware**:

- SCSI controller. You can add up to four virtual SCSI devices. Each controller supports up to 64 disks.
- Network adapter. A single virtual machine can have a maximum of eight Hyper-V–specific network adapters.
- Legacy network adapter. Legacy network adapters allow the use of network adapters with operating systems that do not support integration services. You also can use legacy network adapters to allow network deployment of operating system images. A single virtual machine can have up to four legacy network adapters. Generation 2 virtual machines do not support this feature.
- Virtual Fibre Channel adapter. This adapter allows a virtual machine to connect directly to a Fibre Channel SAN. This adapter needs the Hyper-V host to have a Fibre Channel HBA that also has a Windows Server 2012 driver that supports Virtual Fibre Channel.
- Microsoft RemoteFX 3D video adapter. The RemoteFX 3D video adapter allows Generation 1 virtual machines to take advantage of DirectX and graphics processing power on the host server that is running Windows Server 2012 and later versions to display high-performance graphics. RemoteFX 3D video adapters were unavailable on Generation 2 virtual machines before Windows Server 2016. But Windows Server 2016 host servers can now add RemoteFX 3D to the virtual machine.

Management

Use management settings to configure how a virtual machine behaves on a Hyper-V host. The following virtual machine management settings are configurable:

- **Name.** Use this setting to configure a virtual machine's name on a Hyper-V host. Doing this does not alter the virtual machine's host name.
- **Integration services.** Use this setting to configure which virtual machine integration settings are enabled.
- **Checkpoint file location.** Use this setting to specify a location for storing virtual machine checkpoints.
- **Smart Paging file location.** This is the location that is used when Smart Paging is required to start a virtual machine.
- **Automatic start action.** Use this setting to handle how a virtual machine responds when a Hyper-V host is powered on.
- **Automatic stop action.** Use this setting to handle how a virtual machine responds when a Hyper-V host shuts down gracefully.

What are virtual machine configuration versions?

Virtual machine configuration versions represent a virtual machine's Hyper-V compatibility settings for its configuration, saved states, and checkpoint files. In previous versions of Hyper-V, when you upgraded your host to a new operating system, the virtual machine would upgrade to the same configuration version as the host as soon as you moved the virtual machine.

With Windows Server 2016, a virtual machine's configuration version does not upgrade automatically. Instead, upgrading is now a manual process. With rolling upgrades, it is highly likely that you might have a Hyper-V failover cluster that will have both Windows Server 2012 R2 and Windows Server 2016. Windows Server 2012 R2 version 5.0 will run on both Windows Server 2012 R2 and Windows Server 2016 hosts. This allows administrators to leave virtual machines unchanged until the upgrade of all failover cluster nodes is completed.

After all the hosts are upgraded or when you feel that you will not need to move your virtual machines to legacy hosts, you can shut down the virtual machine and upgrade the configuration version when needed.

- The Windows Server 2012 R2 virtual machine configuration is compatible on Windows Server 2016 hosts
- Check the version
- Update the version

Check the virtual machine configuration version

To check a virtual machine's configuration version, run the following command in an elevated Windows PowerShell command prompt:

```
Get-VM * | Format-Table Name, Version
```

Update a single virtual machine

To update the version of a single virtual machine, run the following command from an elevated Windows PowerShell command prompt:

```
Update-VMVersion "vmname"
```

Update all virtual machines on all cluster nodes

To update all virtual machines' versions on all cluster nodes, run the following command from an elevated Windows PowerShell command prompt:

```
Get-VM -ComputerName (Get-ClusterNode) | Stop-VM
Get-VM -ComputerName (Get-ClusterNode) | Update-Version -confirm $false
Get-VM -ComputerName (Get-ClusterNode) | Start-VM
```



Note: New Hyper-V features in Windows Server 2016 are not available until the virtual machine configuration version has been upgraded to the Windows Server 2016 version. This includes hot add/remove of memory, production checkpoints, and resizing of live shared drives.

Demonstration: Creating a virtual machine

In this demonstration, you will see how to create a virtual machine with the traditional method of using Hyper-V Manager. You also will see how you can automate the process by using Windows PowerShell.

Demonstration Steps

1. Use Hyper-V Manager to create a virtual machine with the following properties:
 - o Name: **LON-GUEST1**
 - o Location: **E:\Program Files\Microsoft Learning\20743\Drives\LON-GUEST1**
 - o Generation: **Generation 1**
 - o Memory: **1024 MB**
 - o Use Dynamic Memory: **Yes**
 - o Networking: **Private Network**
 - o Connect Virtual Hard Disk: **E:\Program Files\Microsoft Learning\20743\Drives\LON-GUEST1\lon-guest1.vhd**
2. Open **Windows PowerShell**, import the Hyper-V module, and then run the following command:

```
New-VM -Name LON-GUEST2 -MemoryStartupBytes 1024MB -VHDPATH "E:\Program
Files\Microsoft Learning\20743\Drives\LON-GUEST2\LON-GUEST2.vhd" -SwitchName "Private
Network"
```

3. Use the Hyper-V Manager console to configure **LON-GUEST2** with the following settings:
 - o Automatic Start Action: **Nothing**
 - o Automatic Stop Action: **Shut down the guest operating system**

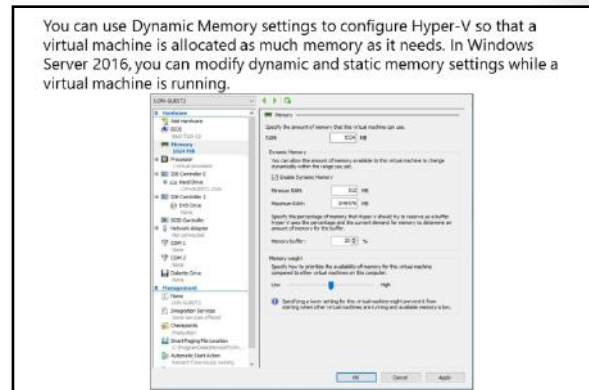
How memory works in Hyper-V

In the first release of Hyper-V with Windows Server 2008, you could assign virtual machines only a static amount of memory. Unless you took special precautions to measure the precise amount of memory that a virtual machine required, you were likely to under-allocate or over-allocate memory, just as would occur on a physical computer.

Windows Server 2008 R2 Service Pack 1 (SP1) introduced dynamic memory, which you can use to allocate the startup amount of memory for a virtual machine. You then can allow the virtual machine to request additional memory as necessary, up to a specified limit. Rather than attempting to guess how much memory a virtual machine requires, dynamic memory allows you to configure Hyper-V so that a virtual machine is allocated as much memory as it needs. You can choose a minimum value, which is always allocated to the virtual machine. You can choose a maximum value, which a virtual machine does not exceed even if more memory is requested. Virtual machines must support Hyper-V integration services through the deployment or inclusion of integration services components to be able to use dynamic memory.

With Windows Server 2012, you can modify certain dynamic memory settings while a virtual machine is running. This was not possible in Windows Server 2008 R2 SP1.

Windows Server 2016 will now let you modify static memory while a virtual machine is online. In previous versions of Hyper-V, if the virtual machine was set to static memory, downtime was required to add or remove memory.



Smart Paging

Windows Server 2016 supports Smart Paging, which provides a solution to the problem of minimum memory allocation as it relates to virtual machine startup. Virtual machines can require more memory during startup than they would require during normal operations. Before Windows Server 2012, it was necessary to allocate the minimum required memory for startup to ensure that startup occurred, even though that value could be more than what the virtual machine needed during normal operation.

Smart Paging uses disk paging for additional temporary memory when additional memory beyond the minimum allocation is required to restart a virtual machine. Smart Paging allows you to allocate a minimum amount of memory based on the amount needed when the virtual machine is operating normally, rather than the amount that is required during startup. One drawback of Smart Paging is a decrease in performance during virtual machine restarts.

You can configure virtual machine memory by using the **Set-VMemory** Windows PowerShell cmdlet.



Additional Reading: For more information, refer to: "Hyper-V Dynamic Memory Overview" at: <http://aka.ms/rb0nbx>

Checkpoints and checkpoint management in Hyper-V

Checkpoints allow administrators the ability to make a snapshot of a virtual machine at a specific time. Windows Server 2016 also provides production checkpoints and standard checkpoints, with the default being production checkpoints. It is important to know when to use a standard checkpoint and when to use a production checkpoint.



Note: Ensure that you only use checkpoints with server applications that support the use of checkpoints. If you revert to a previous checkpoint on a computer that hosts an application that does not support virtual machine checkpoints, it might lead to data corruption or loss. Some applications might only support production checkpoints.

- Checkpoints allow administrators to make a snapshot of a virtual machine at a particular point in time
- Checkpoints do not replace backups
- Standard checkpoints create differencing disks, .avhd files, which merge back into the previous checkpoint when the checkpoint is deleted
- Use VSS to create production checkpoints, and start from an offline state to restore them
- Checkpoints were termed *snapshots* in previous versions of Hyper-V

Creating a checkpoint

You can create a checkpoint in the **Actions** pane of the **Virtual Machine Connection** window or in the Hyper-V Manager console. Each virtual machine can have a maximum of 50 checkpoints.

When creating checkpoints for multiple virtual machines that have dependencies, you should create them at the same time. This ensures synchronization of items such as computer account passwords. Remember that when you revert to a checkpoint, you are reverting to a computer's state at that specific time. If you take a computer back to a point before it performed a computer password change with a domain controller, you will need to rejoin that computer to the domain.

Checkpoints do not replace backups

Checkpoints are not a replacement for backups. Checkpoint data is stored on the same volume as the virtual hard disks. If the volume that hosts these files fails, both the checkpoint and the virtual hard disk files are lost. You can perform a virtual machine export of a checkpoint. When you export the checkpoint, Hyper-V creates full virtual hard disks that represent the state of the virtual machine when you created the checkpoint. If you choose to export an entire virtual machine, all checkpoints that are associated with the virtual machine also are exported.

Standard checkpoints

When you create a standard checkpoint, Hyper-V creates an .avhd file that stores the data that differentiates the checkpoint from either the previous checkpoint or the parent virtual hard disk. When you delete standard checkpoints, this data merges into the previous checkpoint or parent virtual hard disk. If you delete the second-to-last checkpoint of a virtual machine, the content of the differencing virtual hard disk merges with its parent, so that the earlier and latter checkpoint states of the virtual machine retain their integrity.

Production checkpoints

When you create a production checkpoint, Windows Server 2016 uses Volume Shadow Copy Service (VSS) or File System Freeze for Linux. This places a virtual machine in a safe state to create a checkpoint that can be recovered the same as any VSS or application backup. Unlike standard checkpoints that will save all memory and processing in the checkpoint, production checkpoints are closer to a state backup. Production checkpoints require a virtual machine to start from an offline state to restore the checkpoint.

Managing checkpoints

When you apply a checkpoint, the virtual machine reverts to the configuration that existed at the time that the checkpoint was created. Reverting to a checkpoint does not delete any existing checkpoints. If you revert to a checkpoint after making a configuration change, Hyper-V Manager prompts you to create a checkpoint. It only is necessary to create a new checkpoint if you want to return to that current configuration.

Creating checkpoint trees with different branches is possible. For example, if you create a checkpoint of a virtual machine on Monday, Tuesday, and Wednesday, apply the Tuesday checkpoint, and then make changes to the virtual machine's configuration, you create a new branch that diverts from the original Tuesday checkpoint. You can have multiple branches if you do not exceed the 50-checkpoint limit per virtual machine.

Demonstration: Creating checkpoints

In this demonstration, you will see how to create a production checkpoint and a standard checkpoint by using the traditional Hyper-V Manager method.

Demonstration Steps

1. Use Hyper-V Manager to create a production checkpoint on **LON-GUEST1**.
2. Use Hyper-V Manager to change the settings and create a standard checkpoint for **LON-GUEST1**.
3. Delete checkpoints, and then merge them into the current virtual hard disk.

Importing, exporting, and moving virtual machines in Hyper-V

You can use Hyper-V import and export functionalities to transfer virtual machines between Hyper-V hosts and to create point-in-time backups of virtual machines.

Importing virtual machines

The virtual machine import functionality in Windows Server 2016 can identify configuration problems such as missing hard disks or virtual switches. This was more difficult to determine in older operating systems before Windows Server 2012.

- When importing virtual machines in Hyper-V:
 - You can get access to detailed diagnostic information
 - You can import copied virtual machine files
- When exporting virtual machines, there are two options:
 - Export a checkpoint for point-in-time export
 - Export virtual machine to export all checkpoints
- When moving virtual machines:
 - You can relocate virtual machine files while a virtual machine is online
 - You can perform a live migration

In Windows Server 2016, you can import virtual machines from copies of virtual machine configurations, checkpoints, and virtual hard disk files rather than specially exported virtual machines. Doing this is beneficial in recovery situations where an operating system volume might have failed but the virtual machine files remain intact.

To import a virtual machine by using Hyper-V Manager, perform the following steps:

1. In the **Actions** pane of the Hyper-V Manager console, click **Import Virtual Machine**.
2. On the **Before You Begin** page of the **Import Virtual Machine** wizard, click **Next**.
3. On the **Locate Folder** page, specify the folder that hosts the virtual machine files, and then click **Next**.

4. On the **Select Virtual Machine** page, select the virtual machine that you want to import, and then click **Next**.
5. On the **Choose Import Type** page, choose from the following options, and then click **OK**:
 - **Register the virtual machine in-place (use the existing unique ID)**
 - **Restore the virtual machine (use the existing unique ID)**
 - **Copy the virtual machine (create a new unique ID)**
6. If you select the **Copy the virtual machine (create a new unique ID)** or **Restore the virtual machine (use the existing unique ID)** option, you will need to go through the **Choose Destination** and **Choose Storage Folders** menus. Otherwise click **Finish** to complete the import.

You can import virtual machines by using the **Import-VM** cmdlet.

Exporting virtual machines

When performing an export, you can choose one of the following options:

- Export a checkpoint. You can do this by right-clicking the checkpoint in the Hyper-V Manager console, and then clicking **Export**. Doing this creates an exported virtual machine because it existed at the point of checkpoint creation. The exported virtual machine will have no checkpoints.
- Export virtual machine with checkpoint. You can do this by selecting the virtual machine and then clicking **Export**. This exports the virtual machine and all checkpoints that are associated with the virtual machine.

Exporting a virtual machine does not affect the existing virtual machine. However, you cannot import the virtual machine again unless you use the **Copy the Virtual Machine** option, which creates a new, unique ID.

You can export virtual machines by using the **Export-VM** cmdlet.

Windows Server 2016 Hyper-V supports exporting virtual machines and checkpoints while a virtual machine is running.

Moving virtual machines

You can perform two types of moves by using the Hyper-V move function: a live migration and a move of the actual virtual machine. You can move virtual machines from one Windows Server 2016 Hyper-V server to another if you have enabled live migrations. Live migration of virtual machines occurs when you move a virtual machine from one host to another while keeping the virtual machine online and available to clients. For more information on migrating virtual machines, refer to Module 12, "Implementing failover clustering with Windows Server 2016 Hyper-V."

You can use the move functionality to move some or all virtual machine files to a different location. For example, if you want to move virtual machines from one volume to an SMB 3.0 share while keeping the virtual machine hosted in the same location, you have the following options:

- Move all the virtual machine's data to a single location. This moves all configuration files, checkpoints, and virtual hard-disk files to the destination location.
- Move the virtual machine's data to different locations. This moves the virtual machine's configuration files, checkpoints, and virtual hard disks to separate locations.
- Move the virtual machine's virtual hard disks. This moves the hard disks to a separate location, while keeping the checkpoint and configuration files in the same location.

You can move virtual machines in Windows PowerShell by using the **Move-VM** cmdlet.

Best practices for configuring virtual machines

When creating new virtual machines, keep the following best practices in mind:

- Use dynamic memory. The only time you should avoid dynamic memory is if you have an application that does not support it. For example, some Microsoft Exchange 2013 roles keep requesting memory if it is available. In such cases, set static memory limits. You should monitor memory utilization and set the minimum memory to the server's minimum memory utilization. Also, set a maximum amount of memory. The default maximum is more memory than most host servers have available.
- Avoid differencing disks. Differencing disks reduce the amount of space required, but they decrease performance as multiple virtual machines access the same parent virtual hard disk file.
- Use multiple Hyper-V-specific network adapters that are connected to different external virtual switches. Configure virtual machines to use multiple virtual network adapters that are connected to host network adapters, which in turn are connected to separate physical switches. Using multiple virtual network adapters means that network connectivity is retained if a network adapter or a switch fails.
- Store virtual machine files on their own volumes if you are not using shared storage. This minimizes the chance of one virtual machine's virtual hard disk growth affecting other virtual machines on the same server.

- Use Dynamic Memory unless an application does not support it
- Avoid using differencing disks
- Configure multiple synthetic network adapters
- Store each virtual machine's files on a separate volume

Demonstration: Using Windows PowerShell Direct

In this demonstration, you will see how to restart a virtual machine by using Windows PowerShell Direct.

Demonstration Steps

1. Remove network connection of guest virtual machine.
2. Connect to **LON-GUEST1**, and then set the password for the virtual machine.
3. Use Windows PowerShell Direct to restart the **LON-GUEST1** without any remoting configuration to the virtual machine.
4. Re-enable network connection.

Question: You need to run guest virtual machines on both Windows Server 2012 R2 and Windows Server 2016 servers. What should you avoid doing until you no longer need to run these virtual machines on Windows Server 2012 R2?

Question: When building a guest virtual machine running Windows Server 2016 that will require the ability to expand the startup .vhd file while the server is running, what generation version should you use?

Question: You need to build a virtual machine guest that has static memory that you can resize while the server is online. What requirements should you keep in mind?

Lab: Implementing server virtualization with Hyper-V

Scenario

IT management at A. Datum Corporation is concerned about the low utilization of many of the physical servers that are deployed in the London datacenter. A. Datum also is exploring options for expanding into multiple branch offices and deploying servers in public and private clouds. For this purpose, the company is exploring the use of virtual machines.

You will deploy the Hyper-V server role, configure virtual machine storage and networking, and deploy the virtual machines.

Objectives

After completing this lab, you will be able to:

- Install the Hyper-V server role.
- Configure virtual networking.
- Create and configure a virtual machine.

Lab Setup

Estimated Time: **60 minutes**

Virtual machine: **20743A-LON-HOST1**

User name: **Adatum\Administrator**

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

For this lab, you need to use Windows Boot Manager to boot to a virtual hard disk.

1. Restart the classroom computer, and then in Windows Boot Manager, select **20743A-LON-HOST1**.
2. Sign in to **LON-HOST1** with the following credentials:
 - o User name: **Administrator**
 - o Password: **Pa\$\$w0rd**
 - o Domain: **Adatum**

Exercise 1: Installing the Hyper-V server role

Scenario

The first step in exploring a virtualized environment is for A. Datum Corporation to install the Hyper-V server role on a new server.

The main tasks for this exercise are as follows:

1. Install the Hyper-V server role.
2. Complete Hyper-V role installation and verify settings.

► Task 1: Install the Hyper-V server role

1. On **LON-HOST1**, open **Server Manager**, and then use the **Add Roles and Features Wizard** to add the Hyper-V role to **LON-HOST1** with the following options:
 - Do not create a virtual switch.
 - Use the Default Stores locations.
 - Select to allow the server to restart automatically if required.
2. After a few minutes, the server will automatically restart. Ensure that you restart the machine by using the **Boot** menu and then selecting **20743A-LON-HOST1**. The computer will restart several times.

► Task 2: Complete Hyper-V role installation and verify settings

1. Sign in to **LON-HOST1** as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. Start **Server Manager**.
3. When the installation of the Hyper-V tools completes, click **Close**.
4. Open **Hyper-V Manager**, and then click **LON-HOST1**.
5. Open the Hyper-V settings, and then configure or verify the following settings:
 - Keyboard: **Use on the virtual machine**
 - Virtual Hard Disks: **C:\Users\Public\Documents\Hyper-V\Virtual Hard Disks**

Results: After completing this exercise, you should have deployed the Hyper-V role to a physical server.

Exercise 2: Configuring virtual networking

Scenario

After installing the Hyper-V server role on the new server, you need to configure the virtual networks. You need to create a network that connects to the physical network and a private network that you can use only for communication between virtual machines. The private network is used when virtual machines are configured for high availability. You also need to configure a specific range of media access control (MAC) addresses for the virtual machines.

The main tasks for this exercise are as follows:

1. Configure the external network.
2. Create a private network.
3. Create an internal network.
4. Configure network settings on **LON-HOST1**.

► Task 1: Configure the external network

- On **LON-HOST1** in Hyper-V Manager, use Virtual Switch Manager to create a new external virtual network switch with the following properties:
 - Name: **Corporate Network**
 - Connection type: **External Network**: Mapped to the host computer's physical network adapter. This will vary depending on the host computer.

► **Task 2: Create a private network**

- On **LON-HOST1** in Hyper-V Manager, use the Virtual Switch Manager to create a new virtual switch with the following properties.
 - Name: **Private Network**
 - Connection type: **Private network**

► **Task 3: Create an internal network**

- On **LON-HOST1** in Hyper-V Manager, use Virtual Switch Manager to create a new virtual switch with the following properties:
 - Name: **Internal Network**
 - Connection type: **Internal network**

► **Task 4: Configure network settings on LON-HOST1**

1. Open **Server Manager**, click **Local Server**, and then configure the following network settings for the **vEthernet (Internal Network)** adapter:
 - IP Address: **172.16.0.31**
 - Subnet mask: **255.255.0.0**
 - Default gateway: **172.16.0.1**
 - Preferred DNS server: **172.16.0.10**
2. Disable the host computers physical network adapter. It should be labelled Ethernet 2.

Results: After completing this exercise, you should have configured virtual switch options on a physically deployed Windows Server 2016 server that is running the Hyper-V role.

Exercise 3: Creating and configuring a virtual machine

Scenario

Your organization has asked you to deploy two virtual machines and to import a third virtual machine. You have copied a generalized .vhd file that hosts a Windows Server 2016 Hyper-V host.

To minimize the disk space use, at the cost of performance, you are going to create two differencing files based on the generalized virtual hard disk. You will use these differencing files as the hard disk files for the new virtual machines.

You also will import a specially prepared virtual machine to run on the host.

The main tasks for this exercise are as follows:

1. Import virtual machines.
2. Configure virtual machine storage.
3. Create virtual machines.
4. Configure virtual LANs (VLANs) and network bandwidth settings.

5. Configure virtual machine static memory.
6. Configure and test virtual machine checkpoints.
7. Prepare for the next module.

► Task 1: Import virtual machines



Note: Perform the following steps only on **LON-HOST1**.

1. On **LON-HOST1**, run Windows PowerShell Integrated Scripting Environment (ISE) as an administrator, go to **E:\Program Files\Microsoft Learning\20743\Drives**, and then open the **LON-HOST1_VM-Pre-Import-20743A.ps1** file.



Note: The drive letter might depend on the number of drives on the physical host machine.

2. Run **LON-HOST1_VM-Pre-Import-20743A.ps1**. When prompted, be sure to provide the drive letter that represents the host computer's **Base** and **20743** folders.
3. Close the Windows PowerShell ISE.



Note: The drive letter might depend on the number of drives on the physical host machine.

► Task 2: Configure virtual machine storage

1. On **LON-HOST1** use File Explorer to create the following folders on the physical host drive:
 - **E:\Program Files\Microsoft Learning\20743\Drives\LON-GUEST1**
 - **E:\Program Files\Microsoft Learning\20743\Drives\LON-GUEST2**



Note: The drive letter might depend on the number of drives on the physical host machine.

2. In Hyper-V Manager, create a virtual hard disk with the following properties:
 - Disk Format: **VHD**
 - Disk Type: **Differencing**
 - Name: **LON-GUEST1.vhd**
 - Location: **E:\Program Files\Microsoft Learning\20743\Drives\LON-GUEST1**
 - Parent Location: **E:\Program Files\Microsoft Learning\20743\Drives\20743A-BASE.vhd**
3. Open **Windows PowerShell**, import the Hyper-V module, and then run the following command:

```
New-VHD "E:\Program Files\Microsoft Learning\20743\Drives\LON-GUEST2\LON-GUEST2.vhd"
-ParentPath "E:\Program Files\Microsoft Learning\20743\Drives\20743A-BASE.vhd"
```

4. Inspect disk **E:\Program Files\Microsoft Learning\20743A\Drives\LON-GUEST2\LON-GUEST2.vhd** and verify that **LON-GUEST2.vhd** is configured as a differencing virtual hard disk with **E:\Program Files\Microsoft Learning\20743\Drives\20743A-BASE.vhd** as a parent.

► Task 3: Create virtual machines

1. On **LON-HOST1** use Hyper-V Manager to create a virtual machine with the following properties:
 - Name: **LON-GUEST1**
 - Location: **E:\Program Files\Microsoft Learning\20743\Drives\LON-GUEST1**
 - Generation: **Generation 1**
 - Memory: **1024 MB**
 - Use Dynamic Memory: **Yes**
 - Networking: **Private Network**
 - Connect Virtual Hard Disk: **E:\Program Files\Microsoft Learning\20743\Drives\LON-GUEST1\lon-guest1.vhd**
2. Open **Windows PowerShell**, and then run the following command:

```
New-VM -Name LON-GUEST2 -MemoryStartupBytes 1024MB -VHDPATH "E:\Program Files\Microsoft Learning\20743\Drives\LON-GUEST2\LON-GUEST2.vhd" -SwitchName "Private Network"
```

3. Use the **Hyper-V Manager** console to configure the **LON-GUEST2** with the following settings:
 - Automatic Start Action: **Nothing**
 - Automatic Stop Action: **Shut down the guest operating system**

► Task 4: Configure virtual LANs (VLANs) and network bandwidth settings



Note: Perform the following steps only on **LON-HOST1**.

1. On **LON-HOST1**, in Hyper-V Manager, start **20743A-LON-DC1-B**.
2. In Hyper-V Manager, in the **Actions** pane, click **Virtual Switch Manager**.
3. Click **Internal Network**, and then select the **Enable virtual LAN identification for management operating system** check box.
4. In the **VLAN ID** text box, type **4**, and then click **OK**.
5. Click **LON-GUEST2**, and then click **Settings**.
6. In settings for **LON-GUEST2**, click **Network Adapter**, change the **Virtual switch** to **Internal Network**, and then click **Enable virtual LAN identification**.
7. In the **VLAN identifier** text box, type **4**, and then click **OK**.
8. In Hyper-V Manager, right-click **LON-GUEST2**, and then click **Start**.



Note: You should see the machine start up and enter a **running** state.

9. Right-click **LON-GUEST2**, and then click **Connect**.
10. Click **Next**, on the next page, click **Do this later**, and then click **Accept**.
11. Set the password as **Pa\$\$w0rd**, and then click **Finish**.
12. Sign in to **LON-GUEST2** as **Administrator** with the password **Pa\$\$w0rd**.

13. Open a command prompt, and then ping **172.16.0.10** by running the following command:

```
Ping 172.16.0.10
```



Note: You should not be able to reach the machine. You will receive a **General failure** message.

14. Repeat steps 5–7 on **20743A-LON-DC1-B**.
15. Switch to **LON-GUEST2**.
16. At the command prompt, run the following command:

```
Ipconfig /renew
```

17. At the command prompt, ping **172.16.0.10** again from **LON-GUEST2** by running the following command:

```
Ping 172.16.0.10
```



Note: You should now see the message, **Reply from 172.16.0.10**.

► **Task 5: Configure virtual machine static memory**

- Using Hyper-V Manager for **LON-GUEST2**, change the static memory to **2048 MB/2GB** and confirm its expansion.

► **Task 6: Configure and test virtual machine checkpoints**

1. On **LON-GUEST2**, right-click the desktop, click **New**, and then click **Folder**. Name the folder **Sydney**.
2. Repeat step 1, and create a second folder named **Melbourne**.
3. Repeat step 1, and create a third folder named **Brisbane**.
4. On the **Action** menu of the **Virtual Machine Connection** window, click **Checkpoint**.
5. In the text box, type **Before Change**, click **Yes**, and then click **OK**.
6. Drag the **Sydney** and **Brisbane** folders to **Recycle Bin**.
7. Right-click **Recycle Bin**, and then click **Empty Recycle Bin**.
8. In the **Delete Multiple Items** dialog box, click **Yes**.
9. On the **Action** menu of the **Virtual Machine Connection** window, click **Revert**.
10. In the **Revert Virtual Machine** dialog box, click **Revert**.
11. Start and connect to **LON-GUEST2**, and then sign in with the user name as **Administrator** and the password as **Pa\$\$w0rd**.

12. Verify that the following folders are present on the desktop:

- **Sydney**
- **Melbourne**
- **Brisbane**

Results: After completing this exercise, you should have deployed two separate virtual machines on a virtual hard disk file to act as a parent disk for two differencing disks. The System Preparation Tool (Sysprep) has generalized these virtual machines. You also should have imported a specially prepared virtual machine.

► **Task 7: Prepare for the next module**

- When you finish the lab, shut down **LON-GUEST2**. If you are continuing to the next module, keep **LON-DC1** running. Otherwise you can also shut down **LON-DC1**.

Question: What type of virtual network switch would you create if you want to allow a virtual machine to communicate with the LAN that is connected to the Hyper-V host?

Question: How can you ensure that a single virtual machine does not use all the available bandwidth that the Hyper-V host provides?

Question: What dynamic memory configuration task can you perform on a virtual machine that is hosted on Windows Server 2012 Hyper-V or later?

Module Review and Takeaways

Review Questions

Question: In which situations should you use static memory allocation rather than dynamic memory?

Question: In which situations must you use virtual hard disks in .vhdx format rather than virtual hard disks in .vhd format?

Question: You want to deploy a Hyper-V virtual machine's virtual hard disk on a file share. What operating system must the file server be running to support this configuration?

Real-world Issues and Scenarios

You need to ensure that you provision a virtual machine host with adequate RAM. Having multiple virtual machines paging a hard disk drive because they are provisioned with inadequate memory will decrease performance for all virtual machines on the Hyper-V host.

Additionally, you should monitor virtual machine performance carefully. One virtual machine that uses a disproportionate amount of server resources can adversely affect the performance of all other virtual machines that the Hyper-V server hosts.

Tools

The following table includes the tools that are needed for this module:

Tool	Used for	Where to find it
Sysinternals Disk2vhd	Converts physical hard disks to .vhd format	For more information, refer to: "Sysinternals Suite" at: http://aka.ms/kx5ojf
Microsoft System Center 2012 R2 Virtual Machine Manager	<ul style="list-style-type: none"> Manages virtual machines across multiple Hyper-V servers Performs online physical-to-virtual conversions, but does not support physical-to-virtual conversions 	For more information, refer to: "Virtual Machine Manager" at: http://aka.ms/qc0v35

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
A virtual machine does not use dynamic memory.	

MCT USE ONLY. STUDENT USE PROHIBITED

Module 7

Configuring Advanced Networking Features

Contents:

Module Overview	7-1
Lesson 1: Overview of high-performance networking features	7-2
Lesson 2: Configuring advanced Hyper-V networking features	7-12
Lab: Configuring advanced Hyper-V networking features	7-23
Module Review and Takeaways	7-27

Module Overview

Windows Server 2016 introduces advanced high-performance networking features, such as Server Message Block (SMB) 3.1.1, new Quality of Service (QoS) options, and several enhancements on the receiving end of network packet processing. Additionally, new networking features are available to the Microsoft Hyper-V role and to the virtual machines running under Hyper-V, such as expanded virtual switch functionality and extensibility, single-root I/O virtualization (SR-IOV), dynamic virtual machine queuing, and NIC Teaming for virtual machines.

In this module, you will learn how to deploy and configure the advanced networking enhancements in Windows Server 2016 and the new features in Hyper-V networking.

Objectives

After completing this module, you will be able to:

- Describe the high-performance networking enhancements in Windows Server 2016.
- Configure the advanced Hyper-V networking features.

Lesson 1

Overview of high-performance networking features

Datacenters are becoming increasingly connected to the cloud, to other datacenters, and to servers within the datacenters themselves. This connectivity can slow down the overall performance of the servers. Microsoft has introduced several high-performance networking features to enhance connectivity performance. In this lesson, you will learn about the new and improved networking technologies that Windows Server 2016 introduces.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe NIC Teaming.
- Describe how to configure NIC Teaming.
- Describe how to implement SMB 3.1.1 shared folders.
- Describe advanced SMB 3.1.1 functionality.
- Explain how to provide SMB 3.1.1 high availability in remote storage.
- Describe QoS.
- Describe Receive Side Scaling (RSS).
- Describe Receive Segment Coalescing (RSC).

What is NIC Teaming?

NIC Teaming allows you to combine up to 32 network adapters and then use them as a single network interface. NIC Teaming provides redundancy, allowing network communication to occur over the combined network interface even when one or more of the network adapters fail. The combination of network adapters also increases the bandwidth available to the combined network interface. NIC Teaming is a feature that is available in the Windows Server 2016 operating system. Both the Hyper-V host and the Hyper-V virtual machines can use the NIC Teaming feature. A NIC team can contain only one network adapter, but when it has only one network adapter, the NIC team cannot provide load balancing and failover. Still, you can use a NIC team with only one network adapter in it for the separation of network traffic when you are also using virtual local area networks (VLANs).

- **NIC Teaming:**
 - Provides redundancy and aggregates bandwidth
 - Is supported at the host and virtual machine levels
- **Considerations for NIC Teaming:**
 - Deploy multiple network adapters on a physical host
 - Configure separate teams on different switches for fault tolerance

Considerations for NIC Teaming

If you want to enhance the connectivity fault tolerance and performance of your Hyper-V host, you should deploy multiple network adapters to the host and then configure those adapters as part of a team. This helps to ensure that network connectivity will be retained if individual network adapters fail. Configure multiple teams with network adapters that are connected to different switches to help ensure that connectivity remains if a hardware switch fails.



Note: NIC Teaming within a virtual machine is discussed later in this module.

With Windows Server 2016, you can now use Switch Embedded Teaming (SET) within a Microsoft Hyper-V virtual switch to team up to eight physical network adapters into one or more software-based virtual network adapters. These virtual network adapters deliver fast performance and fault tolerance in the event of a network adapter failure. You must install SET member network adapters in the same physical Hyper-V host to be placed in a SET team. You can also use Remote Direct Memory Access (RDMA)-capable network adapters within a SET team, which allows you to use both the RDMA and SET teams while utilizing fewer network adapters in your servers. This also means you do not have to team at the host level, which provides the significant benefit of managing RDMA at the virtual switch.

Dynamic NIC Teaming was introduced as a new load balancing option in Windows Server 2012 R2. Dynamic NIC Teaming should be compared to the address hash method used prior to Dynamic NIC Teaming in Windows Server 2012 R2. With address hash, when a new data flow is detected, that flow is assigned statically to a team member. The assignment is not based on existing traffic on any of the members of the team. Once assigned, a flow will never move to another team member. This means it is possible for several very large flows to all be on the same team member, while other team members have little traffic. This can result in delayed or dropped packets for these over-used members. Dynamic NIC Teaming constantly watches flows and any time there is a pause, when the flow resumes the traffic on all members is evaluated and the flow moved to the members with less traffic. This means traffic is constantly rebalanced to avoid any one member having significantly more traffic than others.

Demonstration: Implementing NIC Teaming

In this demonstration, you will learn how to implement NIC Teaming.

Demonstration Steps

1. On **LON-HOST1**, open **Server Manager**, and then select the **Local Server** node.
2. In the **Local Server** node, create a NIC team that uses the **vEthernet (Corporate Network)** network adapter, and then name it **Host NIC Team**.
3. In the **NIC Teaming** dialog box, in the **Teams** pane, note the following details:
 - o Team: **Host NIC Team**
 - o Status: **OK**
 - o Teaming Mode: **Switch Independent**
 - o Load Balancing: **Dynamic**
 - o Adapters: **1**

Implementing SMB 3.1.1 shared folders

The latest version of SMB is SMB 3.1.1, which was introduced in Windows 10 and Windows Server 2016. SMB 3.1.1 supports Advanced Encryption Standard (AES) 128 Galois/Counter Mode (GCM) encryption in addition to the AES 128 Counter with CBC-MAC (CCM) encryption that is included in SMB 3.0, and it applies a preauthentication integrity check by using the Secure Hash Algorithm (SHA) 512 hash. SMB 3.1.1 also requires a security-enhanced negotiation when connecting to devices that use SMB 2.x and later.

- SMB 3.1.1 is available only in Windows Server 2016; SMB 3.0 is available in Windows Server 2012; both have similar functionality
- Hyper-V 10.0 can store the following on SMB 3.1.1 file shares:
 - XML-based configuration files
 - Virtual hard disk files (in .vhd or .vhdx format)
 - Checkpoint files

Hyper-V supports storing virtual machine data, such as virtual machine configuration files, checkpoints, and virtual hard disk files, on SMB 3.0 and later file shares. The file share must support SMB 3.0. This limits the placement of virtual hard disks on file shares that are hosted on file servers that are running Windows Server 2012 or later. Earlier Windows Server versions do not support SMB 3.0.



Note: We recommend that the bandwidth for network connectivity to the file share be 1 gigabit per second (Gbps) or more.

An SMB 3.0 file share provides an alternative to storing virtual machine files on Internet Small Computer System Interface (iSCSI) or Fibre Channel storage area network (SAN) devices. When creating a virtual machine in Hyper-V on Windows Server 2012 or later, you can specify a network share when choosing the virtual machine location and the virtual hard disk location. You also can attach disks stored on SMB 3.0 and later file shares. You can use both .vhd and .vhdx disks with SMB 3.0 or later file shares.



Additional Reading:

For more information, refer to: "Server Message Block Overview" at: <http://aka.ms/obywww0>

Since Windows Server 2012 R2, Microsoft has improved SMB 3.0 to allow shared storage for guest clustering that is stored on an SMB 3.0 file server. SMB 3.1.1 continues to support this functionality on Windows Server 2016.

Using advanced SMB 3.1.1 functionality

In Windows Server 2012 R2, several enhancements were made to the SMB 3.0 functionality. Windows Server 2016 continues to support the SMB 3.0 enhancements as well as several advanced functions that you can employ by using SMB 3.1.1. For example, you can store virtual machine files on a highly available SMB 3.1.1 file share. This is referred to as a Scale-Out File Server. By using this approach, you achieve high availability not by clustering Hyper-V nodes but by using file servers that host virtual machine files on their file shares. With this capability, Hyper-V can store all virtual machine files, including configuration files, .vhd files, and checkpoints, on highly available SMB file shares.

SMB 3.0 features that are introduced in Windows Server 2012:

- SMB Transparent Failover
- SMB Scale Out
- SMB Multichannel
- SMB Direct
- SMB Encryption.
- VSS for SMB file shares
- SMB Directory Leasing
- Windows PowerShell commands for managing SMB

The SMB 3.0 features that are introduced in Windows Server 2012 include:

- SMB Transparent Failover. This feature allows you to perform the hardware or software maintenance of nodes in a clustered file server without interrupting server applications that are storing data on file shares.
- SMB Scale Out. By using Cluster Shared Volumes (CSV) version 2, you can create file shares that provide simultaneous access to data files, with direct I/O, through all the nodes in a file server cluster.
- SMB Multichannel. This feature allows you to aggregate network bandwidth and network fault tolerance if multiple paths are available between the SMB 3.0 client and server.
- SMB Direct. This feature supports network adapters that have the Remote Direct Memory Access (RDMA) capability and can perform at full speed with very low data latency and by using very little CPU processing time.
- SMB Encryption. This feature provides the end-to-end encryption of SMB data on untrusted networks and helps to protect data from eavesdropping.
- Volume Shadow Copy Service (VSS) for SMB file shares. To take advantage of VSS for SMB file shares, both the SMB client and the SMB server must support SMB 3.0 at a minimum.
- SMB Directory Leasing. This feature improves branch office application response times. It reduces the number of round trips from client to server as metadata is retrieved from a longer living directory cache.
- Windows PowerShell commands for managing SMB. You can manage file shares on the file server, end-to-end, from the command line.

The new SMB 3.1.1 features that are introduced in Windows Server 2016 are:

- Preauthentication integrity. Preauthentication integrity provides improved protection from a man-in-the-middle attack that might tamper with the establishment and authentication of SMB connection messages.
- SMB Encryption improvements. SMB Encryption, introduced with SMB 3.0, used a fixed cryptographic algorithm: AES-128-CCM. However, AES-128-GCM performs better in most modern processors, so SMB 3.1.1 uses GCM as its first encryption option.
- Cluster Dialect Fencing. Cluster Dialect Fencing provides support for cluster rolling upgrades for the Scale-Out file Servers feature.

- The removal of the **RequireSecureNegotiate** setting. Because some third-party implementations of SMB do not correctly perform this negotiation, Microsoft provides a switch to disable **Secure Negotiate**. However, the default for SMB 3.1.1 servers and clients is to use preauthentication integrity, as described earlier.
- The x.y.z notation for dialects with a nonzero revision number. Windows Server 2016 uses three separate digits to notate the version of SMB. This information is then used to negotiate the highest level of SMB functionality.

Providing highly available remote storage by using SMB 3.1.1

SMB in Windows Server 2016 provides a collection of enhancements that are designed to improve availability, performance, and reliability at the single-server and multiple-server (scale-up and scale-out) levels. These features significantly enhance the availability of remote storage. The SMB remote storage enhancements include:

SMB remote storage enhancements:

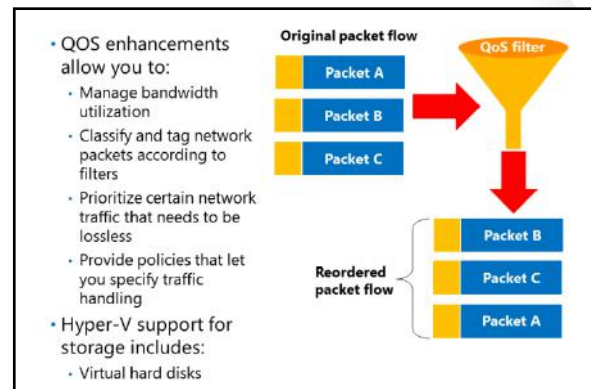
- Hyper-V over SMB
- SMB hardening improvements for SYSVOL and NETLOGON connections
- SMB Multichannel
- SQL Server over SMB
- Storage Spaces Direct
- Storage Replica
- QoS

- Hyper-V over SMB. You can use SMB 3.0 and later file shares as shared storage for Hyper-V in Windows Server. This allows Hyper-V to store virtual machine files, including configuration files, .vhd files, and snapshot files, on SMB file shares.
- SMB hardening improvements for SYSVOL and NETLOGON connections. Client connections to the Active Directory Domain Services (AD DS) default SYSVOL and NETLOGON shares on domain controllers now require SMB signing and mutual authentication in Windows 10 and Windows Server 2016.
- SMB Multichannel. SMB Multichannel allows file servers to simultaneously use multiple network connections. It allows for the aggregation of network bandwidth and network fault tolerance when multiple paths are available between the SMB 3.0 or later client and server. This capability allows server applications to take full advantage of all the available network bandwidth and makes them more resilient to network failures.
- SQL Server over SMB. SQL Server can store user database files on SMB file shares, and this feature adds support for clustered servers running SQL Server and system databases.
- Storage Spaces Direct. Storage Spaces Direct allows you to build highly available and scalable storage systems with local storage. This is a significant advancement in Windows Server software-defined storage for two reasons. First, it makes the deployment and management of software-defined storage systems easier. Second, it unlocks the use of new classes of disk devices, such as Serial ATA and Non-Volatile Memory Host Controller Interface Specification-Enhanced disk devices, that were previously not possible to use with clustered Storage Spaces with shared disks.

- **Storage Replica.** Storage Replica is a new feature that supports storage-agnostic, block-level, synchronous replication between servers or clusters for disaster recovery as well as the stretching of a failover cluster between sites. Synchronous replication provides the mirroring of data in physical sites with crash-consistent volumes that helps to ensure no data loss at the file-system level. Asynchronous replication permits site extension outside metropolitan ranges when a possibility of data loss exists. The Storage Replica functionality:
 - Allows for a single-vendor disaster recovery solution for planned and unplanned outages.
 - Uses SMB 3-level transport, which provides enhanced reliability, scalability, and performance.
 - Stretches failover clusters to metropolitan distances.
 - Uses Microsoft software for end-to-end storage and clustering. Such software includes Hyper-V, Storage Replica, Storage Spaces, Failover Clustering, Scale-Out File Server, SMB 3-level transport, Data Deduplication, Resilient File System, and New Technology File System (NTFS).
 - Helps to reduce the cost and complexity in the following ways:
 - The functionality is hardware agnostic, so no requirement exists for a specific storage configuration, such as direct-attached storage or SAN.
 - The functionality permits commodity storage and networking technologies.
 - Failover Cluster Manager provides ease of graphical management for individual nodes and clusters.
 - Windows PowerShell now includes comprehensive, large-scale scripting options.
 - Helps to decrease downtime and increase the reliability and productivity fundamental to Windows Server.
 - Provides supportability, performance metrics, and diagnostic abilities.
- **Storage QoS.** You use QoS to centrally monitor end-to-end storage performance and make policies by using Hyper-V and Scale-Out File Server in Windows Server 2016.

What is QoS?

QoS is a collection of technologies that allows you to meet the service requirements of a workload or an application by measuring network bandwidth; detecting changing network conditions, such as congestion or the availability of bandwidth; and then prioritizing or throttling network traffic. This means your priority traffic takes precedence over noncritical traffic, and priority traffic processes first. For instance, you can use QoS to prioritize traffic such as voice or video streaming, which are very latency-sensitive applications, and to control the impact of latency-insensitive traffic, such as bulk data transfers.



The following lists several important QoS feature benefits:

- **Bandwidth management.** Hyper-V administrators can use the QoS functionality to manage bandwidth for converging multiple traffic types through a virtual machine network adapter, which allows a predictable service level for each traffic type. You also can allocate minimum and maximum bandwidth allocations on a per-virtual machine basis.
- **Classification and tagging.** Before you can manage the bandwidth for a workload, you need to classify or filter out that workload so that either the QoS Packet Scheduler or a Data Center Bridging (DCB)–capable network adapter can act on it. Windows Server 2016 has an advanced traffic classification capability. A classification can be based on 5-tuples, user types, or Uniform Resource Identifiers (URIs). Windows Server 2016 streamlines the management task so that you can use built-in filters in Windows PowerShell to classify some of the more common workloads.
- **Priority-based flow control (PFC).** Certain workloads, such as RDMA, need lossless transport. When RDMA is built directly on top of Ethernet, it is known as RDMA over Converged Ethernet (RoCE). In this case, the Ethernet transport must be lossless. Traditional link-level flow control, relying on the 802.3 Pause frame, is a solution for this. However, link-level flow control can cause problems—for example, head of line blocking. This issue is resolved by PFC, one of the standards defined by the Institute of Electrical and Electronics Engineers (IEEE) DCB workgroup. Windows Server allows you to enable PFC as long as the physical network adapter supports it. When you enable PFC for RoCE on both ends of the Ethernet link, only the virtual link selected for RoCE, which is designated by a priority value, becomes lossless, and other workloads on the same physical link do not have head of line blocking.
- **Policy-based QoS and Hyper-V QoS.** You use policy-based QoS to manage network traffic on a physical network. This allows you to specify what network bandwidth control measure to use based on application types, users, and computers. You use policy-based QoS to manage traffic, which helps to control bandwidth costs, negotiate service levels with bandwidth providers or business departments, and offer better end-user experiences. Policy-based QoS is configurable through the AD DS Group Policy, is part of your existing management infrastructure, and is consequently a cost-effective solution. You can use a new function in QoS, called Hyper-V QoS, to manage traffic on the virtual network.

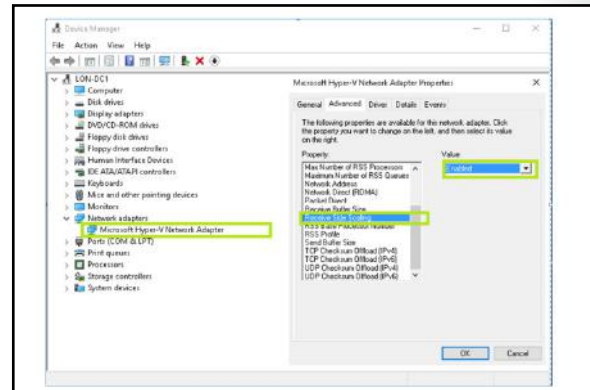
Storage QoS

Starting in Windows Server 2012, Hyper-V includes the ability to set QoS parameters for storage on virtual machines. Virtual hard disks support the configuration of QoS parameters. When you configure the QoS parameters, you can specify the maximum number of input/output operations (IOPS) for the virtual hard disk, which minimizes the chance that a single virtual hard disk will consume the majority of the IOPS capacity of the underlying storage. You also can configure a virtual hard disk to trigger an alert if the number of IOPS falls below a threshold value. IOPS are measured in 8-kilobyte increments. You cannot configure storage QoS when you are using shared virtual hard disks.

Windows Server 2016 now uses storage QoS to manage QoS policies for Hyper-V and Scale-Out File Server. This allows the deployment of QoS policies for SMB 3.1.1 storage.

What is RSS?

Windows Server 2016 supports virtual RSS on the virtual machine network path. This allows virtual machines to support greater network traffic loads. Virtual RSS accomplishes this by spreading the processing load across multiple processor cores on both the Hyper-V host and the virtual machine. A virtual machine can take advantage of virtual RSS improvements only if the processor on the Hyper-V host supports RSS and you have configured the virtual machine to use multiple processor cores.



Virtual RSS allows network adapters to balance the network processing load across the processor cores that are assigned to a virtual machine. Virtual RSS allows a virtual machine to process greater amounts of network traffic than it could process if only a single CPU core was responsible for processing traffic. You can implement virtual RSS by allocating a virtual machine multiple cores through the advanced network. To use virtual RSS, the host's processor must support RSS and the host's network adapters must support Virtual Machine Queue (VMQ).

Enabling virtual RSS

You can use Device Manager or Windows PowerShell to enable virtual RSS. To enable RSS by using Device Manager, perform the following steps:

1. On the virtual machine, open **Device Manager**.
2. Expand **Network adapters**, right-click the network adapter you want to configure virtual RSS on, and then click **Properties**.
3. On the **Advanced** tab, in the network adapter's properties, locate the setting for RSS, and then make sure that it is enabled.



Note: Some network adapters advertise the number of RSS queues they support on the Advanced tab.

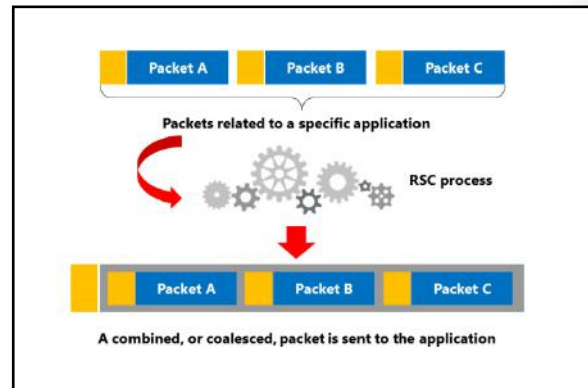
To enable RSS by using Windows PowerShell, perform the following steps:

1. On the virtual machine, open **Windows PowerShell**.
2. At the command prompt, type the following command, and then press Enter.

```
Enable-NetAdapterRSS -Name "AdapterName"
```

What is RSC?

RSC is an offload technology that helps you to reduce how much CPU time is used in network processing. RSC works by having the network adapter look at the incoming data packets and strip them before joining the combined payloads, or coalescing the segments into a single packet. The network adapter then sends the coalesced packet to an application, which results in much less CPU time on the receive side. The CPU can then take care of other important tasks, resulting in increased productivity and scalability support. RSC supports only incoming packets; so it does not affect outgoing packets at all, which the CPU processes normally.



To use RSC, the server must have an RSC-capable network adapter. If you want to use RSC in a virtualized environment, the network adapter must also support SR-IOV.

RSC provides multiple benefits, including:

- Hosted cloud deployments. RSC reduces the number of CPU cycles used for network storage and live migration.
- Faster processing. I/O-heavy database applications and database replication are processed faster.
- Enhanced performance on file servers that are deployed with the Windows Server File Services server role. If your file server is also configured as a BranchCache-enabled content server, BranchCache performance is improved by RSC.
- Improvement on any server workloads that are I/O intensive. I/O intensive workloads are significant consumers of network traffic, and by coalescing the segments, cut down on I/O processing time.

You can use Windows PowerShell to manage RSC. You can use the cmdlets **Get-NetAdapterRsc** and **Get-NetAdapterStatistics** to see the network adapter's RSC configuration. Use the cmdlet **Enable-NetAdapterRsc** to enable RSC.



Additional Reading:

For more information on the preceding Windows PowerShell cmdlets, refer to: "Network Adapter Cmdlets in Windows PowerShell" at: [https://technet.microsoft.com/en-us/library/jj134956\(v=wps.630\).aspx](https://technet.microsoft.com/en-us/library/jj134956(v=wps.630).aspx)

Categorize Activity

Categorize each item into the appropriate category. Indicate your answer by writing the category number to the right of each item.

Items	
1	This allows you to combine up to 32 network adapters and then use them as a single network interface.
2	This configuration can be deployed with only one network adapter but does not offer fault tolerance.
3	To use this, the host must have at least two external virtual switches.
4	To use this, you must configure a virtual machine to use multiple CPU cores.
5	You can configure this through Device Manager or Windows PowerShell
6	You can use this to prioritize traffic such as voice or video streaming.
7	This can help you to implement bandwidth management.
8	You can implement this by allocating a virtual machine's multiple cores through the advanced network.
9	This is a collection of technologies that allow you to meet the service requirements of a workload.

Category 1		Category 2		Category 3
NIC Teaming		QoS		RSS

Lesson 2

Configuring advanced Hyper-V networking features

Hyper-V provides several options for allowing network communication among virtual machines. You can use Hyper-V to configure virtual machines that communicate with an external network in a manner similar to that for physical hosts that you deploy traditionally. You also can use Hyper-V to configure virtual machines that can communicate only with a limited number of other virtual machines that are hosted on the same Hyper-V host. Windows Server 2016 provides several advanced networking features for Hyper-V and virtual machines. This lesson describes the various advanced features that are available for Hyper-V virtual networks, which you can use to best meet your organization's needs.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the virtual switch expanded functionality.
- Describe virtual switch extensibility.
- Describe SR-IOV.
- Describe dynamic virtual machine queuing.
- Describe the network adapter advanced features
- Describe NIC Teaming in virtual machines.
- Describe how to configure the network adapter advanced features.

Virtual switch expanded functionality

Virtual switches are virtual devices that you can manage through the Virtual Switch Manager, which allows you to create three types of virtual switches. Virtual switches control how network traffic flows among the virtual machines that are hosted on a Hyper-V server and how network traffic flows between the virtual machines and the rest of the organizational network.

Hyper-V in Windows Server 2012 and in Windows Server 2016 supports three types of virtual switches, which the following table details.

The virtual switch improvements in Windows Server 2016 include:

- Extended port ACLs
- Dynamic load balancing
- Coexistence with third-party forwarding extensions
- RSS support on the virtual machine network path
- Network tracing enhancements
- Router guarding
- DHCP guarding
- Trunk mode for virtual machine
- Port mirroring
- VLAN isolation through a Private VLAN
- Extended bandwidth management

Type	Description
External	You use this type of switch to map a network to a specific network adapter or network adapter team. Windows Server 2012 supports mapping an external network to a wireless network adapter if you have installed the Wireless LAN service on the host Hyper-V server and if the Hyper-V server has a compatible network adapter.
Internal	You use internal virtual switches to communicate among the virtual machines on a Hyper-V host and to communicate between the virtual machines and the Hyper-V host itself.

Type	Description
Private	You use private switches only to communicate among the virtual machines on a Hyper-V host. You cannot use private switches to communicate between the virtual machines and the Hyper-V host.

When configuring a virtual network, you also can configure a VLAN ID to associate with the network. You can use this to extend the existing VLANs on an external network to VLANs within the Hyper-V host's network switch. You can use VLANs to partition network traffic. VLANs function as separate logical networks. Traffic can pass from one VLAN to another only if it passes through a router.

You can configure the following extensions for each virtual switch type:

- Microsoft Network Driver Interface Specification (NDIS) Capture. This extension allows for the capture of data that travels across a virtual switch.
- Microsoft Windows Filtering Platform (WFP). This extension allows filtering of data that travels across a virtual switch.

Windows Server 2012 introduced many new features that are now available in the virtual switch expanded functionality. Several more features were added in Windows Server 2012 R2. These features remain an important part of Windows Server 2016 and continue to improve network performance and the flexibility of virtual machines in private and public cloud environments.

Features included in Windows Server 2012 Hyper-V networking

The features that were added in Windows Server 2012 Hyper-V networking include:

- Network virtualization. This feature allows IP addresses to be virtualized in hosting environments so that virtual machines that migrate to the host can keep their original IP addresses, rather than being allocated IP addresses on the Hyper-V server's network.
- Bandwidth management. You can use this feature to specify a minimum and a maximum bandwidth that Hyper-V will allocate to the adapter. Hyper-V reserves the minimum bandwidth allocation for the network adapter even when other virtual network adapters on virtual machines that are hosted on the Hyper-V host are functioning at capacity.
- Dynamic Host Configuration Protocol (DHCP) guard This feature drops DHCP messages from virtual machines that are functioning as unauthorized DHCP servers. This might be necessary in scenarios where you are managing a Hyper-V server that hosts virtual machines for others but where you do not have direct control over the virtual machines' configurations.
- Router guard. This feature drops router advertisement and redirection messages from virtual machines that are configured as unauthorized routers. This might be necessary in scenarios where you do not have direct control over the configuration of the virtual machines.
- Port mirroring. You can use this feature to copy incoming and outgoing packets from a network adapter to another virtual machine that you have configured for monitoring.
- NIC Teaming. You can use this feature to add a virtual network adapter to an existing team on the host Hyper-V server.
- VMQ. This feature requires the host computer to have a network adapter that supports the feature. VMQ uses hardware packet filtering to deliver network traffic directly to a guest. This improves performance because the packet does not need to be copied from the host operating system to the virtual machine. Only network adapters that are specific to Hyper-V support this feature.

- SR-IOV. This feature requires that specific hardware and special drivers be installed on the guest operating system. SR-IOV enables multiple virtual machines to share the same peripheral component interconnect (PCI) Express physical hardware resources. If sufficient resources are not available, network connectivity fallback occurs so that the virtual switch provides that connectivity. This feature is supported only on network adapters that are specific to Hyper-V.
- Internet Protocol security (IPsec) task offloading. This feature requires that the guest operating system and network adapter are supported. This feature allows a host's network adapter to perform calculation-intensive security-association tasks. If sufficient hardware resources are not available, the guest operating system performs these tasks. You can configure a maximum number of offloaded security associations from 1 through 4,096. This feature is supported only on network adapters that are specific to Hyper-V.
- Private VLANs. A VLAN ID is a 12-bit number in the range 1 through 4,095. The configuration of multiple, isolated VLANs is complex and difficult. However, when you deploy Hyper-V Network Virtualization, many of these complex and difficult issues are solved, but not completely. A simpler solution is to use a private VLAN. A private VLAN tackles some of the scalability issues of VLANs. A private VLAN is a property of a switch port. With a private VLAN two VLAN IDs exist: a primary VLAN ID and a secondary VLAN ID. A private VLAN can exist in one of three modes:
 - Isolated. Communicates only with promiscuous ports in the private VLAN.
 - Promiscuous. Communicates with all ports in the private VLAN.
 - Community. Communicates with ports in the same community and with any promiscuous ports in the private VLAN.
- Trunk mode. Trunk mode allows network services or network appliances on a virtual machine to see traffic from multiple VLANs. In trunk mode, a switch port receives traffic from all the configured VLANs in an allowed VLAN list. You can also configure a switch port that is connected to a virtual machine, but it is not bound to the underlying network adapter.

Features included in Windows Server 2012 R2 Hyper-V networking

The features that were added in Windows Server 2012 R2 Hyper-V networking include:

- Extended port access control lists (ACLs). You can use extended port ACLs in a Hyper-V virtual switch to help enforce security policies and firewall protection at the switch level for virtual machines. The differences between ACLs in Windows Server 2012 and Windows Server 2012 R2 Hyper-V include:
 - Administrators can now include socket port numbers when developing ACLs.
 - Hyper-V switches support unidirectional, stateful rules with a timeout parameter.
- The dynamic load balancing of network traffic. When you map a virtual network to a network adapter team on a Windows Server 2012 R2 Hyper-V host, the network traffic will be continuously load balanced across network adapters, with traffic streams moved as necessary to maintain this balance. In Windows Server 2012 Hyper-V, a traffic stream remained with the network adapter in the team that it was initially assigned to, and traffic streams were not dynamically moved to other network adapters in the team.
- Coexistence with non-Microsoft forwarding extensions. The Hyper-V Network Virtualization module forwards network traffic that is encapsulated through Network Virtualization Generic Routing Encapsulation (NVGRE). Non-Microsoft switch extensions are supported in coexistence scenarios with Hyper-V virtual switches. When a non-Microsoft extension is present, any non-NVGRE network traffic is forwarded via the non-Microsoft forwarding extensions.

- RSS on the virtual machine network path. Windows Server 2012 R2 supports virtual RSS on the virtual machine network path. This allows virtual machines to support greater network traffic loads. Virtual RSS accomplishes this by spreading the processing load across multiple processor cores on both the Hyper-V host and the virtual machine. A virtual machine can take advantage of virtual RSS improvements only if the processor on the Hyper-V host supports RSS and, if you configure the virtual machine to use multiple processor cores.
- Network tracing improvements. You use **Netsh Trace** commands to trace packets. The improvements in Windows Server 2012 R2 allow you to view port and switch information as you trace network traffic through Hyper-V virtual switches.

Features added in Windows Server 2016 Hyper-V networking

The features that were added in Windows Server 2016 Hyper-V networking include:

- Network function virtualization. In most datacenters, hardware appliances handle some network functions or services, such as software load balancing and network address translation, services provided by datacenter firewalls, and Remote Access Service gateway services. However, with software-defined networking, more appliances are becoming virtualized. All three functions are available in Windows Server 2016.
- Network Controller. By using Network Controller, you can have a central location to monitor, manage, troubleshoot, and configure both your physical and your virtual environment.
- Switch Embedded Teaming (SET). SET is a new NIC Teaming option that you can use for Hyper-V networks. SET has some integrated functionality with Hyper-V that provides faster performance and better fault tolerance than traditional teaming.
- RDMA with Hyper-V. RDMA services can now use Hyper-V switches. You can enable this feature with or without SET.
- Multiple queues for virtual machines. This feature allocates multiple hardware queues for each virtual machine, thereby improving throughput as compared to Windows Server 2012 R2.
- Converged network adapters. A converged network adapter supports using a single network adapter or a team of network adapters to handle multiple forms of traffic, management, RDMA, and virtual machine traffic. This reduces the number of specialized adapters that are needed on each host.

Understanding virtual switch extensibility

The new Hyper-V Extensible Switch feature supports isolation policies, allows extensibility, and lets third-party vendors add filters to provide their own forwarding rules.

The Hyper-V Extensible Switch is a layer 2 virtual network switch that is used to connect virtual machines to the physical network by providing programmatically managed and extensible capabilities. The Hyper-V Extensible Switch permits policy enforcement for security enhancement, isolation, and service levels. The Hyper-V Extensible Switch provides for non-Microsoft extensible plug-ins that can provide enhanced networking and security capabilities by using NDIS filter drivers and WFP callout drivers for support.

- Virtual switch extensions allow third-party vendors to create virtual switches
- You can manage virtual switches by using the same tool set that you use to manage physical switches



You can implement and manage virtualized datacenters with the Hyper-V Extensible Switch through:

- An open platform. Built on an open platform, The Hyper-V Extensible Switch allows third-party software vendors to add or extend the capabilities that are natively provided by Microsoft. The abilities of the Hyper-V Extensible Switch can combine with the added capabilities of vendor extensions, which can then be applied to implement and manage virtualized datacenters.
- A standard API. The programming model uses the same NDIS and WFP application programming interface (API) that was used for network filters and drivers in earlier versions of Windows. Several new API functions and parameters have been added for virtual switch ports.
- Windows reliability and quality. The Hyper-V Extensible Switch uses the Windows operating system and the Windows Hardware Quality Logo program to set high standards for extension quality.
- Policy and configuration integration. The management of extensions provides a standard management approach by integrating Windows management through Windows Management Instrumentation calls and Windows PowerShell cmdlets. You can automatically migrate policies of extensions with the virtual machine configuration during a live migration.
- Easy troubleshooting options. Included with the Hyper-V Extensible Switch are various event logs and unified tracing, which makes it easier to diagnose and troubleshoot any issues.

You can extend or replace three aspects of the switching process with extensions: ingress filtering, destination lookup and forwarding, and egress filtering. Additionally, you can use extensions to gather statistical data by monitoring traffic at different layers of the Hyper-V Extensible Switch. Multiple monitoring and filtering extensions can be supported at the ingress and egress portions of the Hyper-V Extensible Switch. However, only one instance of the extension can be used per switch instance if you use a forwarding extension. In this case, it overrides the default forwarding option of the Hyper-V Extensible Switch.

The following table shows the types of extensions, their purposes, the components used to implement them, and examples.

Extension	Purpose	Component	Examples
Intrusion detection or firewall	Allows filtering and modifying TCP/IP packets, monitoring or authorizing connections, and filtering traffic that is protected by IPsec and filter remote procedure calls.	WFP callout driver	Virtual firewall, connection monitoring
Network forwarding	Provides a forwarding extension per Hyper-V Extensible Switch instance, which bypasses the default forwarding option (with a maximum of one per Hyper-V Extensible Switch instance).	NDIS filter driver	OpenFlow, Virtual Ethernet Port Aggregator, proprietary network fabrics

Extension	Purpose	Component	Examples
Network packet filter	Creates, filters, and modifies network packets that are entering or leaving the Hyper-V Extensible Switch and that exist in virtual machine-to-virtual machine traffic.	NDIS filter driver	Security enhancement
Network packet inspection	Views network packets for virtual machine-to-virtual machine traffic per Hyper-V Extensible Switch instance. This extension cannot alter network packets.	NDIS filter driver	sFlow, network monitoring

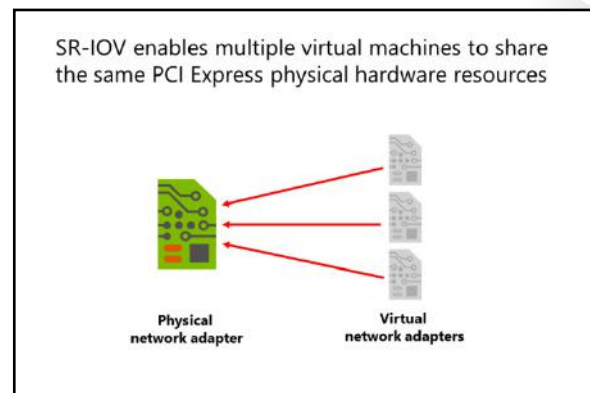
What is SR-IOV?

SR-IOV allows multiple virtual machines to share the same PCI Express physical hardware resources. If sufficient resources are not available, network connectivity fallback occurs, and the virtual switch provides connectivity. SR-IOV requires that you install specific hardware and special drivers on the guest operating system, and you might need to enable it in the computer BIOS.

Note, however, that only 64-bit Windows and Windows Server guest virtual machines, starting with Windows Server 2012 and Windows 8, support SR-IOV. In this case, you should disable

SR-IOV on all the virtual machines that do not support SR-IOV. For those guest operating systems that can use SR-IOV, the physical network adapter on the parent partition—that is the Hyper-V host—appears on the guest operating system. You can use Device Manager to see the physical network adapter and even to upgrade or manage the device driver for it within the guest operating system. In other words, the virtual machine communicates with the physical hardware.

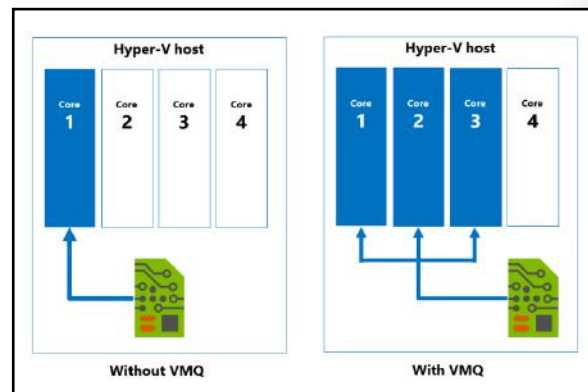
SR-IOV uses Virtual Functions (VF). VFs are associated with a Physical Function (PF). The PF is what the parent partition uses in Hyper-V and is equivalent to the regular bus-addressed, device-addressed, or function-addressed PCI device. The responsibility for arbitration relating to policy decisions, such as those for link speed or media access control (MAC) addresses in use by virtual machines and for I/O from the parent partition, is handled by the PF. Although the parent partition can use a VF, in Windows Server, only virtual machines use VFs. A single PCI Express device can expose multiple VFs, such as a multiple-port networking device, with each port independent and with its own set of VF resources.



When using SR-IOV, a part of the network adapter's hardware is made available to the virtual machine. Because the guest operating system's networking code doesn't know how to directly manipulate that hardware, you will need to load a vendor-supplied driver in the virtual machine. Note that the VF is not a complete device or autonomous. It cannot make any decisions about policy and control. The VF can only read and write the parts of the device configuration that the PF lets it handle, and it can only see the parts of networking hardware in memory space that are allocated to the VF. VFs are transient, because the guest operating system is also transient, in the sense that you can start, stop, or even delete it. However, the PF is always available, and is the arbiter for all policy decisions.

What is dynamic VMQ?

VMQ was developed to be a hardware virtualization technology for the efficient transfer of network traffic to a virtualized guest operating system. A VMQ-capable network adapter categorizes incoming frames to be routed to a receive queue based on filters that associate the queue with a virtual machine's virtual network adapter. A queue is assigned to each virtual machine device buffer, which avoids needless packet copies and route lookups in the virtual switch. VMQ makes a single network adapter on a physical host appear as multiple network adapters to the virtual machines. This, in turn, allows each virtual machine to have its own dedicated network adapter. VMQ provides separate queues for the hardware device. In static VMQ, the Hyper-V administrator can manually set the processor affinity of the hardware queues to different CPU cores, which creates RSS on a per-virtual machine network adapter.



Dynamic VMQ dynamically distributes incoming network traffic processing to physical host CPU cores based on processor usage and network load. During periods of heavy network loads, dynamic VMQ automatically employs more processors. When the network load is light, dynamic VMQ relinquishes those same processors. Dynamic VMQ spreads interrupts for network traffic across the available processors. In Windows Server 2012 and later, dynamic VMQ allows an adaptive algorithm to modify the CPU affinity of queues without requiring the removal and re-creation of queues. This results in a better network load-to-processor use match, which helps to increase network performance.

Dynamic VMQ requires the host computer to have a network adapter that supports the feature. Dynamic VMQ uses hardware packet filtering to deliver network traffic directly to a guest. This helps to improve performance, because the packet does not need to be copied from the host operating system to the virtual machine. Only network adapters that are specific to Hyper-V support this feature.

Dynamic VMQ is enabled by default in Windows Server 2016. You can enable or disable it by using the **Enable-NetAdapterVmq** and **Disable-NetAdapterVmq** Windows PowerShell cmdlets respectively.

Dynamic VMQ is very similar to RSS, which was mentioned in Lesson 1. On a physical host, RSS processes incoming network traffic so that a single CPU core does not slow it down. RSS does this by spreading the calculations across multiple CPU cores. For a Hyper-V host that has several virtual machines with significant incoming traffic, dynamic VMQ is similar to RSS. Dynamic VMQ hashes the destination MAC address, puts the traffic for a particular virtual machine in a specific queue, and distributes the interrupts to the CPU cores. Dynamic VMQ handles this by offloading these functions to the network adapters. In dynamic VMQ, a rare circumstance can occur when processing that is happening on a CPU core generates a large amount of inbound traffic. This triggers dynamic VMQ to use another, less-busy CPU core, and

because the traffic load has not changed, it jumps back to the original or another CPU core. This process continues and is referred to as the *ping-pong effect*. Although dynamic VMQ is more automatic, RSS can better avoid the ping-pong effect in this situation.

Network adapter advanced features

Windows Server 2016 continues to improve the software-defined network infrastructure. A software-defined network is a primary building block of a software-defined datacenter. You can easily manage many of these features through Microsoft System Center Virtual Machine Manager. However, you can also use Windows PowerShell commands to configure the implementations of these features. The default settings are adequate in most small-scale environments. Some of the new or improved features are:

- Network function virtualization
- Network Controller
- SET
- RDMA
- VMQ
- Converged network adapters
- QoS for software-defined networks

- Network function virtualization. In most datacenters, hardware appliances handle some network functions, such as software load balancing and network address translation. However, with software-defined networking, more appliances are becoming virtualized. All three functions are available in Windows Server 2016.



Note: To use containers and build out Hyper-V virtualized networks more efficiently, it is important that you have the ability to use network address translation with Windows Server 2016 as a built in feature of a virtual switch. You can create a virtual switch in a virtual machine container host by running the following command:

```
New-VMSwitch -Name "Virtual Switch Name" -SwitchType NAT
```

- Network Controller. By using Network Controller, you can have a central location to monitor, manage, troubleshoot, and configure both your physical and your virtual environment.
- SET. SET is a new NIC Teaming option that you can use for Hyper-V networks. SET has some integrated functionality with Hyper-V that provides faster performance and better fault tolerance than traditional teaming.
- RDMA with Hyper-V. RDMA services can now use Hyper-V switches. You can enable this feature with or without SET.
- Multiple queues for virtual machines. This feature allocates multiple hardware queues for each virtual machine, thereby improving throughput as compared to Windows Server 2012 R2.
- Converged network adapters. A converged network adapter supports using a single network adapter or a team of network adapters to handle multiple forms of traffic, management, RDMA, and virtual machine traffic. This reduces the number of specialized adapters that are needed on each host.
- QoS for software-defined networks. This feature manages the default class of traffic through the virtual switch within the default class bandwidth.

Hardware acceleration features

The hardware acceleration features specify network tasks that can be offloaded to a physical network adapter. Many of the hardware acceleration features are enabled by default in a virtual network adapter, but that does not mean the virtual machine actually uses them all. All the hardware acceleration settings require hardware support. To configure the hardware acceleration settings for a virtual machine network adapter:

1. In the **Hyper-V Manager** console, right-click the virtual machine, and then click **Settings**.
2. In the **Settings** window for the virtual machine, select and expand the network adapter that you want to manage.
3. Note two subnodes: **Hardware Acceleration** and **Advanced Features**. Click the **Hardware Acceleration** node.
4. In the **details** pane, note the various settings. Some are already selected. You can enable or disable the various features on this page.

The features that you can enable and disable are:

- VMQ. VMQ requires a physical network adapter that supports this feature.
- IPsec task offloading. This technology supports hardware-equipped network adapters to reduce the CPU load by performing the computationally intensive work of encryption and decryption. You can also specify the maximum number of offloaded security associations in the range 1 through 4096. The default is 512.

NIC Teaming in virtual machines

When used with virtual machines, NIC Teaming allows the virtual machines to team the virtual network adapters that connect to separate virtual switches. To get the benefit of NIC Teaming, the host must have at least two external virtual switches. When you have multiple virtual network adapters attached to the same switch, if the physical network adapter that the virtual switch is connected to fails, those virtual network adapters will lose connectivity. When configuring NIC Teaming for virtual machines, the network adapters connected to virtual switches can use SR-IOV.

- **NIC Teaming in virtual machines:**
 - Requires multiple virtual network adapters
 - Must be enabled on the virtual network adapters
 - Allows you to then implement it in the virtual machine's operating system (if supported)
- **SET:**
 - Allows you to group from one through eight physical network adapters into one or more virtual network adapters

Enable virtual machine NIC Teaming for virtual machines on the **Advanced Features** page of the virtual network adapter in Hyper-V Manager. You can also enable NIC Teaming for virtual machines by using the **Set-VMNetworkAdapter** Windows PowerShell cmdlet. To enable NIC Teaming within the virtual machine's operating system, you must enable NIC Teaming on the virtual network adapter or configure the virtual network adapter to allow MAC address spoofing. After you enable virtual NIC Teaming on the virtual network adapter or enable MAC address spoofing, you can configure NIC Teaming within the virtual machine.

Dynamic NIC Teaming was first introduced in Windows Server 2012. It allows new traffic to be assigned to a particular network adapter, and the traffic flow remains with that network adapter throughout the session. Dynamic NIC Teaming balances the traffic flow across all the available network adapters in a team.

SET

SET allows you to use fewer network adapters when you want to use RDMA and SET at the same time. SET is an alternative to NIC Teaming that you can use in environments that include Windows Server 2016 Hyper-V and the Software-Defined Networking stack. SET incorporates some of the NIC Teaming functions into a Hyper-V virtual switch.

SET allows you to group from one through eight physical Ethernet network adapters into one or more virtual network adapters. These virtual network adapters then help to provide faster performance and fault tolerance in the event of a failure of any network adapter. To place the member network adapters in a SET team, you must install them in the same physical Hyper-V host.

Demonstration: Configuring network adapter advanced features

In this demonstration, you will learn how to implement advanced features for network adapters on virtual machines.

Demonstration Steps

Use Windows PowerShell to enable DHCP guarding

- On **LON-HOST1**, open the **Windows PowerShell** window, and then run the following cmdlet:

```
Set-VMNetworkAdapter -VMName 20743A-LON-DC1-B -DhcpGuard On
```

Turn off DHCP guarding

- On the physical host computer, at the Windows PowerShell prompt, type the following command, and then press Enter:

```
Set-VMNetworkAdapter -VMName 20743A-LON-DC1-B -DhcpGuard Off
```

Check Your Knowledge

Question	
What is the ping-pong effect”?	
Select the correct answer.	
<input type="checkbox"/>	The ping-pong effect occurs when multiple physical network adapters from the host are matched to several virtual network adapters. They continuously swap physical addresses.
<input type="checkbox"/>	The ping-pong effect occurs when a virtual switch extension applies network forwarding. It bypasses the default forwarding, which causes network packets to loop back and forth to the router.
<input type="checkbox"/>	The ping-pong effect results from a rare circumstance that can occur in dynamic VMQ when a CPU core is being used, and the processing happens to generates a large amount of inbound traffic. Because of this, another, less-busy CPU core is dynamically selected, and because the traffic load has not changed, it jumps back to the original or another CPU core. This process continues.
<input type="checkbox"/>	When you use Remote Direct Memory Access (RDMA), a network adapter can switch repeatedly between Switch Embedded Teaming (SET) and RDMA functionality.
<input type="checkbox"/>	The ping-pong effect occurs when a NIC team switches repeatedly among team member adapters.

Lab: Configuring advanced Hyper-V networking features

Scenario

A. Datum Corporation has implemented the Hyper-V virtualization platform in one of their subsidiaries. You have created several test virtual machines and familiarized yourself with many of the configuration options. The next step is to implement and test network connectivity for the virtual machines.

Objectives

After completing this lab, you will be able to:

- Create and use Hyper-V virtual switches.
- Configure bandwidth management and DHCP guarding.

Lab Setup

Estimated Time: 30 minutes

Physical computer: Restart to **20743A-LON-HOST1**

Virtual machines: **20743A-LON-DC1-B**, **LON-GUEST1**, and **LON-GUEST2**

User name: **Adatum\Administrator**

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

For this lab, you will need to start **20743A-LON-HOST1**. Restart the physical computer, and in the boot menu that appears, select **20743A-LON-HOST1**. Sign into **LON-HOST1** as the **Adatum\Administrator** with a password of **Pa\$\$w0rd**. You use the available virtual machine environment. To start the lab, complete the following steps:

1. On the **LON-HOST1**, click **Start, All Apps**, scroll down to **Windows Administrative Tools**, and then click **Hyper-V Manager**.
2. In Hyper-V Manager, click **20743A-LON-DC1-B**, 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:
 - o User name: **Adatum\Administrator**
 - o Password: **Pa\$\$w0rd**
5. Do not start **LON-GUEST1** and **LON-GUEST2** until directed to do so. These virtual machines were created in the previous module.

Note that to perform this lab successfully, you need to first complete all lab tasks in Module 6.

Exercise 1: Creating and using Hyper-V virtual switches

Scenario

A. Datum Corporation has the Hyper-V virtualization platform already installed. Before deploying Hyper-V and virtual machines in the production environment, you need to ensure that you understand the different networking options that you can configure in Hyper-V. First, you will review the current networking configuration of the Hyper-V host. Then you will create new virtual network adapters in the parent partition. You will also create different types of Hyper-V virtual switches and explore the connectivity options that exist when using each of the switches.

The main tasks for this exercise are as follows:

1. Verify the current Hyper-V network configuration.
2. Create virtual network adapters.
3. Use the Hyper-V virtual switches.
4. Add NIC Teaming.

► **Task 1: Verify the current Hyper-V network configuration**

1. On **LON-HOST1**, open **Hyper-V Manager**.
2. In Hyper-V Manager, open the **Virtual Switch Manager**, and then note the **Corporate Network**, **Private Network**, and **Internal Network** switches that have been created for **LON-HOST1**.
3. Select **Internal Network** and then remove the check mark next to **Enable virtual LAN identification for management operating system**. Click **OK**.
4. Right-click **20743A-LON-DC1-B** and then click **Settings**.
5. In the **Settings for 20743A-LON-DC1-B on LON-HOST1** dialog box, select **Network Adapter**.
6. In the details pane, under **VLAN ID**, remove the check mark next to **Enable virtual LAN identification**.
7. Click **OK**.
8. Repeat steps 4-7 for **LON-GUEST2**.

► **Task 2: Create virtual network adapters**

1. On **LON-HOST1**, open **Windows PowerShell** and then type the following commands. Press Enter after each line:

```
Add-VMNetworkAdapter -VMName LON-GUEST2 -Name "Network Adapter 2"
```

```
Connect-VMNetworkAdapter -VMName LON-GUEST2 -Name "Network Adapter 2" -SwitchName "Corporate Network"
```

```
Add-VMNetworkAdapter -VMName LON-GUEST2 -Name "Network Adapter 2"
```

```
Connect-VMNetworkAdapter -VMName LON-GUEST2 -Name "Network Adapter 3" -SwitchName "Corporate Network"
```

► **Task 3: Use the Hyper-V virtual switches**

1. On **LON-HOST1**, open **Hyper-V Manager**.
2. In Hyper-V Manager, in the **Virtual Machines** pane, find and then right-click **LON-GUEST2**, and then click **Settings**.
3. In the **Settings for LON-GUEST2 on LON-HOST1** window, in the console tree, select **Network Adapter 2**.
4. Note that the virtual switch assigned is **Corporate Network**.
5. In the **Settings** window, click **Cancel**.
6. In the **Hyper-V Manager** console, start and then connect to **LON-GUEST2**.

7. Sign in as **Administrator** with the password **Pa\$\$w0rd**.
8. In the **Server Manager** console tree, select the **Local Server** node.
9. Notice the network adapters; **Ethernet**, **Ethernet 2**, and **Ethernet 3**, which are configured to use DHCP to obtain IP address information.

► **Task 4: Add NIC Teaming**

1. On **LON-GUEST2**, in **Server Manager**, select the **Local Server** node.
2. In the **Local Server** node, create a NIC team that uses the **Ethernet 2** and **Ethernet 3** virtual network adapters, and then name it **LON-GUEST2 NIC Team**.
3. In the **NIC Teaming** dialog box, in the **Teams** pane, note the following:
 - Team: **LON-GUEST2 NIC Team**
 - Status: **OK** (This may show **fault** depending upon if you have external connectivity)
 - Teaming Mode: **Switch Independent**
 - Load Balancing: **Address Hash**
 - Adapters: **2**

Results: After completing this exercise, you should have successfully configured the Hyper-V virtual switch.

Exercise 2: Configuring and using the advanced features of a virtual switch

Scenario

One of your managers wants to see how the Hyper-V virtual switch can help to protect the network clients from unauthorized DHCP servers. You plan to demonstrate how to configure DHCP guarding and, at the same time, show the manager how simple it is to configure bandwidth management.

The main tasks for this exercise are as follows:

1. Configure DHCP Guard.
2. Configure and use bandwidth management.
3. Prepare for the next module.

► **Task 1: Configure DHCP Guard**

1. On **LON-HOST1**, open **Hyper-V Manager**.
2. In Hyper-V Manager, open the settings for the **LON-GUEST1** virtual machine.
3. Under **Settings**, in the **Network Adapter** node, open **Advanced Features**.
4. Enable **DHCP Guard**.
5. Repeat steps 2-4 for **LON-GUEST2**.

► **Task 2: Configure and use bandwidth management**

1. While still on **LON-HOST1**, in Hyper-V Manager, in the **Settings** for the virtual machine **LON-GUEST2**, in the **Network Adapter 2** settings, click **Enable Bandwidth Management**, and then set the **Maximum Bandwidth** to **100** megabits per second (Mbps).

Results: After completing this exercise, you should have successfully configured the advanced features of the Hyper-V virtual switch.

► **Task 3: Prepare for the next module**

1. Shut down the following virtual machines:
 - **LON-GUEST2**
 - **20743A-LON-DC1-B**
2. Restart the host computer and then at the boot menu select Windows Server 2012.

Question: In the "NIC Teaming" task, you created **LON-GUEST2 NIC Team**. Is this fault tolerant?

Question: In the task named "Create virtual network adapters," you the **LON-GUEST2** virtual machine was shut down. Why?

Module Review and Takeaways

Review Question

Question: You want to deploy a Windows Server 2016 Hyper-V virtual machine's virtual hard disk on a file share. What operating system must the file server be running to support this configuration?

Best Practices

When implementing advanced networking features for Hyper-V, use the following best practices:

- Deploy multiple network adapters to a Hyper-V physical host, and then configure those adapters as part of a team. This helps to ensure that network connectivity will be retained if individual network adapters fail. Configure multiple teams with network adapters that are connected to different switches to help ensure that connectivity will remain if a hardware switch fails.
- Use bandwidth management to allocate a minimum and a maximum bandwidth allocation on a per-virtual network adapter basis. You should configure bandwidth allocation to help guarantee that each virtual machine will have a minimum bandwidth allocation. This helps to ensure that if another virtual machine that is physically hosted on the same Hyper-V server experiences a traffic spike, other virtual machines will be able to communicate normally with the network.
- Provision a Hyper-V physical host with an adapter that supports VMQ. VMQ uses hardware packet filtering to deliver network traffic directly to a virtual machine. This helps to improve performance because the packet does not need to be copied from the physical host operating system to the virtual machine. When you do not configure virtual machines to support VMQ, the physical host operating system can become a bottleneck when it processes large amounts of network traffic.
- If you are physically hosting large numbers of virtual machines and need to isolate them, use network virtualization rather than VLANs. Network virtualization is complicated to configure, but it has an advantage over VLAN—it is not necessary to configure VLANs on all the switches that are connected to the Hyper-V physical host. You can perform all the necessary configurations when you need to isolate servers on a Hyper-V physical host without needing to involve the network team.

MCT USE ONLY. STUDENT USE PROHIBITED

Module 8

Implementing Software Defined Networking

Contents:

Module Overview	8-1
Lesson 1: Overview of Software Defined Networking	8-2
Lesson 2: Implementing network virtualization	8-11
Lesson 3: Implementing Network Controller	8-16
Lab: Deploying Network Controller	8-29
Module Review and Takeaways	8-34

Module Overview

Software Defined Networking (SDN) bypasses the limitations imposed by physical network devices and allows organizations to dynamically manage their networks. SDN uses an abstraction layer in software to manage your network dynamically. When you implement SDN, you can virtualize your network, define policies to manage network traffic, and manage your virtualized network infrastructure.

Objectives

After completing this module, you will be able to:

- Describe Software Defined Networking.
- Implement network virtualization.
- Implement Network Controller.

Lesson 1

Overview of Software Defined Networking

SDN enables you to centrally configure and manage both the physical and virtual network devices in your datacenter, such as switches, routers, and gateways, so that you can provide an automated means of responding to application and workload requirements. In Windows Server 2016, virtualization features, including Hyper-V Virtual Switch, Hyper-V Network Virtualization (HNV), and Remote Access Service (RAS) Gateway, are integrated into your SDN infrastructure.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe SDN in Windows Server 2016.
- Explain when to use SDN.
- Plan to implement SDN.
- Describe how to deploy SDN by using scripts.

What is Software Defined Networking?

Although Software Defined Networking still requires a physical network layer, with SDN, you can:

- Virtualize the network so that you can break the direct connection between the applications and virtual servers and the underlying physical network. To do this, you need to virtualize network management by creating virtual abstractions for network elements such as IP addresses, ports, and switches.
- Define policies that will manage traffic flow across both physical and virtual networks. You define these policies in the management system but apply them at the physical layer.
- Manage the virtualized network infrastructure by providing the tools to configure the virtual network objects and policies.

- Software Defined Networking enables you to:
 - Virtualize the network layer in a datacenter
 - Define policies for the physical and virtual networks
 - Manage the virtualized network infrastructure
- The Microsoft Software Defined Networking solution includes:
 - Network Controller
 - Hyper-V Network Virtualization
 - Hyper-V Virtual Switch
 - RRAS Multitenant Gateway
 - NIC Teaming
 - System Center Operations Manager
 - System Center Virtual Machine Manager
 - Windows Server Gateway

Microsoft has implemented Software Defined Networking in Windows Server 2012, Windows Server 2012 R2, and Windows Server 2016 Hyper-V by providing the following components:

- Network Controller. Provides centralized management, configuration, monitoring, and troubleshooting of both your virtual and physical network infrastructure.



Note: Network Controller is a new feature in Windows Server 2016.

- Hyper-V Network Virtualization (HNV). Helps you abstract your applications and workloads from the underlying physical network by using virtual networks.

- Hyper-V Virtual Switch. Gives you the ability to connect virtual machines to both virtual networks and your physical network. Hyper-V Virtual Switch also provides security, isolation, and service-level policy enforcement.
- RRAS Multitenant Gateway. Gives you the ability to extend your network boundaries to Microsoft Azure or another provider to deliver on-demand hybrid infrastructure.
- NIC Teaming. Gives you the ability to configure multiple network adapters as a team for bandwidth aggregation and traffic failover to guard against loss of connectivity following a network component failure.

You can integrate Microsoft System Center with SDN to extend your SDN capabilities.



Note: System Center is a powerful enterprise datacenter management system that you can use to monitor, provision, configure, automate, and maintain your IT infrastructure.

Microsoft System Center provides a number of SDN technologies in the following components:

- System Center Operations Manager. Provides infrastructure monitoring for your datacenter and both the private and public cloud.
- System Center Virtual Machine Manager. Gives you the ability to provision and manage virtual networks. Provides for central control of virtual network policies that link to your applications or workloads.
- Windows Server Gateway. A virtual software router and gateway that allows you to route datacenter and cloud traffic between your virtual and physical networks.

These components are discussed in more detail throughout this module.

Benefits of SDN

Many customers with extensive network infrastructure face several common problems, which SDN can help to solve. The four main challenges are:

- Resources are finite. Finding the tools and resources to address the needs of business groups is difficult. This results in the IT department becoming a bottleneck to organizational growth.
- Resources are inflexible. After you address a business need by deploying IT infrastructure components and services, it is difficult to shift it around to address other needs.
- Mistakes are expensive. If the infrastructure fails to deliver, then the cost to the business can be huge.
- Networks are not always secure. The more software and hardware you have to address your business needs, the greater is the security risk. Managing the security of a distributed and disparate network infrastructure can be difficult.

• Challenges faced by many IT departments today include:

- Resources are finite
- Resources are inflexible
- Mistakes are expensive
- Networks are not always secure

• SDN overcomes these challenges and enables you to be:

- Flexible
- Efficient
- Scalable

SDN enables you to take advantage of a cloud-based infrastructure to overcome the limitations in your on-premises infrastructure, regardless of whether those limitations are short-term or persistent. This facility enables you to be:

- Flexible. You can move traffic from your on-premises infrastructure to your private or public cloud infrastructure.
- Efficient. You can abstract the hardware components of your network infrastructure with software components.
- Scalable. Your on-premises infrastructure has a finite capacity. Your cloud-based infrastructure has far broader limits that let you scale up your infrastructure when needed.

Planning for Software Defined Networking

Requirements

Before you can deploy a Software Defined Network, you must ensure that your network infrastructure meets the following prerequisites. These prerequisites fall into two categories:

- Physical network. You must be able to access all of your physical networking components. These will include:
 - Virtual LANs (VLANs)
 - Routers
 - Border Gateway Protocol (BGP) devices
 - Data Center Bridging with Enhanced Transmission Selection if using a Remote Direct Memory Access (RDMA) technology
 - Data Center Bridging with Priority-based Flow Control if using an RDMA technology that is based in RDMA over Converged Ethernet (RoCE)
- Physical compute hosts. These computers run the Hyper-V role and host the SDN infrastructure and tenant virtual machines. These hosts must:
 - Have Windows Server 2016 installed.
 - Have the Hyper-V role enabled.
 - Have an external Hyper-V Virtual Switch created with at least one physical adapter.
 - Be reachable with a Management IP address assigned to the Management Host virtual NIC (vNIC).

You must plan the following aspects of your Software Defined Networking configuration:

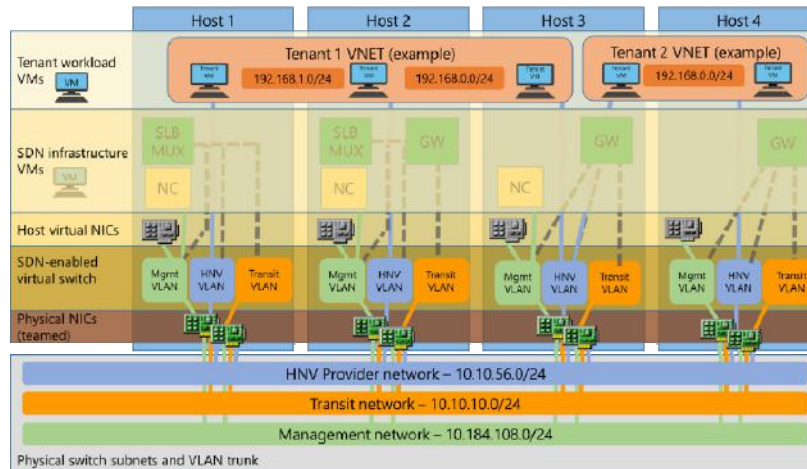
- Management and HNV Provider logical networks
- Logical networks for gateways and the software load balancer
- Logical networks required for RDMA-based storage
- Routing infrastructure
- Default gateways
- Network hardware



Software Defined Networking configuration

After you have ensured that your infrastructure meets these requirements, you must plan your SDN configuration. The components of a typical SDN deployment are shown in the following diagram.

A sample Software Defined Network architecture showing four Hyper-V hosts and two tenants:



A typical SDN deployment consists of the following components:

- Management and HNV Provider logical networks. All physical compute hosts must be able to access the Management logical network and the HNV Provider logical network.

Each physical compute host must be assigned at least one IP address from the Management logical network. You can use Dynamic Host Configuration Protocol (DHCP) for this assignment or else manually assign static IP configurations.



Note: The Management logical network is used by compute hosts to communicate with one another. All physical compute hosts need to have access to the Management logical network. All compute hosts must be reachable by using a Management IP address assigned to the Management Host vNIC.

- Logical networks for gateways and the software load balancer. You must create and provision additional logical networks for gateway and Software Load Balancing (SLB) usage. These include:
 - Transit logical network. Used by the RAS Gateway and SLB multiplexer (MUX) to exchange BGP peering information and North/South (external-internal) tenant traffic.



Note: Only physical compute hosts that run HNV Gateway or SLB MUX virtual machines must have connectivity to the Transit logical network subnet.

- Public virtual IP address (VIP) logical network. Required to have IP subnet prefixes that are Internet-routable outside of the cloud environment. These are the front-end IP addresses that external clients use to access resources in the virtual networks.
- Private VIP logical network. Used for VIPs that are only accessed from internal cloud clients, such as Generic Route Encapsulation (GRE) gateways or private services, and therefore do not need to be routable outside of the cloud.

- GRE VIP logical network. Exists solely for defining VIPs that are assigned to gateway virtual machines running on your SDN fabric for a server-to-server (S2S) GRE connection type.
- Logical networks required for RDMA-based storage. If you are using RDMA-based storage, then you must define a VLAN and a subnet for each physical adapter in your compute and storage hosts.
- Routing infrastructure. Routing information for the VIP subnets is advertised by the SLB MUX and HNV Gateways into the physical network by using internal BGP peering. You must create a BGP peer on the router that your SDN infrastructure uses to receive routes for the VIP logical networks advertised by the SLB MUXs and HNV Gateways. Typically, you configure BGP peering in a managed switch or router as part of the network infrastructure.
- Default gateways. You must configure only one default gateway on computers that are configured to connect to several networks, such as the physical compute hosts and gateway virtual machines. You usually configure the default gateway on the adapter that is used to reach all the way to the Internet.
- Network hardware. Your network hardware has a number of requirements, including those for network interface cards, switches, link control, availability and redundancy, and monitoring.



Note: For more information, refer to: “Plan a Software Defined Network Infrastructure” at: <http://aka.ms/N4y63g>

Deploying Software Defined Networking by using scripts

After you have planned your Software Defined Networking and configured your compute hosts, you can deploy a Software Defined Network. You can do this by using Virtual Machine Manager or by using scripts. This topic describes the process of using scripts to deploy Software Defined Networks.



Note: For more information on how to deploy a Software Defined Network by using Virtual Machine Manager, refer to: <http://aka.ms/Jnv4ko>

Use the following high-level procedure to deploy Software Defined Networking:

1. Install host networking and validate the configuration
2. Run SDN Express scripts and validate setup
3. Deploy a sample tenant workload and validate deployment

Use the following high-level procedure to deploy a Software Defined Network:

1. Install host networking and validate the configuration.
2. Run SDN Express scripts and validate setup.
3. Deploy a sample tenant workload and validate deployment.

Install host networking and validate the configuration

To install host networking, complete the following tasks:

1. On your compute hosts:
 - a. Install the latest network drivers for all NICs.
 - b. Add the **Hyper-V** role.
 - c. Create the **Hyper-V Virtual Switch**.

2. Obtain the VLAN ID of your Management VLAN and attach the Management vNIC of the newly created virtual switch to the Management VLAN.



Note: The Management VLAN ID is determined during the planning phase.

3. Assign a valid IP configuration to the Management vNIC of the newly created virtual switch.



Note: The decision whether to use DHCP or static configuration is made during the planning phase.

4. Deploy a virtual machine to host the Active Directory Domain Services (AD DS) and Domain Name System (DNS) roles and then join your Hyper-V hosts to this AD DS domain.

To validate host networking setup, complete the following tasks:

1. Ensure that the virtual switch was created correctly by using the **Get-VMSwitch "Switch_name"** cmdlet.
2. Verify that the Management vNIC on the virtual switch is connected to the VLAN by using the **Get-VMNetworkAdapterIsolation -ManagementOS** cmdlet.
3. Verify that all Hyper-V hosts are accessible by testing connectivity to their Management IP address and the fully qualified domain name (FQDN).
4. Ensure that the Kerberos credentials that are used provide access to all servers:
 - o To do this, at a command prompt, run the **winrm id -r:<Hyper-V Host FQDN>** command.

Run SDN Express scripts and validate setup

To set up Software Defined Networking by using SDN Express scripts, complete the following tasks:

1. Download the required scripts.



Note: You can download the scripts from the Microsoft SDN GitHub Repository at: <http://aka.ms/G7us9e>

2. Set up your deployment computer:
 - a. Install Windows Server 2016 on the deployment computer.
 - b. Extract the scripts and copy the **SDNExpress** folder from the extract to the root of drive C on the deployment computer.
 - c. Verify that the **SDNExpress** folder contains the following subfolders:
 - **AgentConf.** This subfolder stores copies of schemas used by the SDN Host Agent on each Windows Server 2016 Hyper-V host to program network policy.
 - **Certs.** This subfolder is the temporary location for certificate files.
 - **Images.** You use this subfolder to store your Windows Server 2016 vhd/x image file.
 - **Tools.** This subfolder includes utilities and tools for troubleshooting.
 - **TenantApps.** This subfolder is used to deploy tenant workloads.

▪ **Scripts:**

- **SDNExpress.ps1.** This script deploys and configures the SDN fabric, including the Network Controller virtual machines, SLB/MUX virtual machines, gateway pools, and the HNV gateway virtual machines corresponding to the pools.
- **FabricConfig.psd1.** This script is a configuration file template for the SDNExpress script. You customize this for your environment.
- **SDNExpressTenant.ps1.** This script deploys a sample tenant workload on a virtual network with a load-balanced VIP. You can use this script with an **Undo** option to delete the corresponding configuration.
- **TenantConfig.psd1.** This script is a template configuration file for tenant workload and S2S gateway configuration.
- **SDNExpressUndo.ps1.** This script cleans up the fabric environment and resets it to a starting state.
- **SDNExpressEnterpriseExample.ps1.** This script provisions one or more enterprise site environments. You can use this script with an **Undo** option to delete the corresponding configuration.
- **EnterpriseConfig.psd1.** This script is a template configuration file.

3. Share the **C:\SDNExpress** folder.
4. Edit and configure the **SDNExpress\scripts\FabricConfig.psd1** script file by changing the **<< Replace >>** tags with specific values to fit your infrastructure including:
 - Host names
 - Domain names
 - Usernames and passwords
 - Network information for your networks
5. In DNS, create a Host A record for:
 - The NetworkControllerRestName (FQDN)
 - The NetworkControllerRestIP
6. Run the following script as a Domain Admin:

```
SDNExpress\scripts\SDNExpress.ps1 -ConfigurationDataFile FabricConfig.psd1 -Verbose
```



Note: If you have to roll back the configuration, run the following command:

```
SDNExpress\scripts\SDNExpressUndo.ps1 -ConfigurationDataFile FabricConfig.psd1 -Verbose
```

If the script ran without errors, you can proceed to validate the setup. Complete the following procedure to validate your Software Defined Networking setup:

1. Ensure that the Network Controller Host Agent and SLB Host Agent are running on all Hyper-V hosts by using the **Get-Service NCHostAgent** and **Get-Service SlbHostAgent** cmdlets.
2. Verify network connectivity on the Management logical network between all Network Controller node virtual machines and Hyper-V hosts.

3. Use **Netstat.exe** to check that the Network Controller Host Agent is connected to the Network Controller on TCP:6640.
4. Verify that the Dynamic IPs associated with all Hyper-V hosts that are hosting load-balanced tenant workload virtual machines have Layer-3 IP connectivity to the SLB Manager VIP address.
5. Use diagnostic tools to ensure that there are no errors on any fabric resources in the Network Controller. For example, use the **Debug-NetworkControllerConfigurationState** cmdlet.
6. Verify the BGP peering state to ensure that the SLB MUX is peered to the Top-of-Rack switch or RRAS virtual machine (the BGP peer). Run the following cmdlet from a Network Controller node virtual machine: **Debug-SlbConfigState**.

Deploy a sample tenant workload and validate deployment

After you have deployed SDN and verified the configuration, you can deploy a sample tenant workload.



Note: This sample tenant workload consists of two virtual subnets (a web tier and a database tier) that are protected with access control list rules by using the SDN distributed firewall. The web tier's virtual subnet is accessible through the SLB MUX by using a VIP address. The script automatically deploys two web tier virtual machines and one database tier virtual machine and connects these to the virtual subnets.

Validate your SDN deployment by performing the following steps:

1. Edit and configure the **SDNExpress\scripts\TenantConfig.psd1** file by changing the << **Replace** >> tags with specific values, including the:
 - o VHD image name
 - o Network controller representational state transfer (REST) name
 - o vSwitch Name
2. Run the following script:

```
SDNExpress\scripts\SDNExpressTenant.ps1 -ConfigurationDataFile TenantConfig.psd1 -
Verbose
```



Note: If you have to roll back the configuration, run the following command:

```
SDNExpress\scripts\SDNExpressTenant.ps1 -Undo -ConfigurationDataFile TenantConfig.psd1
-Verbose
```

3. Validate the tenant deployment by performing the following steps:
 - a. Sign in to the database tier virtual machine and verify network connectivity to the IP address of one of the web tier virtual machines.
 - b. Check the Network Controller tenant resources for any errors by running the following cmdlet from any Hyper-V host with Layer-3 connectivity to the Network Controller:

```
Get-NetworkControllerConfigurationState -NCIP <FQDN of Network Controller REST
Name>
```

- c. Validate that the policy has been received and persisted in the Network Controller Host Agent by running the following cmdlet:

```
ovsdb-client.exe dump tcp:127.0.0.1:6641 ms_vtep
```

- d. Check that an IP address has been assigned for a provider address (PA) Host vNIC and the Ethernet adapters for the PA Host vNIC by using the **ipconfig /allcompartments /all** command.
- e. Check PA connectivity between two hosts with a ping command. Obtain the compartment ID from the output of the previous command:

```
ping -c <compartment Id> <Remote Hyper-V Host PA IP Address>
```

Verify the correctness of the statement by placing a mark in the column to the right.

Statement	Answer
In Software Defined Networking, each physical compute host must be assigned at least one IP address from the Management logical network. You can use DHCP for this assignment.	

Question: Does the complexity of your organization's network infrastructure suggest the need for Software Defined Networking?

Lesson 2

Implementing network virtualization

Network virtualization is a part of Software Defined Networking in Windows Server 2016 with which you can create virtual networks that are isolated logically on the same physical network infrastructure. This lesson explores the features and technologies in network virtualization.

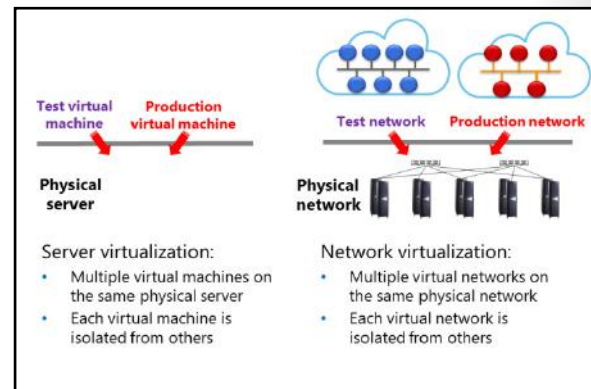
Lesson Objectives

After completing this lesson, you will be able to:

- Explain network virtualization.
- Identify the benefits of network virtualization.
- Describe Network Virtualization Generic Route Encapsulation (NVGRE).
- Explain network virtualization policies.

What is network virtualization?

Network virtualization provides functionality for managing network traffic that is similar to what server virtualization does for managing virtual machines. You can use server virtualization to run multiple virtual machines on the same physical server. The same is true for network virtualization. You can have multiple virtual networks, which are logically isolated, on the same physical network infrastructure. From each virtual network, it seems that only the virtual network is using the physical network infrastructure, even though multiple virtual networks could be using the same physical infrastructure at the same time.



Network virtualization is an implementation of SDN and provides a layer of abstraction over the physical network. To achieve this abstraction, the virtualization platform has to support it. The Hyper-V Virtual Switch in Windows Server 2016 supports network virtualization by using two IP addresses for each virtual machine. By using the two IP addresses, you can use network virtualization to keep the logical network topology, which is virtualized and separated from the actual underlying physical network topology, and addresses used on the physical network. Thus, you can run virtual machines and provide them with the same network access without any modification on any Hyper-V host, assuming that Hyper-V hosts are configured to map between both IP addresses.

Benefits of network virtualization

Network virtualization provides a layer of abstraction between the physical network and network traffic and thereby provides the following benefits:

- Flexible virtual machine placement. Network virtualization provides abstraction and separates IP addresses used in virtual machines from the IP addresses used on the physical network. This way, you can place a virtual machine on any Hyper-V host in the datacenter, and the IP address assignment or VLAN isolation restrictions of the physical network no longer restricts the placement.
- Multitenant network isolation without VLANs. You can define and enforce network traffic isolation without the use of VLANs or the need to reconfigure physical network switches. Because network virtualization uses a 24-bit identifier for virtual networks compared to a 12-bit identifier for VLANs, you are also not limited to 4094 VLAN IDs. Also, with network virtualization, no manual reconfiguration of physical hardware is required when you move existing virtual machines or create new ones.
- IP address reuse. Virtual machines in different virtual networks can use the same or overlapping IP address space, even when deploying those virtual machines on the same physical network. Virtual networks are isolated, and they can use the same address space without any conflict or issue.
- Live migration across subnets. Without network virtualization, virtual machine live migration is limited to the same IP subnet or VLAN, because when a virtual machine moves to different subnets, its IP address has to change to match the new network. With network virtualization, you can move virtual machines by using live migration between two Hyper-V hosts in different subnets without having to change the virtual machine IP address. With the use of network virtualization, the virtual machine location change updates and synchronizes among computers that have ongoing communication with the migrated virtual machine.
- Compatibility with existing network infrastructure. Network virtualization is compatible with existing network infrastructure, and you can deploy it in an existing datacenter. You do not need to redesign the physical network layer to implement network virtualization.
- Transparent moving of virtual machines to a shared infrastructure as a service (IaaS) cloud. With IaaS, the physical platform where the virtual machines run is hosted in a separate datacenter, usually accessible through the Internet. When network virtualization is used, IP addresses, IP policies, and virtual machine configurations remain unchanged, regardless of which Hyper-V host the virtual machine is running on. As a result, you can move virtual machines between Hyper-V hosts in your datacenter, between Hyper-V hosts in different datacenters, and between a Hyper-V host in your datacenter and the shared IaaS cloud.
- Support for resource metering. With Hyper-V in Windows Server 2016, you can enable resource metering. Resource metering provides information about the usage of host and network resources for individual virtual machines. You can use this information to charge the tenants for actual resource usage. You can enable network resource metering for virtual machines that use network virtualization.

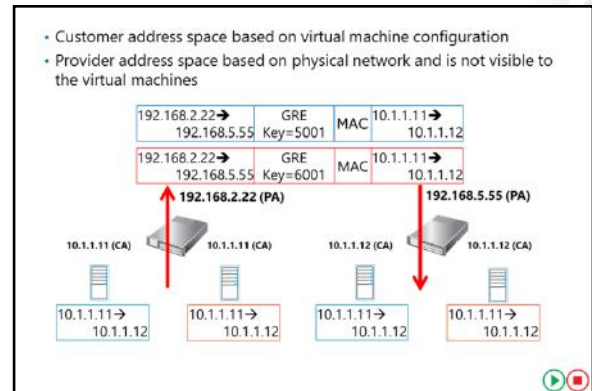
- Flexible virtual machine placement
- Multitenant network isolation without VLANs
- IP address reuse
- Live migration across subnets
- Compatibility with existing network infrastructure
- Transparent moving of virtual machines to a shared IaaS cloud
- Support for resource metering
- Configuration by using Windows PowerShell or by using System Center Virtual Machine Manager

- Configuration by Windows PowerShell. Network virtualization supports Windows PowerShell for configuring the network virtualization and isolation policies. The Hyper-V module includes cmdlets that you can use to configure, monitor, and troubleshoot network virtualization. Configuring network virtualization in Windows PowerShell is complex, so we strongly recommend that you use the Virtual Machine Manager to configure and manage network virtualization.

What is Generic Route Encapsulation?

Windows Server 2016 Hyper-V use *Network Virtualization Generic Route Encapsulation* (NVGRE) to implement network virtualization. When network virtualization is used, each virtual network adapter is associated with two IP addresses. Those two addresses are:

- Customer address (CA). This is the IP address that is configured and used by the virtual machine. This address is configured in the properties of the virtual network adapter in the virtual machine guest operating system, regardless of whether network virtualization is used. A virtual machine uses the CA when communicating with another system, and if you migrate a virtual machine to a different Hyper-V host, the CA can remain the same.
- Provider address (PA). This is the IP address that the virtualization platform assigns to the Hyper-V host and is dependent on the physical network infrastructure where the Hyper-V host is connected. When network virtualization is used and the virtual machine sends network traffic, the Hyper-V host encapsulates the packets and includes the PA as the source address from where packets were sent. The PA is visible on the physical network but not visible to the virtual machine. If you migrate a virtual machine to a different Hyper-V host, the PA changes.



Using NVGRE

When a virtual machine has to communicate over a network and you have configured network virtualization, NVGRE is used to encapsulate its packets. For example, assume that one virtual machine is configured with IP address 10.1.1.11 (CA 1) and is running on a Hyper-V host that uses IP address 192.168.2.22 (PA 1). The second virtual machine is configured with IP address 10.1.1.12 (CA 2) and is running on Hyper-V host with IP address 192.168.5.55 (PA 2). If network virtualization is used, the first Hyper-V host will use NVGRE to encapsulate the virtual machine packets, which contain the source (CA 1) and the destination IP address (CA 2), into the envelope. This envelope uses its own IP address (PA 1) as the source address and the IP address of the Hyper-V host on which the second virtual machine is running (PA 2) as the destination address. Encapsulated packages will be sent on the physical network between the two Hyper-V hosts. The destination Hyper-V host (PA 2) will extract the envelope from the encapsulated packet and pass it to the destination virtual machine (CA 2), which is running on that Hyper-V host.

With NVGRE, you can configure virtual machines with the same IP addresses and deploy them on the same or different host machines. To address this scenario, the GRE envelope header includes a field named Key, which represents a Virtual Subnet ID. When implementing network virtualization, you define a Virtual Subnet ID on the Hyper-V host for each network used by the hosted virtual machines. The Virtual Subnet ID is used to separate and isolate traffic between different virtual networks and enables Hyper-V host to pass the traffic only to virtual machines on the same virtual network.

What are network virtualization policies?

If you configure network virtualization and if two virtual machines have to communicate, the Hyper-V host that the first virtual machine is running on must be aware which Hyper-V host the second virtual machine is running on before the host can encapsulate network packets into GRE envelopes. If both virtual machines are running on the same Hyper-V host, Hyper-V already has this information. But virtual machines usually run on different Hyper-V hosts, so you must configure network virtualization to ensure the virtual machines can communicate. You

configure network virtualization by deploying network virtualization policies. Network virtualization policies define the mappings between IP address spaces used by the virtual machines (CA spaces) and the IP addresses of the Hyper-V hosts that those virtual machines are running on (PA spaces). Before sending traffic on the physical network, the Hyper-V hosts consult the network virtualization policies to determine on which Hyper-V host the target virtual machine is running and encapsulate the traffic with a GRE envelope. The encapsulated traffic is sent on the physical network only after that determination.

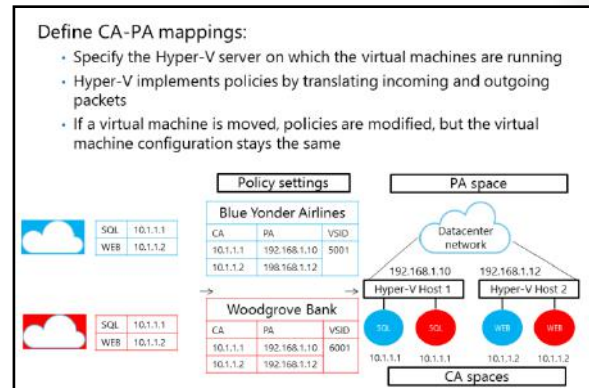
For example, assume that you are hosting two companies, Blue Yonder Airlines and Woodgrove Bank, with the following configuration:

- Blue Yonder Airlines is running Microsoft SQL Server data management software in a virtual machine with the IP address 10.1.1.1 and a web server in a virtual machine with the IP address 10.1.1.2. The web server is using SQL Server as a database for storing transactions.
- Woodgrove Bank is running SQL Server in a virtual machine configured with the same IP address 10.1.1.1 and a web server in a virtual machine with the IP address 10.1.1.2. The web server is using SQL Server as a database for storing transactions.
- The computers running SQL Server for both companies are running on Hyper-V Host 1, which has the IP address 192.168.1.10. Web servers for both companies are running on Hyper-V Host 2, which has the IP address 192.168.1.12.

This means that the virtual machines have the following CAs and PAs.

Organization	CAs	PAs
Blue Yonder Airlines	SQL is 10.1.1.1; WEB is 10.1.1.2	SQL is 192.168.1.10; WEB is 192.168.1.12
Woodgrove Bank	SQL is 10.1.1.1; WEB is 10.1.1.2	SQL is 192.168.1.10; WEB is 192.168.1.12

To enable communication between the virtual machines, you need to configure a virtual network. For example, you could configure a virtual network for Blue Yonder Airlines with the Virtual Subnet ID 5001 and configure a virtual network for Woodgrove Bank with the Virtual Subnet ID 6001. You also create network virtualization policies for both companies and apply policies to Hyper-V Host 1 and Hyper-V Host 2.



When the Blue Yonder Airlines WEB virtual machine on Hyper-V Host 2 queries its SQL Server at 10.1.1.11, the following process happens:

1. Hyper-V Host 2, based on its policy settings, translates the addresses in the packet from the following:
 - Source: 10.1.1.2 (the CA of Blue Yonder Airlines WEB)
 - Destination: 10.1.1.1 (the CA of Blue Yonder Airlines SQL)
2. The addresses are translated into the encapsulated packet that contains the following:
 - GRE header with Virtual Subnet ID: 5001
 - Source: 192.168.2.12 (the PA for Blue Yonder Airlines WEB)
 - Destination: 192.168.1.10 (the PA for Blue Yonder Airlines SQL)



Note: The encapsulated packet also contains the original packet.

When Hyper-V Host 1 receives the packet, based on its policy settings, it will decapsulate the NVGRE packet, determine that it is for the Blue Yonder Airlines virtual network (Virtual Subnet ID 5001), and pass it to the virtual machine with IP 10.1.1.1, as specified in the original (encapsulated) packet.



Note: You can configure network virtualization policies by using Windows PowerShell. It is easier to configure network virtualization policies with tools such as System Center Virtual Machine Manager.

You can use network virtualization and network virtualization policies to move virtual machines between Hyper-V hosts and preserve their network configuration. When you move a virtual machine, you need to update only the network virtualization policies to reflect the new Hyper-V host on which the virtual machine is running; the virtual machine network configuration stays the same and is still connected to the same virtual network.

Question: Does a virtual machine CA change when you move the virtual machine between Hyper-V hosts?

Question: Why are network virtualization policies necessary when using network virtualization?

Lesson 3

Implementing Network Controller

Network Controller, a new feature of Windows Server 2016, gives you the ability to manage, configure, monitor, and troubleshoot virtual and physical network infrastructure in your datacenter by using a centralized, programmable point of automation. Using Network Controller, you can automate the configuration of your network infrastructure without needing to perform manual configuration of network devices and services.


Lesson Objectives

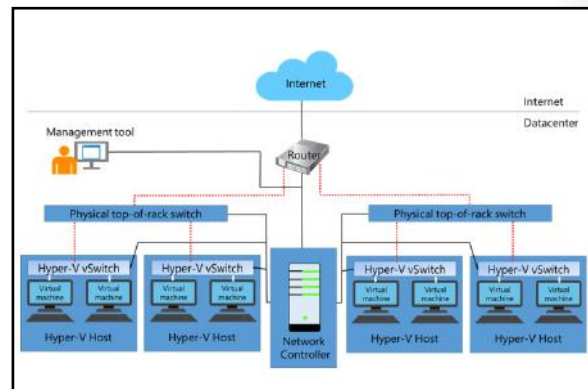
After completing this lesson, you will be able to:

- Describe Network Controller.
- List the requirements for installing Network Controller.
- Prepare to deploy Network Controller.
- Describe the procedure for deploying Network Controller.
- Describe the role of Network Controller in Datacenter Firewall.
- Describe the role of Network Controller in SLB.
- Describe the role of Network Controller in RAS Gateway.
- Deploy Network Controller.

What is Network Controller?

Network Controller is a Windows Server 2016 server role. It provides two application programming interfaces (APIs): the Southbound API and the Northbound API. The first of these enables Network Controller to communicate with the network, while the second API gives you the ability to communicate with Network Controller.

 **Note:** You can deploy Network Controller in AD DS domain and non-domain environments. In AD DS domain environments, Network Controller authenticates users and devices with Kerberos, while in non-domain environments, you must deploy digital certificates to provide authentication.



Southbound API

Network Controller uses the Southbound API to communicate with network devices, services, and components. With the Southbound API, Network Controller can:

- Discover network devices.
- Detect service configurations.

- Gather all of the information you need about the network.
- Send information to the network infrastructure, for example, configuration changes that you have made.

Northbound API

The Network Controller Northbound API provides you with the ability to gather network information from Network Controller with which you can monitor and configure the network. The Network Controller Northbound API enables you to configure, monitor, troubleshoot, and deploy new devices on the network by using:

- Windows PowerShell
- REST API
- A management application with a graphical user interface, for example, System Center Virtual Machine Manager or System Center Operations Manager

Management with Network Controller

You can use Network Controller to manage the following physical and virtual network infrastructure:

- Hyper-V virtual machines and virtual switches
- Datacenter Firewall
- RAS Multitenant Gateways, Virtual Gateways, and gateway pools
- Load balancers

Network Controller provides a number of features with which you can configure and manage both virtual and physical network devices and services. These are:

- Firewall management. You can configure and manage firewall Access Control rules for your workload virtual machines.
- SLB management. You can configure multiple servers to host the same workload, helping to provide high availability and scalability.
- Virtual network management. You can deploy and configure HNV, including Hyper-V Virtual Switch, virtual network adapters on individual virtual machines, and virtual network policies storage and distribution.
- RAS gateway management. You can provide gateway services to your tenants by deploying, configuring, and managing Hyper-V hosts and virtual machines that are members of a RAS gateway pool.

Requirements for deploying Network Controller

You can deploy Network Controller on one or more computers, one or more virtual machines, or a combination of both. Because Network Controller is a Windows Server 2016 server role, the requirements are not complex. They are as follows:


- You can only deploy Network Controller to the Windows Server 2016 Datacenter edition.
- The management client you use must be installed on a computer or virtual machine running Windows 10, Windows 8.1, or Windows 8.
- You must configure dynamic DNS registration to enable registration of required DNS records for Network Controller.
- If the computers or virtual machines running Network Controller or the management client for Network Controller are joined to a domain, you must:
 - Create a security group that holds all the users that have permission to configure Network Controller.
 - Create a security group that holds all of the users who have permission to configure and manage the network by using Network Controller.

- You can deploy Network Controller only to the Windows Server 2016 Datacenter edition
- The management client must be running Windows 10, Windows 8.1, or Windows 8
- You must configure dynamic DNS registration for Network Controller
- If virtual machines running Network Controller are joined to a domain, you must create appropriate AD DS security groups
- If virtual machines running Network Controller are not joined to a domain, you must configure certificate-based authentication




Note: In both these instances, all users added to either of these groups must also belong to the Domain Users group.

- If the computers or virtual machines running Network Controller or the management client for Network Controller are not joined to a domain, you must configure certificate-based authentication by:
 - Creating a certificate for use on the management client. The Network Controller must trust this certificate.
 - Creating a certificate on the Network Controller for computer authentication. The certificate must meet the following requirements:
 - The certificate subject name must match the DNS name of the computer or virtual machine holding the Network Controller role.
 - The server authentication purpose is present in enhanced key usage (EKU) extensions.
 - The certificate subject name should resolve to one of the following addresses:
 - The IP address of the Network Controller, if Network Controller is deployed on a single computer or virtual machine.
 - The REST IP address, if Network Controller is deployed on multiple computers, multiple virtual machines, or both.
 - The certificate must be trusted by all the REST clients.
 - The certificate must be trusted by the SLB MUX and the southbound host computers that Network Controller manages.

 **Note:** A certification authority can enroll the certificate, or the certificate can be self-signed. We do not recommend self-signed certificates for production deployments, but they are acceptable for test lab environments.

- Enrolling this certificate on the Network Controller.

 **Note:** The same certificate must be provisioned on all the Network Controller nodes. After creating the certificate on one node, you can export the certificate (with private key) and import it on the other nodes.

Demonstration: Preparing to deploy Network Controller

In this demonstration, you will see how to:

- Create AD DS security groups.
- Request a certificate.

Demonstration Steps

Create AD DS security groups

1. On **LON-DC1**, open **Active Directory Users and Computers**.
2. Create the following global security groups:
 - a. **Network Controller Admins**
 - b. **Network Controller Ops**
3. Add **Beth Burke** and **Administrator** to both these groups.

These security groups are required for users that will administer Network Controller and for users that will use Network Controller to administer the network devices and services.

Request a certificate

1. On **LON-SVR2**, open the management console, and then add the **Certificates** snap-in with the focus on the local computer.
2. Request a **Computer** certificate.
3. Close the management console without saving changes.

This certificate is required for encryption of communication between the Network Controller and the management clients.

The procedure for deploying Network Controller

You can deploy the Network Controller role by using Windows PowerShell by following these high-level steps:

1. Install the Network Controller server role.
2. Configure the Network Controller cluster.
3. Configure the Network Controller application.
4. Validate the Network Controller deployment.

1. Install the Network Controller server role
2. Configure the Network Controller cluster
3. Configure the Network Controller application
4. Validate the Network Controller deployment

Install the Network Controller server role


To install the Network Controller server role, on the server computer or virtual machine that will host the role, open Windows PowerShell (Admin) and then run the following cmdlet:

```
Install-WindowsFeature -Name NetworkController -IncludeManagementTools
```

After performing this task, restart your computer or virtual machine.

Configure the Network Controller cluster

The Network Controller cluster provides scalability and high-availability for the Network Controller application. To configure the cluster, sign in as a local administrator on the computer or virtual machine where you want to configure the cluster.

 **Note:** If the computer or virtual machine on which you deployed the Network Controller role is a domain member, the user account you use to sign in with must also belong to Domain Users.


To configure the cluster, complete the following steps:

1. Create a node object. You must create a node object for each computer or virtual machine that is a member of the Network Controller cluster. Use the **New-NetworkControllerNodeObject** cmdlet to complete this step. For example, the following command creates a Network Controller node object named Node1. The FQDN of the computer is NCNode1.Adatum.com, and Ethernet is the name of the interface on the computer listening to REST requests.

```
New-NetworkControllerNodeObject -Name "Node1" -Server "NCNode1.Adatum.com" -
FaultDomain "fd:/rack1/host1" -RestInterface "Ethernet"
```

2. Configure the cluster. After you have created the node(s) for the cluster, use the **Install-NetworkControllerCluster** cmdlet to configure the cluster. For example, the following commands install a Network Controller cluster in a test lab. High-availability support is not available because a single node is used. Kerberos authentication is used between the cluster nodes.

```
$NodeObject = New-NetworkControllerNodeObject -Name "Node1" -Server
"NCNode1.Adatum.com" -FaultDomain "fd:/rack1/host1" -RestInterface "Ethernet"
Install-NetworkControllerCluster -Node $NodeObject -ClusterAuthentication Kerberos
```

 **Additional Reading:** For more information on the syntax of these cmdlets, refer to: <http://aka.ms/A6mp6v>

Configure the Network Controller application

The last deployment step involves the configuration of the Network Controller application. Use the **Install-NetworkController** cmdlet to complete this procedure. For example, the following code creates a Network Controller node object, and then stores it in the `$NodeObject` variable.

```
$NodeObject = New-NetworkControllerNodeObject -Name "Node01" -Server "NCNode11" -
FaultDomain "fd:/rack1/host1" -RestInterface Ethernet
```

The following command gets a certificate named `NCEncryption`, and then stores it in the `$Certificate` variable.

```
$Certificate = Get-Item Cert:\LocalMachine\My | Get-ChildItem | where {$_.Subject -imatch
"NCEncryption" }
```

The following command creates a Network Controller cluster by using the **Install-NetworkControllerCluster** cmdlet.

```
Install-NetworkControllerCluster -Node $NodeObject -ClusterAuthentication None
```

The following command deploys the Network Controller in a test environment. Because a single node is used in the deployment, there is no high-availability support. This Network Controller employs no authentication between the cluster nodes, nor between the REST clients and Network Controller. The command specifies the `$Certificate` to encrypt the traffic between the REST clients and Network Controller.

```
Install-NetworkController -Node $NodeObject -ClientAuthentication None -RestIpAddress
"10.0.0.1/24" -ServerCertificate $Certificate
```



Additional Reading: For more information on the syntax of this cmdlet, refer to: <http://aka.ms/bkmtmqo>

Validate the Network Controller deployment

After you have deployed the Network Controller, you can validate the deployment by adding a credential to the Network Controller and then retrieving the credential.



Note: If you are using Kerberos as the `ClientAuthentication` mechanism—that is, if the computers or virtual machines are members of a domain—then membership in the `ClientSecurityGroup` that you created is the minimum required to perform this procedure. You define the `ClientSecurityGroup` when you use the **Install-NetworkController** cmdlet.

Complete this task by performing the following steps:

1. Open Windows PowerShell (Admin), and then run the following commands:

```
$cred=New-Object Microsoft.Windows.Networkcontroller.credentialproperties
$cred.type="usernamepassword"
$cred.username="admin"
$cred.value="abcd"
New-NetworkControllerCredential -ConnectionUri https://networkcontroller -Properties
$cred -ResourceId cred1
```

- To retrieve the credential that you added to Network Controller, run the following command:

```
Get-NetworkControllerCredential -ConnectionUri https://networkcontroller -ResourceId cred1
```


- If everything works, you should receive output similar to the following:

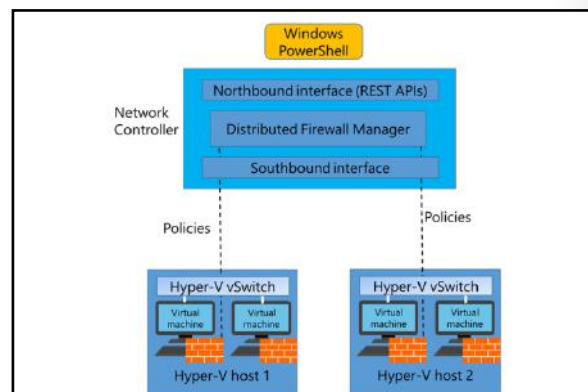
```
Tags                :
ResourceRef         : /credentials/cred1
CreatedTime         : 1/1/0001 12:00:00 AM
InstanceId          : e16ffe62-a701-4d31-915e-7234d4bc5a18
Etag                : W/"1ec59631-607f-4d3e-ac78-94b0822f3a9d"
ResourceMetadata    :
ResourceId          : cred1
Properties           : Microsoft.Windows.NetworkController.CredentialProperties
```

You have successfully verified the deployment.

Datacenter Firewall

Datacenter Firewall in Windows Server 2016 helps you install and configure firewall policies to protect your virtual networks from unwanted network traffic. You manage the Datacenter Firewall policies by using Network Controller Northbound APIs.

 **Note:** Both the cloud service provider admin and the tenant admin can manage Datacenter Firewall policies by using the Network Controller.



Benefits for cloud provider

For cloud service providers, the Datacenter Firewall provides these benefits:

- A software-based firewall solution that is highly scalable and manageable and that can easily be offered to tenants.
- The ability to easily move tenant virtual machines to different compute hosts without disrupting tenant firewall configuration because:
 - It is deployed as a vSwitch port host agent firewall.
 - Tenant virtual machines get the policies assigned to their vSwitch host agent firewall.
 - Firewall rules are configured in each vSwitch port, independent of the host that runs the virtual machine.
- Protection to tenant virtual machines regardless of the tenant guest operating system.

Benefits for tenants

For tenants, the Datacenter Firewall provides the ability to:

- Define firewall rules that can help protect Internet-facing workloads on their virtual networks.
- Define firewall rules that can help protect traffic between virtual machines on the same L2 virtual subnet and also between virtual machines on different L2 virtual subnets.
- Define firewall rules that can help protect and isolate network traffic between tenant on-premises networks and their virtual networks at the service provider.

Software Load Balancing

You can use Software Load Balancing (SLB) in SDN to distribute network traffic across available network resources. Windows Server SLB provides the following features:

- Layer 4 load balancing for both "North-South" and "East-West" Transmission Control Protocol/User Datagram Protocol (TCP/UDP) traffic
- Public and internal network traffic load balancing
- Support for Dynamic IP addresses on VLANs and on Hyper-V virtual networks
- Support for health probe

Windows Server SLB provides the following features:

- Layer 4 load balancing for both "North-South" and "East-West" TCP/UDP traffic
- Public and internal network traffic load balancing
- Support for dynamic IP addresses on VLANs and on Hyper-V virtual networks
- Support for health probe

SLB maps VIPs to dynamic IP addresses that are part of a set of resources in the cloud. In this scenario, VIPs are single IP addresses that map to a pool of available virtual machines. Dynamic IP addresses are assigned to tenant resources within the cloud infrastructure.



Note: VIPs are IP addresses available on the Internet for tenants and tenant customers to connect to tenant resources in the cloud datacenter. Dynamic IP addresses are the IP addresses of the virtual machines that are members of a load-balanced pool.

SLB infrastructure

The SLB infrastructure consists of the following components:

- Virtual Machine Manager. You use Virtual Machine Manager to configure Network Controller, including Health Monitor and SLB Manager.



Additional Reading: You also can use Windows PowerShell cmdlets. For more information on the Windows PowerShell cmdlets that you can use to manage Network Controller, refer to: <http://aka.ms/Wct3o3>

- Network Controller. Before you can deploy SLB in Windows Server 2016, you must first deploy Network Controller. Network Controller performs the following functions in SLB:
 - Processes SLB commands that arrive via the Northbound API from Virtual Machine Manager, Windows PowerShell, or other network management application.
 - Calculates policy for distribution to Hyper-V hosts and SLB MUXs.
 - Provides the health status of the SLB infrastructure.
 - Provides each MUX with each VIP.
 - Configures and controls the behavior of the VIP to dynamic IP mapping in the MUX.



Note: Specifically, you define load balancing policies by using Network Controller, and the MUX maps VIPs to the correct dynamic IP addresses by using these policies. These load balancing policies include Protocol, Front-end port, Back-end port, and distribution algorithm (5-, 3-, or 2-tuples).

- SLB MUX. When network inbound Internet traffic arrives, the SLB MUX maps and rewrites the traffic so that it will arrive at an individual dynamic IP. This is based on an examination of the traffic by the MUX for the destination VIP. Within the SLB infrastructure, the MUX:
 - Holds the VIPs.
 - Uses BGP to advertise each of the VIPs to routers on the physical network.
 - Consists of one or more virtual machines.
- Hosts that run Hyper-V. You use SLB with computers that are running Windows Server 2016 and Hyper-V.
- SLB Host Agent. The SLB Host Agent:
 - Listens for SLB policy updates from Network Controller.
 - Programs rules for SLB into the Software Defined Networking-enabled Hyper-V virtual switches that are configured on the local computer.



Note: When you deploy SLB, you must deploy the SLB Host Agent on every Hyper-V host computer. You can install this agent on all versions of Windows Server 2016 that support the Hyper-V role, including Nano Server.

- Software Defined Networking-enabled Hyper-V Virtual Switch. For a virtual switch to be compatible with SLB, you must use Hyper-V Virtual Switch Manager or Windows PowerShell commands to create the switch, and then you must enable Virtual Filtering Platform for the virtual switch. The virtual switch performs the following actions for SLB:
 - Processes the data path for SLB.
 - Receives inbound network traffic from the MUX.
 - Bypasses the MUX for outbound network traffic, sending it to the router by using direct server return (DSR).
 - Runs on Nano Server instances of Hyper-V.

- BGP-enabled router. BGP allows the routers to:
 - Route inbound traffic to the MUX by using equal-cost multi-path routing (ECMP).
 - For outbound network traffic, use the route provided by the host.
 - Listen for route updates for VIPs from SLB MUX.
 - Remove SLB MUXs from the SLB rotation if Keep Alive fails.

RAS Gateway

When you implement network virtualization by using the Hyper-V Virtual Switch, the switch operates as a router between different Hyper-V hosts in the same infrastructure. The network virtualization policies define how packets will be routed from one host to another. However, a virtual switch cannot route to networks outside the Hyper-V server infrastructure when using network virtualization. If you were not using network virtualization, you would just connect the virtual machine to an external switch, and the virtual machine could connect to the same networks as the host machine.

RAS Gateway provides the following features:

- Site-to-site VPN
- Point-to-site VPN
- GRE tunneling
- Dynamic routing with BGP

Use RAS Gateway in the following scenarios:

- Multitenant-aware VPN gateway
- Multitenant-aware NAT gateway
- Forwarding gateway for internal physical network access



But in a network virtualization scenario, you might have multiple virtual machines running on a Hyper-V host that share the same IP addresses. You might also want to move the virtual machine to any host in the network without disrupting network connectivity. You must 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 correctly routed to the internal addresses that the virtual machines use. Windows Server 2016 provides the RAS Gateway to address these issues.



Note: RAS Gateway is referred to as Windows Server Gateway in System Center.

Overview of RAS Gateway

RAS Gateway is a software-based, multitenant, BGP-capable router. It is designed for cloud service providers and large organizations that host multiple tenant virtual networks using Hyper-V Network Virtualization (HNV). RAS Gateway provides the following features:

- Site-to-site VPN. Gives you the ability to connect two networks at different physical locations across the Internet with a site-to-site VPN connection.
- Point-to-site VPN. Gives organization employees or administrators the ability to connect to your organization's network from remote locations.
- GRE tunneling. Enables connectivity between tenant virtual networks and external networks.
- Dynamic routing with BGP. Reduces the need for manual route configuration on routers because it is a dynamic routing protocol, and automatically learns routes between sites that are connected by using site-to-site VPN connections.

Scenarios for use

You can implement RAS Gateway in several different configurations:

- Multitenant-aware VPN gateway. In this configuration, RAS Gateway is configured as a VPN gateway that is aware of the virtual networks deployed on the Hyper-V hosts. Deploying RAS Gateway with this configuration means that you can connect to the RAS Gateway by using a site-to-site VPN from a remote location or that you can configure individual users with VPN access to the RAS Gateway. The RAS Gateway operates like any other VPN gateway, where it allows the remote users to connect directly to the virtual networks on the Hyper-V servers. The main difference is that the RAS 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 also can use this configuration to provide remote clients direct network access between their on-premises network and the hosted networks.
- Multitenant-aware network address translation (NAT) gateway for Internet access. In this configuration, RAS Gateway provides access to the Internet for virtual machines on virtual networks. The RAS Gateway is configured as a NAT device, which translates addresses that can connect to the Internet to addresses used on the virtual networks. In this configuration, RAS Gateway is also multitenant-aware, so all virtual networks behind the RAS Gateway can connect to the Internet, even if they use overlapping address spaces.
- Forwarding gateway for internal physical network access. In this configuration, RAS 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 configured as a forwarding gateway, RAS Gateway enables computers on the virtual networks to connect to those physical hosts.

Network Controller with RAS Gateway

By using Network Controller you can deploy, configure, and manage Hyper-V hosts and virtual machines that are members of a RAS Gateway pool, so you can provide RAS Gateway services to your tenants. You can use Network Controller to automatically deploy virtual machines running RAS Gateway to support the following features:

- Ability to add and remove gateway virtual machines from the RAS Gateway pool and specify the level of backup required.
- Site-to-site VPN gateway connectivity between remote tenant networks and your datacenter using Internet Protocol security (IPsec).
- Site-to-site VPN gateway connectivity between remote tenant networks and your datacenter using GRE.
- Point-to-site VPN gateway connectivity so that your tenants' administrators can access their resources on your datacenter from anywhere.
- Layer 3 forwarding capability.
- BGP routing, so you can manage the routing of network traffic between your tenants' virtual machine networks and their remote sites.

Demonstration: Deploying Network Controller

In this demonstration, you will see how to:

- Add the Network Controller role.
- Configure the Network Controller cluster.
- Configure the Network Controller application.
- Validate the deployment.

Demonstration Steps

Add the Network Controller role

1. Use **Server Manager** to add the **Network Controller** server role.
2. Restart the computer after the role installation is complete.

Configure the Network Controller cluster

1. Create a new Network Controller node. At a Windows PowerShell (Admin) command prompt, run the following command:

```
$node=New-NetworkControllerNodeObject -Name "Node1" -Server "LON-SVR2.Adatum.com" -  
FaultDomain "fd:/rack1/host1" -RestInterface "Ethernet"
```

2. Retrieve details about a certificate in the local computer store for use in client encryption. At the Windows PowerShell (Admin) command prompt, run the following command:

```
$Certificate = Get-Item Cert:\LocalMachine\My | Get-ChildItem | where {$_.Subject -  
imatch "LON-SVR2" }
```

3. Install the Network Controller cluster. At the Windows PowerShell (Admin) command prompt, run the following command:

```
Install-NetworkControllerCluster -Node $node -ClusterAuthentication Kerberos -  
ManagementSecurityGroup "Adatum\Network Controller Admins" -  
CredentialEncryptionCertificate $Certificate
```



Note: This command can take quite a while to complete.

Configure the Network Controller application

- At the Windows PowerShell (Admin) command prompt, run the following command:

```
Install-NetworkController -Node $node -ClientAuthentication Kerberos -  
ClientSecurityGroup "Adatum\Network Controller Ops" -RestIpAddress "172.16.0.99/24" -  
ServerCertificate $Certificate
```

Validate the deployment

- Run the following commands in sequence to validate the deployment:

```
$cred=New-Object Microsoft.Windows.Networkcontroller.credentialproperties
$cred.type="usernamepassword"
$cred.username="admin"
$cred.value="abcd"
New-NetworkControllerCredential -ConnectionUri https://LON-SVR2.Adatum.com -
Properties $cred -ResourceId cred1
Get-NetworkControllerCredential -ConnectionUri https://LON-SVR2.Adatum.com -
ResourceId cred1
```

Question: What does Network Controller use the Northbound and Southbound APIs for?

Lab: Deploying Network Controller

Scenario

A Datum Corporation intends to deploy and use Network Controller to manage network services and devices. You should set up a trial of the technology in a test lab.

Objectives

After completing this lab, you will be able to:

- Prepare to deploy Network Controller.
- Deploy Network Controller.

Lab Setup

Estimated Time: 30 minutes

Virtual machines: **20743A-LON-DC1** and **20743A-LON-SVR2**

User name: **Adatum\Administrator**

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

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

1. On the host computer, start **Hyper-V Manager**.
2. In Hyper-V Manager, click **20743A-LON-DC1**, and in the **Actions** pane, click **Start**.
3. In the **Actions** pane, click **Connect**. Wait until the virtual machine starts.
4. Sign in by using the following credentials:
 - User name: **Adatum\Administrator**
 - Password: **Pa\$\$w0rd**
5. Repeat steps 2 through 4 for **20743A-LON-SVR2**.

Exercise 1: Preparing to deploy Network Controller

Scenario

You decide to deploy Network Controller on a single virtual machine called LON-SVR2 by using Server Manager. Beth Burke and the domain administrator account will be responsible for managing Network Controller and for managing the network using Network Controller. The first stage in your test deployment is to configure the required security groups in AD DS and to obtain a certificate for encryption on the Network Controller server.

The main tasks for this exercise are as follows:

1. Create the required AD DS security groups.
2. Request a certificate for authenticating Network Controller.

► Task 1: Create the required AD DS security groups

1. On **LON-DC1**, from **Server Manager**, open **Active Directory Users and Computers**.
2. Create the following global security groups:
 - **Network Controller Admins**
 - **Network Controller Ops**
3. Add **Beth Burke** and **Administrator** to both these groups.
4. Close **Active Directory Users and Computers**.

► Task 2: Request a certificate for authenticating Network Controller

1. On **LON-SVR2**, open the management console, and then add the **Certificates** snap-in with the focus on the local computer.
2. Request a **Computer** certificate.
3. Close the management console without saving changes.

Results: After completing this exercise, you should have successfully prepared your environment for Network Controller.

Exercise 2: Deploying Network Controller

Scenario

After creating the required groups and obtaining the relevant certificate on LON-SVR2, you must now use Windows PowerShell to deploy Network Controller.


The main tasks for this exercise are as follows:

1. Add the Network Controller role.
2. Configure the Network Controller cluster.
3. Configure the Network Controller application.
4. Verify the deployment.
5. Prepare for the next module.

► Task 1: Add the Network Controller role

1. On **LON-SVR2**, use **Server Manager** to add the **Network Controller** server role.
2. Restart the computer after the role installation is complete.
3. Sign in as **Adatum\administrator** with the password as **Pa\$\$w0rd**.

► Task 2: Configure the Network Controller cluster

 **Note:** These steps are duplicated in the detailed steps for this lab due to the complexity of the Windows PowerShell cmdlets.

1. On **LON-SVR2**, right-click **Start**, and then click **Windows PowerShell (Admin)**.
2. At the Windows PowerShell (Admin) command prompt, type the following command, and then press Enter:

```
$node=New-NetworkControllerNodeObject -Name "Node1" -Server "LON-SVR2.Adatum.com" -
FaultDomain "fd:/rack1/host1" -RestInterface "Ethernet"
```


3. At the Windows PowerShell (Admin) command prompt, type the following command, and then press Enter:

```
$Certificate = Get-Item Cert:\LocalMachine\My | Get-ChildItem | where {$_.Subject -
imatch "LON-SVR2" }
```

4. At the Windows PowerShell (Admin) command prompt, type the following command, and then press Enter:


```
Install-NetworkControllerCluster -Node $node -ClusterAuthentication Kerberos -
ManagementSecurityGroup "Adatum\Network Controller Admins" -
CredentialEncryptionCertificate $Certificate
```

► Task 3: Configure the Network Controller application

 **Note:** This step is duplicated in the detailed steps for this lab due to the complexity of the Windows PowerShell cmdlets.

- On **LON-SVR2**, at the Windows PowerShell (Admin) command prompt, type the following command, and then press Enter:

```
Install-NetworkController -Node $node -ClientAuthentication Kerberos -
ClientSecurityGroup "Adatum\Network Controller Ops" -RestIpAddress "172.16.0.99/24" -
ServerCertificate $Certificate
```

 **Note:** This command can take quite a while to complete.

► Task 4: Verify the deployment



Note: These steps are duplicated in the detailed steps for this lab due to the complexity of the Windows PowerShell cmdlets.

1. On **LON-SVR2**, at the Windows PowerShell (Admin) command prompt, type the following command, and then press Enter:

```
$cred=New-Object Microsoft.Windows.Networkcontroller.credentialproperties
```

2. At the Windows PowerShell (Admin) command prompt, type the following command, and then press Enter:

```
$cred.type="usernamepassword"
```

3. At the Windows PowerShell (Admin) command prompt, type the following command, and then press Enter:

```
$cred.username="admin"
```

4. At the Windows PowerShell (Admin) command prompt, type the following command, and then press Enter:

```
$cred.value="abcd"
```

5. At the Windows PowerShell (Admin) command prompt, type the following command, and then press Enter:

```
New-NetworkControllerCredential -ConnectionUri https://LON-SVR2.Adatum.com -  
Properties $cred -ResourceId cred1
```

6. Press **Y**, and then press Enter when prompted.

7. At the Windows PowerShell (Admin) command prompt, type the following command, and then press Enter:

```
Get-NetworkControllerCredential -ConnectionUri https://LON-SVR2.Adatum.com -  
ResourceId cred1
```

8. You should receive output that looks similar to the output below:

```
Tags                :  
ResourceRef         : /credentials/cred1  
CreatedTime         : 1/1/0001 12:00:00 AM  
InstanceId           : e16ffe62-a701-4d31-915e-7234d4bc5a18  
Etag                 : W/"1ec59631-607f-4d3e-ac78-94b0822f3a9d"  
ResourceMetadata    :  
ResourceId           : cred1  
Properties           : Microsoft.Windows.NetworkController.CredentialProperties
```

► Task 5: Prepare for the next module

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

1. On the host computer, start **Hyper-V Manager**.
2. In the **Virtual Machines** list, right-click **20743A-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20743A-LON-SVR2**.

Results: After completing this exercise, you should have successfully deployed Network Controller.

Question: In the lab, you used Windows PowerShell to manage Network Controller. What other tools could you use?

Question: In the lab, you deployed Network Controller in a domain environment. In a non-domain environment, what steps must you take to provide for authentication?

Module Review and Takeaways

Review Questions

Question: You decide to deploy Network Controller in your AD DS domain environment. What steps must you take to prepare for the deployment?

Question: What are the reasons to consider implementing Software Defined Networking with Windows Server 2016?

Question: How do you install the Network Controller feature in Windows Server 2016 by using Windows PowerShell?

Module 9

Implementing remote access

Contents:

Module Overview	9-1
Lesson 1: Remote access overview	9-2
Lesson 2: Implementing DirectAccess	9-9
Lesson 3: Implementing VPN	9-24
Lab: Implementing DirectAccess	9-36
Module Review and Takeaways	9-41

Module Overview

You can use various remote access technologies to provide secure access to your organization's infrastructure from disparate locations. While organizations normally entirely own and protect local area networks (LANs) by themselves, remote connections to servers, shares, and apps must often travel across unprotected and unmanaged networking infrastructure, such as the Internet.

Any method of using public networks for the transit of organizational data must include a way to protect the integrity and confidentiality of that data. For users that want to connect to enterprise resources remotely, the features in Windows 10 and Windows Server 2016 provide for the functionality and security required.

For example, you can use virtual private networks (VPNs) or DirectAccess to enable your users to access their work environments from anywhere they are connected. In order to support these remote working needs, you must understand these remote access features.

Objectives

After completing this module, you will be able to:

- Describe common remote access solutions and technologies.
- Implement DirectAccess.
- Implement VPNs.

Lesson 1

Remote access overview

You can implement remote access for your users by using a number of technologies, including DirectAccess and VPNs. The type of remote access technology that you choose to implement depends on your organization's business requirements. Some organizations might deploy several remote access technologies on different servers, while other organizations might deploy them on the same server. In this lesson, you will learn about these different remote access technologies.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe remote access technologies.
- Describe the remote access features in Windows Server 2016.
- Explain remote application access in Windows Server 2016.
- Explain when it is necessary to deploy a public key infrastructure (PKI).

Overview of remote access technologies

Choosing the right type of remote access technology enables you to deploy the best solution according to your organizational needs, infrastructure, and technological level. For example, organizations that need administrators to manage servers from the Internet will deploy DirectAccess, and they will deploy Web Application Proxy at the same time if they want to publish internal applications to the Internet. For organizations that still need older technology support, a VPN solution might still be the best choice.

- Remote infrastructure solutions provide access to an internal LAN infrastructure
- Remote application access solutions provide access to applications or services remotely
- You must provide integrity and confidentiality of data and means of communication
- You can deploy a combination of different technologies to accomplish secure and robust solutions

Organizations that have integrated cloud services, such as Microsoft Azure, can use Azure RemoteApp or the Web Apps feature in the Azure App Service to provide a secure and robust app solution for end users. Finally, many organizations might use a combination of remote access technologies to provide an overall, comprehensive remote access solution to suit their needs.

As these technologies developed and grew, remote connectivity to the Internet became standard. Today, when your organization and remote users have connectivity to the Internet, you can provide remote users access to your organizational network; however, this use is inherently not secure. You must choose a technology that encrypts data and control packets to provide secure communications. This can be done with a VPN, DirectAccess, RemoteApp, Web Application Proxy, and Azure, for example, and all of these technologies are used in many situations today.

Remote access features in Windows Server 2016

In order to provide secure network access, you must meet your organization's needs for users to use remote access to organizational resources by using VPNs, DirectAccess, RemoteApp, Web Application Proxy, or Azure. Each of these options represents a technology that you can use in different scenarios to enable access to internal resources from offices in remote site locations or from the Internet.

DirectAccess

DirectAccess enables remote users to securely access organizational resources such as email servers, shared folders, or internal websites. DirectAccess also increases productivity for a mobile workforce because, from users' perspectives, their computers connect to the organization the same way both inside and outside of the office. Users can even make a remote desktop connection to their workstations at their organization's office from anywhere an Internet connection is available. With the unified management experience, you can configure both DirectAccess and older VPN connections from one place. DirectAccess provides the following benefits:

- **DirectAccess advantages include:**
 - Always-on connectivity
 - Seamless connectivity
 - Bidirectional access
 - Remote management
 - Improved security
 - Integrated solution
- **VPN:**
 - Can use with older operating systems
 - Often requires users to establish connections
 - Encrypts and protects data and communications

- **Always-on connectivity.** Whenever a user connects a client computer to the Internet, the client computer is also connected to the intranet. This connectivity enables remote client computers to access and update applications more easily. It also makes intranet resources always available and enables users to connect to an organization's intranet from anywhere at any time, thereby improving their productivity and performance.
- **Seamless connectivity.** DirectAccess provides a consistent connectivity experience, regardless of whether a client computer is local or remote. This enables users to focus more on productivity and less on connectivity options and processes. This consistency can reduce support incidents and training costs for users.
- **Bidirectional access.** You can configure DirectAccess in a way that DirectAccess clients have access to intranet resources, and you can have access from the intranet to those DirectAccess clients. Therefore, DirectAccess can be bidirectional. This ensures that client computers receive recent security updates, that domain Group Policy is enforced, and that there is no difference whether users are on the organizational intranet or on a public network. This bidirectional access also results in:
 - Decreased update time.
 - Increased security.
 - Decreased rate of missed updates.
 - Improved compliance monitoring.
- **Remote management of DirectAccess clients.** Provides the ability to enable only remote management functionality in a DirectAccess client. This new option in the DirectAccess Client Configuration Wizard automates policy deployments for managing a client computer. Remote management of DirectAccess clients does not implement any policy options that allow users to connect to a network for file or application access. Remote management of DirectAccess clients is unidirectional, and it provides incoming-only access for administration purposes only.

- Improved security. Unlike VPNs, DirectAccess offers many levels of access control to network resources. This gives security architects tighter, more precise control over remote users who access specified resources. You can use a detailed policy to define which specific user can use DirectAccess, and the location from which the user can access it. You can use Internet Protocol security (IPsec) encryption for protecting DirectAccess traffic so that users can help ensure that their communication is safe.
- Integrated solution. DirectAccess integrates with server isolation and domain isolation, resulting in the integration of security, access, and health requirement policies between an intranet and remote computers.

VPN

VPN connections enable users who are working off-premises—for example, at home, at a customer site, or from a public wireless access point—to access a server on an organization's private network by using the infrastructure that a public network provides, such as the Internet. The difference between DirectAccess and VPN from a user's perspective is that when using DirectAccess, clients connect automatically from the Internet to the internal network. However, with many VPN solutions, clients must start VPN client software, initiate a VPN connection, and then provide the credentials that the VPN server requires for authentication and authorization. The exact infrastructure of the shared or public network is irrelevant because it appears logically as if the data is sent over a dedicated private link.

Organizations might consider replacing VPN technology with DirectAccess after they migrate their remote access servers to the Windows Server 2016 operating system and their client computers to Windows 10. However, VPN connections are still included in Windows 10 as an option for organizations that prefer to use VPN technology for remote access.

Windows 10 provides a VPN auto-connect feature. This feature enables users to establish VPN connections without the need to manually trigger the VPN connections first. In addition, Windows 10 also includes the VPN Reconnect feature, first introduced with Windows 8. With this feature, users can access an organization's data by using a VPN connection, which automatically reconnects if connectivity is interrupted.

Overview of remote applications access

VPNs and DirectAccess provide a secure means of connecting an end user to an organizational network infrastructure. However, after establishing overall secure communications, there are technologies that connect and run remote software. Many of these technologies simply provide an app to a user remotely, rather than an entire environment. Another option is to use software as a service (SaaS) through Microsoft Azure.

- Enterprise and organizational solutions:
 - RDS
 - Web Application Proxy
 - RemoteApp by using RDS
- Cloud-based subscription solutions:
 - Azure
 - Azure RemoteApp

Remote desktops

Remote Desktop Services (RDS), formerly *Terminal*

Services, provide users with access to a full remote desktop experience. In this scenario, users securely connect to a remote session by using their local Remote Desktop Connection (RDC) client. After they authenticate, users are presented with a full desktop just as if they were signed in locally. The client devices send keystrokes and mouse movements to the server, and the server sends screen images back to the client devices. Users have access to applications as if the applications are running locally, even though they are running on a Remote Desktop Session Host (RD Session Host) server. Each user establishes their

own private session that does not affect any other users who are connected to the same RD Session Host server.

To access any remote desktop, you must add the user account (or domain global group) of the connecting user to the Remote Desktop Users group on the computer to which they are connecting. By default, with the exception of the domain administrator account, this group has no members and, therefore, users cannot make a remote desktop connection until their account has been added to the local Remote Desktop Users group. However, you can configure this during the initial RDS deployment.



Note: Standard users do not have the right to sign in to domain controllers either locally or remotely. Being added to the Remote Desktop Users group on a domain controller does not allow a standard user to sign in to domain controllers. A standard user still needs to be given the right to sign in to a domain controller and must be added to the Remote Desktop Users group to connect to a domain controller remotely.

Installing the RD Session Host role on a server automatically enables remote desktop connections to the local computer and adds users who have been granted access to the local Remote Desktop Users group. If you do not install the RD Session Host role, you can still enable remote desktop access to any Windows-based operating system by modifying the system properties to allow remote connections. This method is limited to Administrators by default, and only two concurrent connections are allowed. You can allow remote connections and select the users who can connect remotely by using the System Properties item in Control Panel.

Remote desktops are well-suited to single-task workers, such as point-of-sale terminals or data-entry workers. In such scenarios, it is important to provide a consistent desktop experience for all workers. Remote desktops also perform well over limited bandwidth, making this a suitable solution for branch offices where information technology (IT) support might be limited. Remote desktops are typically employed with thin clients. Another common use for remote desktops is to enable users to access their organizational desktop. For example, users can work from home by connecting to their workstations.

Normally, you provide remote access to RDS by using the Remote Desktop Gateway (RD Gateway). With the RD Gateway role service, authorized remote users can connect to resources on an internal organizational network from any Internet-connected device by encapsulating RDP traffic into HTTPS envelopes. Access is controlled by configuring the Remote Desktop connection authorization policies (RD CAPs) and the Remote Desktop resource authorization policies (RD RAPs). An RD CAP specifies who is authorized to make a connection, and an RD RAP specifies to which resources authorized users might connect.

Web Application Proxy

The Web Application Proxy role service, by functioning as a reverse web proxy, provides access to internal organizational web applications for users who remotely connect to the organization's network. Web Application Proxy uses Active Directory Federation Services (AD FS) to preauthenticate Internet users, and it acts as an AD FS proxy for publishing claims-aware applications.

AD FS provides users with single sign-on (SSO) functionality. With SSO, users who enter their credentials once to access an organizational web application are not asked to enter their credentials again for subsequent access to the organizational web application for the remainder of that session. After Web Application Proxy configuration is complete, you can publish both claims-aware applications that use AD FS preauthentication and web applications that use pass-through preauthentication.



Note: If you implement pass-through authentication with WAP, you must additionally use AD FS to enable SSO for claims-aware apps.

A typical scenario for Web Application Proxy server placement is in the perimeter network between two firewall devices. The AD FS server and applications that are published are located in the organizational network with domain controllers and other internal servers, and they are protected by the second firewall. This scenario helps provide secure access to organizational applications for users on the Internet. At the same time, this scenario helps protect the organization's IT infrastructure from security threats from the Internet.

You can use Web Application Proxy to publish the RD Gateway. To do this, you must publish the RD Web Access and RD Gateway to fully qualified domain names (FQDNs). If the RD Web Access and RD Gateway roles are on the same server, publish the root FQDN in Web Application Proxy. For example, you can publish at <https://rdg.adatum.com/>. If the roles are on separate servers, you must publish the two virtual directories separately. You can use the same or different external FQDNs. For example you can use, <https://rdweb.adatum.com/rdweb/> and <https://gateway.adatum.com/rpc/>.



Note: You can find out more at the Microsoft TechNet website at: <http://aka.ms/Oappc0>

RemoteApp

RemoteApp programs are accessed remotely through RDS, but they appear as if they are running on the end user's local computer. These applications can appear on the **Start** menu like any locally installed application. On a computer that is running Windows 10, RemoteApp programs can be pinned to the taskbar and are identified by an icon overlay of the Remote Desktop logo. You configure RemoteApp by providing the URL of the web feed to the RemoteApp and Desktop Connections Control Panel app. The format of this address is <https://ServerFQDN/rdweb/feed/webfeed.aspx>, where *ServerFQDN* is the FQDN of the Remote Desktop Web Access (RD Web Access) server.



Note: Secure Sockets Layer (SSL) certificates are required to configure the RemoteApp and Desktop Connections Control Panel app with the secure URL.

RemoteApp is a good choice for users who have laptop computers or mobile devices, such as tablets; users can interact with remote applications in the same manner that they interact with locally installed applications. Running an application on a server avoids compatibility issues that might prevent the application from installing locally. RemoteApp is suited to applications that need to be managed centrally or that have higher computing requirements than users' PCs can satisfy—for example, a business application that requires large amounts of RAM or that requires intensive graphics processing.



Note: When users are outside the corporate network, you must publish the RemoteApp URL through a reverse proxy such as WAP.

Azure RemoteApp

RemoteApp is also a good choice for users with mobile devices. As part of the overall Bring Your Own Device (BYOD) functionality of the modern work force, RemoteApp can successfully bring business software to a device itself. However, there are numerous issues involving security and availability when not on-premises and on the same LAN as the servers that provide the RemoteApp.

Azure RemoteApp can make apps available anywhere that a device has Internet connectivity. You can greatly reduce the often laborious and complex process of making software available from off-premises when you move the process to the cloud. Azure also offers many benefits. With Windows 10, you can join a device to Microsoft Azure Active Directory (Azure AD) and then apply SSO to the various Azure RemoteApp programs on that device. There is no need to sign in to each RemoteApp because the first use

of Azure AD saves a user's account credentials in the local credential store, similar to the email app on the **Start** menu that saves a user's account credentials.

When to deploy a PKI for remote access

When employees of an organization access internal resources from the Internet, it is very important that the communication and data in transit are protected from interception by unauthorized users. Therefore, the communication between the employees located on the Internet and the internal resources should be encrypted. Furthermore, users that connect from the Internet and their computers should be authenticated. Remote access technologies in Windows Server 2016 use PKI for authenticating users and computers and encrypting data and communication when users are remotely accessing internal resources.

- Will you use PKI for the encryption of data between the client computer and the server?
- Will you use PKI both for encryption and for authenticating users and their computers?
- Will you use self-signed certificates, certificates provided by internal private CAs, or external public CAs?

When planning for using PKI for remote access in your organizations, you should ask the following questions:

- Will you use PKI for encrypting the data between the client computer and the server? In this scenario, the certificate is installed on the Remote Access server only, and users are authenticated with their user name and password.
- Will you use PKI both for encryption and for authenticating users and their computers? In this scenario, you should use PKI for encryption and for issuing certificates to users and computers. Note that some organizations choose to issue certificates to only users or computers.
- Which type of certificates will you use? You can use self-signed certificates or certificates issued by a private certification authority (CA) or by a public CA.
 - Self-signed certificates are issued by the server itself. By default, they are trusted only by the issuing server, and not by other computers in the organization. You use self-signed certificates in small- and medium-sized organizations that use DirectAccess. You configure DirectAccess with the Getting Started Wizard, which provides easy setup and configuration.
 - You use certificates issued by a private CA, such as a CA installed on a server in a domain by using Windows Server 2016 Active Directory Certificate Services (AD CS). You use private CAs in organizations that want to manage their own PKI infrastructure, and where PKI is used for purposes, such as remote access, client authentication, and server authentication. These organizations have significant cost benefits because a large number of certificates are not purchased, but are issued by the private CA. However, certificates issued by a private CA are trusted only by computers that are members of an Active Directory domain where private CA is installed.
 - You use certificates issued by a public CA in organizations that deploy certificates for applications, which need to be trusted by many different operating systems, computers, and devices. Public CAs also are used by organizations that do not have a PKI infrastructure deployed or that need smaller number of certificates.

When deploying DirectAccess infrastructure, organizations can choose between using a private CA, a public CA, or self-signed certificates. However, in organizations that need to deploy an advanced DirectAccess solution, it is not a best practice to use self-signed certificates.

The following table includes the advantages and disadvantages of certificates issued by a private CA or a public CA.

CA type	Advantages	Disadvantages
Private CA	<ul style="list-style-type: none"> Provides greater control over certificate management Lower cost when compared to a public CA Customized templates Automatic enrollment 	<ul style="list-style-type: none"> By default, not trusted by external clients (web browsers, operating systems) Requires greater administration
Public CA	<ul style="list-style-type: none"> Trusted by many external clients (web browsers, operating systems) Requires minimal administration 	<ul style="list-style-type: none"> Higher cost when compared to a private CA Cost is based per certificate Certificate procurement is slower

Some organizations have started using a hybrid approach for their PKI architecture. A hybrid approach uses an external public CA for the root CA, and a hierarchy of internal CAs for distribution of certificates. This gives organizations the advantage of having their internally issued certificates trusted by external clients, while still providing the advantages of an internal CA. The only disadvantage to this method is the cost. A hybrid approach is typically the most expensive approach because public certificates for CAs are very expensive.

Check Your Knowledge

Question	
What is the main benefit from using DirectAccess over a VPN? (Choose two answers.)	
Select the correct answer.	
<input type="checkbox"/>	Faster
<input type="checkbox"/>	A user does not have to initiate a connection
<input type="checkbox"/>	DirectAccess requires more user configuration
<input type="checkbox"/>	With DirectAccess, a user does not have to remember one connection for an internal connection and another for an external connection
<input type="checkbox"/>	VPNs provide internal and external connectivity

Lesson 2

Implementing DirectAccess

The DirectAccess feature in Windows Server 2016 enables remote access to intranet resources without first establishing a user-initiated VPN connection. Consequently, DirectAccess ensures seamless connectivity to the application infrastructure, for both internal users and remote users.

Unlike traditional VPNs that require user intervention to initiate a connection to an intranet, DirectAccess enables any application that supports IPv6 on the client computer to have complete access to intranet resources. DirectAccess also enables you to specify resources and client-side applications that are restricted for remote access.

Understanding how best to implement DirectAccess enables you to select the most appropriate remote infrastructure solution within your organization and to optimize its configuration to support your users' needs.

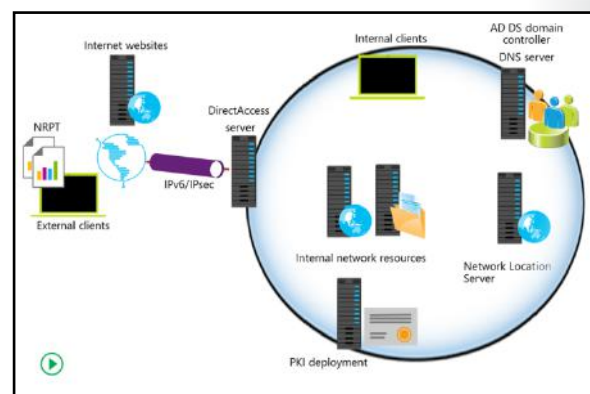
Lesson Objectives

After completing this lesson, you will be able to:

- Describe the components that are required to implement DirectAccess.
- Explain how DirectAccess works for internal clients.
- Explain how DirectAccess works for external clients.
- List the requirements for deploying DirectAccess.
- Explain how to deploy DirectAccess by running the Getting Started Wizard.
- Use the Getting Started Wizard to deploy DirectAccess.
- Explain the limitations of the Getting Started Wizard.
- Address the limitations of using the Getting Started Wizard.
- Monitor DirectAccess.
- Troubleshoot DirectAccess.

Components of DirectAccess

The DirectAccess architecture consists of a number of components that you must deploy and configure within your network infrastructure. Although not based on IPv6, DirectAccess requires IPv6 be enabled. Understanding how DirectAccess uses IPv6 and how best to deploy the components of DirectAccess enables you to deploy an optimized DirectAccess solution.



Components

To deploy and configure DirectAccess, your organization must support the following infrastructure components:

- **DirectAccess server.** The DirectAccess server can be any computer running Windows Server 2016 that you join to a domain that accepts connections from DirectAccess clients, and that establishes communication with intranet resources.
- **DirectAccess clients.** A DirectAccess client can be any domain-joined computer that is running the Enterprise edition of Windows 10, Windows 8.1, Windows 8, or Windows 7.
- **Network location server.** A DirectAccess client uses the network location server to determine its location. If the client computer can securely connect to the network location server by using HTTPS, then the client computer assumes it is on the intranet, and the DirectAccess policies are not enforced. If the client computer cannot contact the network location server, the client assumes it is on the Internet.
- **Internal resources.** These are the server-based resources to which users want to connect.
- **An AD DS domain.** You must deploy at least one AD DS domain running, at a minimum, Windows Server 2003 domain functional level.
- **Group Policy.** You need to use Group Policy for the centralized administration and deployment of DirectAccess settings.
- **Public key infrastructure (PKI).** This is optional for the internal network. It provides the security infrastructure (in terms of certificates) for authentication in some configurations of DirectAccess.
- **DNS server.** You use the DNS server to enable name resolution of the servers in the DirectAccess topology.
- **Name Resolution Policy Table (NRPT).** DirectAccess Group Policy Objects (GPOs) create NRPT entries for client computers. The NRPT has an entry for each DNS namespace that has been configured for DirectAccess.

IPv6 in DirectAccess

DirectAccess uses IPv6 and IPsec when clients connect to internal resources. However, many organizations do not have native IPv6 infrastructure. Therefore, DirectAccess uses transitioning tunneling technologies and communication through IPv4-based Internet to connect IPv6 clients to IPv4 internal resources.

DirectAccess tunneling protocols include:

- **ISATAP.** ISATAP enables DirectAccess clients to connect to the DirectAccess server over the IPv4 networks for intranet communication. By using ISATAP, an IPv4 network emulates a logical IPv6 subnet to other ISATAP hosts, where ISATAP hosts automatically tunnel to each other for IPv6 connectivity. ISATAP does not need changes on IPv4 routers because IPv6 packets are tunneled within an IPv4 header. To use ISATAP, you have to configure DNS servers to answer ISATAP queries, and enable IPv6 on network hosts.
- **6to4.** 6to4 enables DirectAccess clients to connect to the DirectAccess server over IPv4-based Internet. You can use 6to4 when clients have a public IP address. IPv6 packets are encapsulated in an IPv4 header and sent over the 6to4 tunnel adapter to the DirectAccess server. You can use a GPO to configure the 6to4 tunnel adapter for DirectAccess clients and the DirectAccess server.
- **Teredo.** Teredo enables DirectAccess clients to connect to the DirectAccess server across the IPv4 Internet, when clients are located behind an IPv4 network address translation (NAT) device. Clients that have a private IPv4 address use Teredo to encapsulate IPv6 packets in an IPv4 header and send them over IPv4-based Internet. You can use a GPO to configure Teredo for DirectAccess clients and the DirectAccess server.

- 11-11
or
n
s
PT,
his
r's
he

11-11
or
n
s
PT,
his
r's
he

11-11
or
n
s
PT,
his
r's
he



11-11

or

n

CS

PT,

his

r's

he

11-11

or

n

CS

PT,

his

r's

he

11-11

or

n

CS

PT,

his

r's

he

11-11

or

n

CS

PT,

his

r's

he


- 11-11
- or
- n
- CS
- PT,
- his
- r's
- he

6. The DirectAccess client computer attempts to locate and sign in to the AD DS domain by using its computer account. Because the client no longer references any DirectAccess rules in the NRPT for the rest of the connected session, all DNS queries are sent through interface-configured DNS servers (intranet-based DNS servers). With the combination of network location detection and computer domain sign-in, the DirectAccess client configures itself for normal intranet access.
7. Based on the computer's successful sign-in to the domain, the DirectAccess client assigns the domain (firewall network) profile to the attached network. By design, DirectAccess connection security tunnel rules are scoped for public and private firewall profiles, and they are disabled from the list of active connection security rules.
8. The DirectAccess client has successfully determined that it is connected to its intranet, and it does not use DirectAccess settings (NRPT rules or connection security tunnel rules). The DirectAccess client can access intranet resources normally. It also can access Internet resources through normal means, such as a proxy server.

How DirectAccess works for external clients

When a DirectAccess client cannot reach the URL address specified for the network location server, the DirectAccess client assumes that it is not connected to an intranet and that it is located on the Internet.

When the client computer cannot communicate with the network location server, it starts to use NRPT and connection security rules. NRPT has DirectAccess-based rules for name resolution, and connection security rules define DirectAccess IPsec tunnels for communication with intranet resources. Internet-connected DirectAccess clients use the following process to connect to intranet resources:



1. The DirectAccess client attempts to access the Network Location Server.
2. The client attempts to locate a domain controller.
3. The client attempts to access intranet resources first, and then Internet resources.

DirectAccess client attempts to access the network location server

DirectAccess clients attempt to access the network location server as follows:

1. The client tries to resolve the FQDN of the network location server URL. Because the FQDN of the network location server URL corresponds to an exemption rule in the NRPT, the DirectAccess client does not send the DNS query to a locally configured DNS server (an Internet-based DNS server). An external, Internet-based DNS server would not be able to resolve the name.
2. The DirectAccess client processes the name resolution request as defined in the DirectAccess exemption rules in the NRPT.
3. Because the network location server is not found on the same network where the DirectAccess client is currently located, the DirectAccess client applies a public or private firewall network profile to the attached network.

4. The connection security tunnel rules for DirectAccess, scoped for public and private profiles, provide a public or private firewall network profile.
5. The DirectAccess client uses a combination of NRPT rules and connection security rules to locate and access intranet resources across the Internet through the DirectAccess server.

DirectAccess client attempts to locate a domain controller

After starting up and determining its network location, the DirectAccess client attempts to locate and sign in to a domain controller. This process creates an IPsec tunnel or an infrastructure tunnel by using the IPsec tunnel mode and Encapsulating Security Payload (ESP) to the DirectAccess server. The process is as follows:

1. The DNS name for the domain controller matches the intranet namespace rule in the NRPT, which specifies the IPv6 address of the intranet DNS server. The DNS client service constructs the DNS name query that is addressed to the IPv6 address of the intranet DNS server and forwards it to the DirectAccess client's TCP/IP stack for sending.
2. Before sending the packet, the TCP/IP stack checks to determine if there are Windows Firewall outgoing rules or connection security rules for the packet.
3. Because the destination IPv6 address in the DNS name query matches a connection security rule that corresponds with the infrastructure tunnel, the DirectAccess client uses Authenticated Internet Protocol (AuthIP) and IPsec to negotiate and authenticate an encrypted IPsec tunnel to the DirectAccess server. The DirectAccess client (both the computer and the user) authenticates itself with its installed computer certificate and its NTLM credentials, respectively.



Note: AuthIP enhances authentication in IPsec by adding support for user-based authentication with the Kerberos protocol or SSL certificates. AuthIP also supports efficient protocol negotiation and using multiple sets of credentials for authentication.

4. The DirectAccess client sends the DNS name query through the IPsec infrastructure tunnel to the DirectAccess server.
5. The DirectAccess server forwards the DNS name query to the intranet DNS server. The DNS name query response is sent back to the DirectAccess server and back through the IPsec infrastructure tunnel to the DirectAccess client.

Subsequent domain sign-in traffic goes through the IPsec infrastructure tunnel. When a DirectAccess client user signs in, the domain sign-in traffic goes through the IPsec infrastructure tunnel.

DirectAccess client attempts to access intranet resources

The first time a DirectAccess client sends traffic to an intranet location that is not on the list of destinations for the infrastructure tunnel, such as an email server, the following process occurs:

1. The application or process that attempts to communicate constructs a message or payload, and hands it off to the TCP/IP stack for sending.
2. Before sending the packet, the TCP/IP stack checks to determine if there are Windows Firewall outgoing rules or connection security rules for the packet.
3. Because the destination IPv6 address matches the connection security rule that corresponds with the intranet tunnel, which specifies the IPv6 address space of the entire intranet, the DirectAccess client uses AuthIP and IPsec to negotiate and authenticate an additional IPsec tunnel to the DirectAccess server. The DirectAccess client authenticates itself with its installed computer certificate and the user account's Kerberos credentials.

4. The DirectAccess client sends the packet through the intranet tunnel to the DirectAccess server.
5. The DirectAccess server forwards the packet to the intranet resources. The response is sent back to the DirectAccess server and back through the intranet tunnel to the DirectAccess client.

Any subsequent intranet access traffic that does not match an intranet destination in the infrastructure tunnel connection security rule goes through the intranet tunnel.

DirectAccess client attempts to access Internet resources

When a user or process on a DirectAccess client attempts to access an Internet resource, such as an Internet web server, the following process occurs:

1. The DNS client service passes the DNS name for the Internet resource through NRPT. If there are no matches, the DNS client service constructs the DNS name query that is addressed to the IP address of an interface-configured Internet DNS server and hands it off to the TCP/IP stack for sending.
2. Before sending the packet, the TCP/IP stack checks to determine if there are Windows Firewall outgoing rules or connection security rules for the packet.
3. Because the destination IP address in the DNS name query does not match the connection security rules for the tunnels to the DirectAccess server, the DirectAccess client sends the DNS name query normally.
4. The Internet DNS server responds with the IP address of the Internet resource.
5. The user application or process constructs the first packet to send to the Internet resource. Before sending the packet, the TCP/IP stack checks to determine if there are Windows Firewall outgoing rules or connection security rules for the packet.
6. Because the destination IP address in the DNS name query does not match the connection security rules for the tunnels to the DirectAccess server, the DirectAccess client sends the packet normally.

Any subsequent Internet resource traffic that does not match a destination in either the infrastructure intranet tunnel or connection security rules is sent and received normally.

Accessing the domain controller and intranet resources is very similar to the connection process because both of these processes use NRPT to locate an appropriate DNS server to resolve name queries. However, the main difference is in the IPsec tunnel that establishes between the client and the DirectAccess server. When accessing the domain controller, all the DNS queries are sent through the IPsec infrastructure tunnel, and when accessing intranet resources, a second IPsec tunnel establishes to access intranet resources.

Requirements and prerequisites

Organizations could choose different DirectAccess server deployment options depending on their business requirements. Deployment options might vary from using the Getting Started Wizard for a simple deployment to using advanced configuration options for a more complex deployment. However, the server on which you plan to install the DirectAccess server role should meet the following prerequisites:

- The DirectAccess server must be a domain member. You cannot deploy the DirectAccess server role on workgroup server computers.
- The DirectAccess server must have at least one network adapter connected to the domain network.
- Windows Firewall must be enabled on all profiles. You should not turn off Windows Firewall on the DirectAccess server and the DirectAccess clients, because turning off Windows Firewall will disable DirectAccess connectivity.
- The DirectAccess server cannot be a domain controller. Deploying the DirectAccess server role on a domain controller is not supported.
- You must deploy the DirectAccess server in one of the following network topologies:
 - Edge. You use this topology in organizations where firewall software is deployed on an edge computer that is running Windows Server 2016. The edge computer must have two network adapters: one network adapter that connects to the internal network, and the other network adapter that connects to the Internet.
 - Behind the firewall with two network adapters. You use this topology in organizations that use an edge device as a firewall solution. In this scenario, the DirectAccess server is located in a perimeter network, behind the edge device. The DirectAccess server must have two network adapters: one network adapter that connects to the internal network, and the other network adapter that connects to the perimeter network.
 - Behind the firewall with one network adapter. You use this topology in organizations that use an edge device as a firewall solution where the DirectAccess server has one network adapter connected to the internal network.

- The DirectAccess server:
 - Must be a domain member
 - Must have at least one network adapter connected to the domain network
 - Must have Windows Firewall enabled on all profiles
 - Cannot be a domain controller
- You must deploy the DirectAccess server in one of the following network topologies:
 - Edge
 - Behind the firewall with two network adapters
 - Behind the firewall with one network adapter

Using the Getting Started Wizard

Organizations could choose to deploy and configure DirectAccess server options differently depending on their business requirements. Consequently, the process for configuring DirectAccess might vary from using the Getting Started Wizard for a simple deployment to using advanced configuration options for a more complex deployment.

The Getting Started Wizard makes multiple configuration changes so that DirectAccess clients can connect to the intranet. These changes include:

The Getting Started Wizard makes the following configuration changes:

- GPO settings
 - DirectAccess Server Settings GPO
 - DirectAccess Client Settings GPO
- DNS server settings
- Remote clients
- Remote access server
- Infrastructure servers

- GPO settings. The wizard creates the following two GPOs to determine which computers will be DirectAccess servers and which will be DirectAccess clients:
 - The DirectAccess Server Settings GPO. Defines the settings that will apply to the DirectAccess servers. These settings include:
 - Global Settings. Define the IPsec Internet Control Message Protocol (ICMP) that will be allowed through the local firewall on the DirectAccess server.
 - Inbound Rules. Define inbound IP-HTTPS traffic to provide connectivity across HTTP proxies and firewalls. Inbound rules also allow traffic to the DNS64 server that is deployed on the Remote Access server.
 - Connection Security Settings. Define the IPv6 address prefixes and the Kerberos authentication settings.
 - The DirectAccess Client Settings GPO. Defines the settings that will apply to the DirectAccess clients. These settings include:
 - Public Key Policies/Trusted Root Certification Authorities. The Getting Started Wizard configures the DirectAccess client computers to trust the self-signed certificates that the DirectAccess server issues.
 - Global Settings. Define the IPsec ICMP protocol that the Getting Started Wizard will allow through the local firewall on the DirectAccess clients.
 - Outbound Rules. Define the outbound IP-HTTPS traffic to provide connectivity across HTTP proxies and firewalls.
 - Connection Security Settings. Define the IPv6 address prefixes and the Kerberos authentication settings.
- DNS server settings. In the DNS Manager console, under Forward Lookup Zones, the Getting Started Wizard creates host (A and AAAA) resource records for the following hosts: directaccess-corpConnectivityHost, DirectAccess-NLS, and directaccess-WebProbeHost.
- Remote clients. In the Getting Started Wizard, you can configure the following DirectAccess settings for client computers:
 - Select groups. You can select which groups of client computers will be configured for DirectAccess. By default, the Domain Computers group is configured for DirectAccess. In the Getting Started Wizard, you can edit this setting and replace the Domain Computers group with a custom security group.

- Enable DirectAccess for mobile computers only. This setting is enabled by default, and you can disable it in the Getting Started Wizard.
- Network Connectivity Assistant. Network Connectivity Assistant runs on every client computer and provides DirectAccess connectivity information, diagnostics, and remediation support.
- Resources that validate connectivity to internal network. DirectAccess client computers need information that will help them decide whether they are located on the intranet or Internet. For this reason, they will contact resources you provide in this wizard. You can provide a URL, which the Getting Started Wizard will access by using a HTTP request, or a FQDN that the wizard will contact by the **ping** command. By default, this setting is not configured.
- Helpdesk email address. By default, this setting is not configured.
- DirectAccess connection name. The default name is Workplace Connection.
- Allow DirectAccess clients to use local name resolution. This setting is disabled by default.

After you configure these settings in the Getting Started Wizard, the wizard then suggests options based on your settings.

- Remote Access server. In the Getting Started Wizard, you define the network topology where the DirectAccess server is located:
 - On an edge of the internal corporate network, where the edge server has two network adapters.
 - On a server located behind an edge device, where the server has two network adapters.
 - On a server located behind an edge device, where the server has one network adapter.

The wizard detects the network topology and suggest the recommended topology settings. The public name or IPv4 address where DirectAccess clients connect from the Internet is already entered in the wizard.

You can also define the network adapter to which the DirectAccess clients connect, in addition to the certificates that the IP-HTTPS connections use.

- Infrastructure servers. In the Getting Started Wizard, you define infrastructure servers. DirectAccess clients connect to these servers before they connect to internal corporate resources. By default, two entries are configured: the domain name suffix, and the DirectAccess-NLS name followed by the domain name suffix. For example, if the domain name is contoso.com, then the following entries are configured: contoso.com and DirectAccess-NLS.contoso.com.

Demonstration: Configuring DirectAccess with the Getting Started Wizard

In this demonstration, you will see how to:

- Configure DirectAccess using the Getting Started Wizard.
- Verify that the DirectAccess client is configured.

Demonstration Steps

1. On **LON-RTR**, open **Server Manager**, and then select **Remote Access Management**. Complete the **Getting Started Wizard** with the following settings:
 - a. On the **Configure Remote Access** page, click **Deploy DirectAccess only**.
 - b. Verify that **Edge** is selected, and then in the **Type the public name or IPv4 address used by clients to connect to Remote Access server** text box, type **131.107.0.200**.

- c. On the **Remote Access Review** page, remove the **Domain Users** group, and then add the **DA_Clients** group.
 - d. On the **Remote Access Review** page, clear the **Enable DirectAccess for mobile computers only** check box.
 - e. On the **DirectAccess Client Setup** page, set **Windows 10 Workplace Connection** as the **DirectAccess connection name**.
 - f. Provide the defaults for all other pages, and on the **Configure Remote Access** page, click **Finish**, and then in the **Applying Getting Started Wizard Settings** dialog box, click **Close**.
2. Restart **LON-RTR**.
 3. Switch to **LON-CL1**.
 4. When you configured the DirectAccess server, the wizard created two Group Policies and linked them to the domain. To apply them, restart **LON-CL1**, and then sign in as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
 5. On **LON-CL1**, open the **Command Prompt** window, and then to force apply Group Policy, type **gpupdate /force**.
 6. At the command prompt, type **gpresult /R** to verify that the DirectAccess Client Settings GPO is applied to the computer settings.



Note: If the DirectAccess Client Settings GPO is not applied, restart **LON-CL1**, and then repeat steps 4 through 6 on **LON-CL1**.

7. Type the following command at the command prompt, type the following cmdlet, and then press Enter:

```
netsh name show effectivepolicy
```

8. Verify that the "DNS Effective Name Resolution Policy Table Settings" message appears.

Limitations of the Getting Started Wizard

The Getting Started Wizard is easy to implement, but it is not suitable for large deployments. Specifically, it is not suited for large deployments that need to support multisite access, that require a highly-available infrastructure, or that require support for computers running Windows 7 in a DirectAccess scenario. The following limitations are identified:

- **Self-Signed Certificates.** The Getting Started Wizard creates a self-signed certificate to enable SSL connections to the DirectAccess and network location server servers. In order for DirectAccess to function, you need to ensure that the CRL distribution point for both certificates is available externally. In addition, you cannot use the self-signed certificate in multisite deployments. Because of these limitations, most companies configure either a public certificate for the DirectAccess and network location server servers, or provide certificates generated by an internal CA.

The following limitations of using the Getting Started Wizard are identified:

- Uses self-signed certificates
- Based on network location server design
- No support for Windows 7 and earlier clients

Organizations that have implemented an internal CA can use the web server certificate template to issue a certificate to the DirectAccess and network location server servers. The organizations must also ensure that CRL distribution points are accessible from the Internet.

- **Network location server design.** The network location server is a critical part of a DirectAccess deployment. The Getting Started Wizard deploys the network location server on the same server as the DirectAccess server. If DirectAccess client computers on the intranet cannot successfully locate and access the secure web page on the network location server, they might not be able to access intranet resources. When DirectAccess clients obtain a physical connection to the intranet or experience a network status change on the intranet (such as an address change when roaming between subnets), they attempt an HTTPS connection to the network location server URL. If the client can establish an HTTPS connection to network location server and check the revocation status for the web server's certificate, the client determines that it is on the intranet. As a result, the NRPT will be disabled on the client and Windows Firewall will be configured to use the Domain profile with no IPsec tunnels. The network location server needs to be deployed on a highly-available, high-capacity intranet web server. Larger companies will consider implementing the network location server on a Network Load Balancing (NLB) cluster or by using external hardware balancer.
- **Support for Windows 7.** The Getting Started Wizard configures the Remote Access server to act as a Kerberos proxy to perform IPsec authentication without requiring certificates. Client authentication requests are sent to a Kerberos proxy service running on the DirectAccess server. The Kerberos proxy then sends Kerberos requests to domain controllers on behalf of the client. This configuration is only applicable for clients running Windows 8, Windows 8.1, and Windows 10. If Windows 7 clients must be supported for DirectAccess, you must deploy a PKI to issue computer certificates for backward compatibility.

Addressing the limitations of the Getting Started Wizard

Instead of using the Getting Started Wizard, some organizations choose to deploy DirectAccess by configuring advanced features such as PKI, configuring advanced DNS settings, and configuring advanced settings for network location servers and management servers.

The following list are the advanced options that you can use to configure DirectAccess:

- **Scalable and customized PKI infrastructure.** The DirectAccess deployment can benefit from a custom PKI solution, whether used with a public or private CA. You can configure the PKI components according to the organization's business requirements, for example to provide support for computers running Windows 7.
- **Customized network configurations options.** Organizations can benefit from deploying DirectAccess that meets specific network topology and design, including complex scenarios such as multiple site and multiple domain deployments. You can configure the DirectAccess clients so that they can connect to the corporate network by using multiple Internet connections in different geographical locations as DirectAccess entry points. Customized network configuration options include advanced DNS configurations and firewall settings.

The following list are the advanced options that you can use to configure DirectAccess:

- Scalable and customized PKI infrastructure
- Customized network configurations options
- Scalable and highly available server deployment
- Customized monitoring and troubleshooting

- Scalable and highly available server deployment. While configuring advanced DirectAccess options, organizations can use a variety of solutions for better scalability of the servers. This helps them achieve their business goal of better remote access performance. Additionally, in cases where DirectAccess is a business critical solution, organizations can deploy multiple servers that are highly available so that no single point of failure exists and users can establish DirectAccess connectivity regardless of any potential issue. You can also configure management servers that will perform management tasks, such as deploying Windows updates on DirectAccess clients and servers.
- Customized monitoring and troubleshooting. Advanced DirectAccess options include customized monitoring and troubleshooting options that will help you to diagnose any resolve any potential DirectAccess issues quickly.

Monitoring DirectAccess

To help to ensure that your remote users can connect to internal resources without interruption due to service availability issues, it is important that you monitor DirectAccess.

You can monitor DirectAccess connectivity by using the Remote Access Management console. This console contains information on how DirectAccess server components work. When you use the Remote Access Management console, you can also monitor DirectAccess client connectivity information. When you monitor DirectAccess connectivity, you can obtain information about the DirectAccess role service health that helps you troubleshoot potential connectivity issues.

The Remote Access Management Console includes the following monitoring components:

- Dashboard. The Remote Access Management Console includes a centralized dashboard for multiple DirectAccess monitored components. Information about each of the dashboard components is available in separate windows in the Remote Access Management Console.
 - The dashboard contains the following information:
 - Operation status
 - Configuration status
 - DirectAccess status
 - VPN client status
- Operation status. Operation status provides the following information about the health of each DirectAccess component:
 - DNS
 - DNS64
 - Domain controllers
 - IP-HTTPS
 - Kerberos
 - NAT64

- The Remote Access Management Console monitoring components include:
 - Dashboard
 - Operation status
 - Remote Access client status
 - Remote Access reporting
- You can troubleshoot DirectAccess connectivity by using:
 - A troubleshooting methodology
 - Command-line tools
 - GUI tools

- network adapters
- NLS
- Network security and services

If the DirectAccess component is healthy, it is marked with a green check mark. If there is any issue with the DirectAccess component, it is marked with a blue question mark. When you click a component, you can obtain more detailed information about the related issue, the cause of the issue, and how to resolve it.

- Remote Access Client Status. Remote Access Client Status displays information about the DirectAccess client computers that connect to the DirectAccess server. The information displayed in the Remote Access Client Status window includes the following:
 - User Name
 - Host Name
 - ISP Address, Protocol/Tunnel
 - Duration

Keep in mind that for each DirectAccess client connection you can view more detailed information.

- Remote Access reporting. Remote Access reporting provides the same information as Remote Access Client Status, but in the form of a historical DirectAccess client usage report. You can choose the start date and end date for the report. In addition, Remote Access Reporting displays the following:
 - Server Load Statistics, which is statistical connectivity information on Total DirectAccess sessions
 - Average sessions per day
 - Maximum concurrent sessions
 - Unique DirectAccess clients

Troubleshooting DirectAccess

When problems occur, knowing how to troubleshoot issues quickly can help to ensure your users are not inconvenienced by excessive service unavailability and help you to meet your agreed service level agreement (SLA). You should develop a troubleshooting methodology for DirectAccess connectivity in order to quickly eliminate any problems that DirectAccess client computers might encounter. The troubleshooting methodology should contain step-by-step instructions on how to diagnose the problem.

When a DirectAccess client computer cannot connect to intranet resources from the Internet, there can be a number of possible reasons it cannot connect. Use the following high-level troubleshooting procedure to help pinpoint the problem:

1. Verify that the client is running a supported operating system. A DirectAccess client computer must be running, Windows Server 2016, Windows 10 Enterprise, Windows 10 Education, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2, Windows 7 Ultimate, Windows 7 Enterprise, Windows 8 Enterprise, or Windows 8.1 Enterprise.

To troubleshoot DirectAccess, verify the following:

1. The client is running a supported operating system
2. The client computer is part of an AD DS domain
3. The client computer belongs to a suitable AD DS security group
4. Client GPOs are applying
5. The server configuration GPOs are applying
6. IPv6 connectivity is working
7. The DirectAccess client has IPv6 connectivity to the intranet DNS servers
8. The DirectAccess client has correctly determined its location



Note: If you configured DirectAccess by using the Getting Started Wizard, the DirectAccess client computer must be running Windows Server 2016, Windows Server 2012, Windows Server 2012 R2, Windows 10, Windows 8.1, or Windows 8.

2. Verify that the client computer is part of an AD DS domain within your forest.
3. Verify that the client computer belongs to a suitable AD DS security group for the purposes of applying GPOs for DirectAccess. The appropriate group is configured during the setup of DirectAccess.
4. Verify that the two GPOs are created and configured correctly. Use the Group Policy troubleshooting tools to verify the correct application of GPOs for DirectAccess.
5. Check that the server configuration GPOs are applying to the DirectAccess server. Again, use standard GPO troubleshooting tools and techniques.
6. Check IPv4 and IPv6 connectivity from the DirectAccess client to the DirectAccess server. IP connectivity is required for DirectAccess.



Note: Although IPv6 must be enabled, because IPv4 tunneling is used, you must verify IPv4 connectivity in addition to verifying IPv6 is enabled correctly.

7. Also, check that the DirectAccess client has IPv6 connectivity to the intranet DNS servers. The DirectAccess client must be able to use these servers to resolve intranet FQDNs.
8. Verify that the DirectAccess client has correctly determined its location as being on the Internet. You can use the **netsh dnsclient show state** command to make this determination. The determined network location displays in the Machine Location field.

Useful netsh commands for DirectAccess

In addition to the command shown above, you can use the following **netsh** commands to troubleshoot DirectAccess connectivity issues:

- **Netsh interface Teredo show state.** This command is useful for determining whether the client-side GPOs have successfully applied.
- **Netsh interface httpstunnel show interface.** Displays detailed information about the IP-HTTPS adapter on your computer. Enables you to see the name of the IP-HTTPS listener that runs on your DirectAccess server, in addition to whether the adapter is currently connected.
- **Netsh namespace show policy.** Should display the same information as you entered into the Name Resolution Policy Table during the setup process on your DirectAccess server. If it does not, it means that the GPOs have not applied to the local computer yet.
- **Netsh namespace show effectivepolicy.** When the DirectAccess client is external, the output mirrors the output from the **netsh namespace show policy** command. When the DirectAccess client is internal, the output says "Note: DirectAccess settings would be turned off when computer is inside corporate network."
- **Netsh advfirewall show currentprofile.** Shows which Windows Firewall profile is active. The IPsec tunnels are only enabled on the Public and Private profiles. If the Domain profile is active, then DirectAccess is not enabled.

Windows PowerShell

You can use the following Windows PowerShell cmdlets to investigate DirectAccess client problems:

- **Get-DAClientExperienceConfiguration.** This cmdlet retrieves the DirectAccess client configuration.
- **Get-DACConnectionStatus.** This cmdlet retrieves the status of a DirectAccess connection in Windows 8.1 and Windows 10.

View DirectAccess client-side settings in Windows 10

In Windows 10, you can access DirectAccess client-side settings from the NETWORK & INTERNET app in SETTINGS. This tool enables you to:

- Obtain DirectAccess connectivity information so that you can view DirectAccess connectivity status from a client computer.
- Pick a DirectAccess entry point from those available to you.
- Obtain logging information. You can enable logging for your DirectAccess connection.

Question: How do you configure DirectAccess clients?

Question: How does the DirectAccess client determine if it is connected to the intranet or the Internet?

Lesson 3

Implementing VPN

VPN provides secure access to organizations' internal data and applications to remote clients and devices that are using the Internet. To properly implement and support a VPN environment within your organization, it is important that you understand how to select a suitable tunnelling protocol, configure VPN authentication, and configure the server role to support your chosen configuration.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe various VPN scenarios.
- Describe the tunneling protocols used for a VPN connection.
- Describe the VPN authentication options.
- Explain how to configure a VPN infrastructure.
- Configure a Network Policy Server.
- Identify the process for configuring a VPN client.
- Explain advanced VPN features.
- Configure a VPN infrastructure.

VPN scenarios


A VPN provides a point-to-point connection between components of a private network, through a public network such as the Internet. Tunneling protocols enable a VPN client to establish and maintain a connection to the listening virtual port of a VPN server. To emulate a point-to-point link, the data is encapsulated, or *wrapped*, and prefixed with a header. This header provides routing information that enables the data to traverse the public network to reach its endpoint.

To emulate a private link, the data is encrypted to ensure confidentiality. Packets that are intercepted on the public network are indecipherable without encryption keys. Two types of VPN connections exist:

- Remote access. Remote access VPN connections enable users who are working at home, at customer sites, or from public wireless access points to access a server that exists in your organization's private network. They do so by using the infrastructure that a public network, such as the Internet, provides.

From the user's perspective, the VPN is a point-to-point connection between the computer, the VPN client, and your organization's server. The exact infrastructure of the shared or public network is irrelevant, because it logically appears as if the data is sent over a dedicated private link.

- Site-to-site. Site-to-site VPN connections, which also are known as router-to-router VPN connections, enable your organization to have routed connections between separate offices or with other organizations over a public network, while maintaining secure communications.



Properties of VPN connections

VPN connections in Windows 10 can use:

- Point-to-Point Tunneling Protocol (PPTP)
- Layer Two Tunneling Protocol with IPsec (L2TP/IPsec)
- Secure Socket Tunneling Protocol (SSTP)
- Internet Key Exchange version 2 (IKEv2)



Note: An IKEv2 VPN provides resilience to the VPN client when the client either moves from one wireless hotspot to another or switches from a wireless to a wired connection. This ability is a requirement of VPN Reconnect.

All VPN connections, irrespective of tunneling protocol, share some common characteristics:

- **Encapsulation.** With VPN technology, private data is encapsulated with a header that contains routing information, which allows the data to traverse the transit network.
- **Authentication.** Authentication ensures that the two communicating parties know with whom they are communicating.
- **Data encryption.** To ensure data confidentiality as the data traverses the shared or public transit network, the sender encrypts the data and the receiver decrypts it. The encryption and decryption processes depend on both the sender and the receiver using a common encryption key. Intercepted packets sent along the VPN connection in the transit network will be unintelligible to anyone who does not have the common encryption key.

VPN tunneling protocols

When deploying VPN, you can choose between different tunneling protocols that will be used by clients connecting from public networks. Tunneling protocols for VPN in Windows Server 2016 include PPTP, L2TP, SSTP, and IKEv2. IKEv2 is the preferred VPN tunneling protocol in Windows 10, Windows 8, and Windows 7.

The PPTP, L2TP, and SSTP protocols are based on the Point-to-Point Protocol (PPP) features. For example, IP uses PPP technology; that is, PPP frames encapsulate IP packets and PPP transmits the encapsulated PPP packets across a point-to-point link. Originally, network administrators used PPP between a dial-up client and a network access server for sending data across dial-up or dedicated point-to-point connections.

Windows Server 2016 supports the following four VPN tunneling protocols:

Tunneling protocol	Firewall access	Description
PPTP	TCP port 1723	<ul style="list-style-type: none"> • Provides data confidentiality but not data integrity or data authentication
L2TP/IPsec	UDP port 500, UDP port 1701, UDP port 4500, and IP protocol ID 50	<ul style="list-style-type: none"> • Uses either certificates or preshared keys for authentication • Certificate authentication is recommended
SSTP	TCP port 443	<ul style="list-style-type: none"> • Uses SSL to provide data confidentiality, data integrity, and data authentication
IKEv2	UDP port 500	<ul style="list-style-type: none"> • Supports the latest IPsec encryption algorithms to provide data confidentiality, data integrity, and data authentication

PPTP

You can use PPTP when clients connect via public networks such as the Internet, or for site-to-site VPN connections to encrypt communication between two site locations over the public network.

PPTP enables you to encrypt and encapsulate in an IP header multiprotocol traffic that then is sent across an IP network or a public IP network, such as the Internet:

- Encapsulation. For network transmission, PPTP encapsulates PPP frames in IP datagrams. PPTP uses the following to encapsulate PPP frames for tunneled data:
 - A TCP connection for tunnel management.
 - A modified version of Generic Route Encapsulation (GRE) to encapsulate PPP frames for tunneled data.

You can encrypt, compress, or encrypt and compress payloads of the encapsulated PPP frames.

- Encryption. PPTP uses Microsoft Point-to-Point Encryption (MPPE) to encrypt PPP frames. For this encryption, PPTP uses encryption keys that are generated from the authentication process of Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAP v2) or Extensible Authentication Protocol-Transport Layer Security (EAP-TLS). VPN clients must use the MS-CHAP v2 or EAP-TLS authentication protocol so that the payloads of PPP frames are encrypted. PPTP uses PPP encryption and encapsulates a previously encrypted PPP frame.

L2TP

L2TP is a combination of PPTP and Layer 2 Forwarding and contains the best features of them both. You can use L2TP to encrypt multiprotocol traffic that you want to send over any medium supporting point-to-point datagram delivery, such as IP or asynchronous transfer mode (ATM).

For encryption, the Microsoft implementation of L2TP does not use MPPE; however, it uses IPsec in transport mode. This method is called L2TP/IPsec. To utilize L2TP/IPsec, both the VPN client and server must support L2TP and IPsec.

The encapsulation and encryption methods for L2TP is described as follows:

- Encapsulation. Encapsulation for L2TP/IPsec packets consists of two layers: L2TP encapsulation, and IPsec encapsulation. L2TP encapsulates and encrypts data in the following way:
 - First layer. The first layer is the L2TP encapsulation. A PPP frame (an IP datagram) is wrapped with an L2TP header and a UDP header.
 - Second layer. The second layer is the IPsec encapsulation. The resulting L2TP message is wrapped with:
 - An IPsec ESP header.
 - An IPsec Authentication trailer that provides message integrity and authentication.
 - A final IP header that contains the source and destination IP address that corresponds to the VPN client and server.
- Encryption. Advanced Encryption Standard (AES) or Triple Data Encryption Standard (3DES) uses the encryption keys that the IKE negotiation process generates to encrypt L2TP messages.

SSTP

You can use SSTP to allow PPTP and L2TP/IPsec traffic through firewalls and web proxies. SSTP is a tunneling protocol that uses the HTTPS protocol over TCP port 443. SSTP includes a method to encapsulate PPP traffic over the SSL channel of the HTTPS protocol. Because SSTP uses PPP, strong authentication methods such as EAP-TLS are possible. TLS with enhanced key negotiation, encryption, and integrity checking is possible with SSTP.

When a client tries to establish a SSTP-based VPN connection, the following process happens:

1. SSTP creates a bidirectional HTTPS layer with the SSTP server.
2. On this HTTPS layer, the protocol packets flow as data payload by using the following encapsulation and encryption methods:
 - Encapsulation. SSTP establishes a TCP connection over port 443 encapsulates PPP frames in IP datagrams for transmission over the network.
 - Encryption. The SSL channel of the HTTPS protocol encrypts the SSTP message.

IKEv2

IKEv2 uses the IPsec Tunnel Mode protocol over UDP port 500. IKEv2 supports mobility making it a good protocol choice for a mobile workforce. IKEv2-based VPNs enable users to move easily between wireless hotspots, or between wireless and wired connections.

The use of IKEv2 and IPsec enables support for strong authentication and encryption methods:

- Encapsulation. IKEv2 encapsulates datagrams by using IPsec ESP or Authentication Header (AH) for transmission over the network.
- Encryption. The message is encrypted with one of the following protocols by using encryption keys that are generated from the IKEv2 negotiation process:
 - AES 256
 - AES 192
 - AES 128
 - 3DES encryption algorithms

Authentication options

The authentication of access clients is an important security concern. Authentication methods typically use an authentication protocol that is negotiated during the connection establishment process.

Password Authentication Protocol (PAP) uses plaintext passwords and is the least secure authentication protocol. It typically is negotiated if the remote access client and Remote Access server cannot negotiate a more secure form of validation. PAP is included in Windows Server 2016 to support earlier Windows client operating systems. The following methods are supported by the Remote Access role:

Protocol	Description	Security level
PAP	Uses plaintext passwords; typically used if the remote access client and remote access server cannot negotiate a more secure form of validation	The least secure authentication protocol
CHAP	Uses the industry-standard MD5 hashing scheme; this is a challenge-response authentication protocol	An improvement over PAP; that is, the password is not sent over the PPP link
MS-CHAPv2	Provides two-way authentication, also known as mutual authentication; this is an upgrade of MS-CHAP	The protocol provides stronger security than CHAP
EAP	Allows for arbitrary authentication of a remote access connection through the use of authentication schemes, known as EAP types	The strongest security protocol by providing the most flexibility in authentication variations

CHAP

CHAP is a challenge-response authentication protocol that uses the industry-standard MD5 hashing scheme to encrypt the response. Various vendors of network access servers and clients use CHAP. Because CHAP requires the use of a reversibly encrypted password, you should consider using another authentication protocol, such as MS-CHAP version 2.

MS-CHAP version 2

MS-CHAP v2 is a one-way, encrypted password, mutual-authentication process that works as follows:

1. The authenticator (the remote access server or the computer that is running NPS) sends a challenge to the remote access client. The challenge consists of a session identifier and an arbitrary challenge string.
2. The remote access client sends a response that contains a one-way encryption of the received challenge string, the peer challenge string, the session identifier, and the user password.
3. The authenticator checks the response from the client and sends back a response containing an indication of the success or failure of the connection attempt, and an authenticated response based on the sent challenge string, the peer challenge string, the client's encrypted response, and the user password.
4. The remote access client verifies the authentication response, and if correct, it then uses the connection. If the authentication response is not correct, the remote access client terminates the connection.

EAP

With the Extensible Authentication Protocol (EAP), an arbitrary authentication mechanism authenticates a remote access connection. The remote access client and the authenticator (either the Remote Access server or the Remote Authentication Dial-In User Service [RADIUS] server) negotiate the exact authentication scheme to be used. Routing and Remote Access includes support for EAP-TLS by default. You can plug in other EAP modules to the server that is running Routing and Remote Access to provide other EAP methods.

Other options

In addition to the previously mentioned authentication methods, there are two other options that you can enable when selecting an authentication method:

- Unauthenticated access. This is not actually an authentication method, but rather the lack of one. Unauthenticated access allows remote systems to connect without authentication. You should never enable this option in a production environment, however, because it leaves your network at risk. Nonetheless, this option can sometimes be useful for troubleshooting authentication issues in a test environment.
- Machine certificate for IKEv2. Select this option if you want to use VPN Reconnect.

Configuring a VPN infrastructure

Before deploying your organization's VPN solution, consider the following configuration requirements:

- Your VPN server typically requires two network interfaces. You must determine which network interface will connect to the Internet, and which network interface will connect to your private network. During configuration, you will be asked to choose which network interface connects to the Internet. If you specify the incorrect interface, your remote access VPN server will not operate correctly.

VPN server configuration requirements include:

- Two network interfaces (public and private)
- IP address allocation (static pool or DHCP)
- Authentication provider (NPS/RADIUS or the VPN server)
- DHCP relay agent considerations
- Membership in the Local Administrators group or equivalent

- Determine whether remote clients receive IP addresses from a Dynamic Host Configuration Protocol (DHCP) server on your private network or from the remote access VPN server that you are configuring. If you have a DHCP server on your private network, the remote access VPN server can lease blocks of 10 addresses at a time from the DHCP server, and then assign those addresses to remote clients. If you do not have a DHCP server on your private network, you can configure the remote access VPN server to generate and assign IP addresses to remote clients. If you want the remote access VPN server to assign IP addresses from a range that you specify, you must determine what that range should be.
- Determine whether you want connection requests from VPN clients to be authenticated by a Remote Authentication Dial-In User Service (RADIUS) server or by the remote access VPN server that you are configuring. Adding a RADIUS server is useful if you plan to install multiple remote access VPN servers, wireless access points, or other RADIUS clients to your private network.



Note: To enable a RADIUS infrastructure, install the Network Policy and Access Services server role. The NPS can act as either a RADIUS proxy or a RADIUS server.

- Determine whether VPN clients can send DHCPINFORM messages to the DHCP server on your private network. If a DHCP server is on the same subnet as your remote access VPN server, DHCPINFORM messages from VPN clients will be able to reach the DHCP server after the VPN connection is established. If a DHCP server is on a different subnet from your remote access VPN server, make sure that the router between subnets can relay DHCP messages between clients and the server.
- Ensure that the person who is responsible for the deployment of your VPN solution has the necessary administrative group memberships to install the server roles and configure the necessary services; membership of the local Administrators group is required to perform these tasks.

Configuring a Network Policy Server

Network Policy Server (NPS) is part of the Network Policy and Access Services server role. It enables you to create and enforce organization-wide network access policies for connection request authentication and connection request authorization. You also can use NPS as a RADIUS proxy to forward connection requests to NPS or other RADIUS servers that you configure in remote RADIUS server groups.

You can use NPS to centrally configure and manage network-access authentication, authorization, and client health policies with any combination of the following functions:

- RADIUS server
- RADIUS proxy

RADIUS server

NPS performs centralized connection authentication, authorization, and accounting for wireless, authenticating switch, and dial-up and VPN connections. When using NPS as a RADIUS server, you configure network access servers, such as wireless access points and VPN servers, as RADIUS clients in NPS. You also configure network policies that NPS uses to authorize connection requests, and you can

A Windows Server 2016 Network Policy Server provides the following functions:

- **RADIUS server**
 - NPS performs centralized connection authentication, authorization, and accounting for wireless, authenticating switch, and dial-up and VPN connections
- **RADIUS proxy**
 - You configure connection request policies that indicate which connection requests that the NPS server will forward to other RADIUS servers and to which RADIUS servers you want to forward connection requests



configure RADIUS accounting so that NPS logs accounting information to log files on the local hard disk or in a Microsoft SQL Server database.



Note: NPS is the Microsoft implementation of a RADIUS server. You can use NPS with the Routing and Remote Access service.

When an NPS server is a member of an Active Directory Domain Services (AD DS) domain, NPS uses AD DS as its user-account database and provides single sign-on (SSO), which means that users utilize the same set of credentials for network-access control (authenticating and authorizing access to a network) as they do to access resources within the AD DS domain.

Organizations that maintain network access, such as Internet service providers (ISPs), have the challenge of managing a variety of network-access methods from a single administration point, regardless of the type of network-access equipment they use. The RADIUS standard supports this requirement. RADIUS is a client-server protocol that enables network-access equipment, used as RADIUS clients, to submit authentication and accounting requests to a RADIUS server.

A RADIUS server has access to user-account information, and can verify network-access authentication credentials. If the user's credentials are authentic, and RADIUS authorizes the connection attempt, the RADIUS server then authorizes the user's access based on configured conditions, and logs the network-access connection in an accounting log. Using RADIUS allows you to collect and maintain the network-access user authentication, authorization, and accounting data in a central location, rather than on each access server.

RADIUS proxy

When using NPS as a RADIUS proxy, you configure connection request policies that indicate which connection requests that the NPS server will forward to other RADIUS servers and to which RADIUS servers you want to forward connection requests. You also can configure NPS to forward accounting data for logging by one or more computers in a remote RADIUS server group.

With NPS, your organization also can outsource remote-access infrastructure to a service provider, while retaining control over user authentication, authorization, and accounting.

You can create different NPS configurations for the following solutions:

- Wireless access
- Organization dial-up or VPN remote access
- Outsourced dial-up or wireless access
- Internet access
- Authenticated access to extranet resources for business partners

NPS policies

NPS supports Connection Request Policies and Network Policies. These policies are designed to manage and control connection request attempts for remote access clients and to determine which NPS servers are responsible for managing and controlling connection attempts.

Connection Request Policies. Allow you to designate whether connection requests are processed locally (by the local NPS server) or are forwarded for processing to another RADIUS server.

- With connection request policies, you can use NPS as a RADIUS server or as a RADIUS proxy, based on a variety of factors, including:
 - The time of day and day of the week.
 - The realm name in the connection request.
 - The connection type that you are requesting.
 - The RADIUS client's IP address.

When you install NPS, a default connection request policy is created with the following conditions:

- Authentication is not configured.
 - Accounting is not configured to forward accounting information to a remote RADIUS server group.
 - Attribute manipulation is not configured with rules that change attributes in forwarded connection requests.
 - Forwarding Request is turned on, which means that the local NPS server authenticates and authorizes connection requests.
 - Advanced attributes are not configured.
 - The default connection request policy uses NPS as a RADIUS server.
- Network Policies. Allow you to designate which users are authorized to connect to your network—and the circumstances under which they can, or cannot, connect. A network policy is a set of conditions, constraints, and settings that enable you to designate who is authorized to connect to the network, and the circumstances under which they can or cannot connect.

Each network policy has four categories of properties:

- Overview. Overview properties allow you to specify whether the policy is enabled, whether the policy grants or denies access, and whether a specific network connection method or type of network access server is required for connection requests. Overview properties also enable you to specify whether to ignore the dial-in properties of user accounts in AD DS. If you select this option, NPS uses only the network policy's settings to determine whether to authorize the connection.
- Conditions. These properties allow you to specify the conditions that the connection request must have to match the network policy. If the conditions that are configured in the policy match the connection request, NPS applies the network policy settings to the connection. For example, if you specify the network access server IPv4 address (NAS IPv4 address) as a condition of the network policy, and NPS receives a connection request from a NAS that has the specified IP address, the condition in the policy matches the connection request.
- Constraints. Constraints are additional parameters of the network policy that are required to match the connection request. If the connection request does not match a constraint, NPS rejects the request automatically. Unlike the NPS response to unmatched conditions in the network policy, if a constraint is not matched, NPS does not evaluate additional network policies, and the connection request is denied.
- Settings. The Settings properties allow you to specify the settings that NPS applies to the connection request, provided that all of the policy's network policy conditions are matched and the request is accepted.

When NPS performs authorization of a connection request, it compares the request with each network policy in the ordered list of policies, starting with the first policy and moving down the list. If NPS finds a policy in which the conditions match the connection request, NPS uses the matching policy and the dial-in properties of the user account to perform authorization. If you configure the dial-in properties of the user account to grant or control access through network policy, and the connection request is authorized, NPS applies the settings that you configure in the network policy to the connection:

- If NPS does not find a network policy that matches the connection request, NPS rejects the connection.
- If the dial-in properties of the user account are set to deny access, NPS rejects the connection request anyway.



Note: When you first deploy the NPS role, the two default network policies deny remote access to all connection attempts. You can then configure additional network policies to manage connection attempts.

The process of configuring a VPN client

Once you have completed deployment and configuration of the Network Policy and Access Services role, including NPS, and you have created your Connection Request Policies and Network Policies, you must configure the VPN clients.

To create a VPN connection in Windows 10, use the following procedure:

1. Click the **Network** icon in the notification area, and then click **Network settings**.
2. In **NETWORK & INTERNET**, click the **VPN** tab.
3. Click **Add a VPN connection**.
4. In the **Add a VPN connection** dialog box, in the **VPN provider** list, click **Windows (built-in)**.
5. In the **Connection name** box, enter a meaningful name, such as Office Network.
6. In the **Server name or address** box, type the FQDN of the server to which you want to connect. This is usually the name of the VPN server.
7. In the **VPN type** list, select between **Point to Point Tunneling Protocol (PPTP)**, **L2TP/IPsec with certificate**, **L2TP/IPsec with pre-shared key**, **Secure Socket Tunneling Protocol (SSTP)**, and **IKEv2**. This setting must match the setting and policies configured on your VPN server. In you are unsure, click **Automatic**.
8. In the **Type of sign-in info** list, select either **User name and password**, **Smart card**, **One-time password**, or **Certificate**. Again, this setting must match your VPN server policies.
9. In the **User name (optional)** box, type your user name, and then in the **Password (optional)** box, type your password. Select the **Remember my sign-in info** check box, and then click **Save**.

To manage your VPN connection, from within NETWORK & INTERNET, on the **VPN** tab, click the VPN connection, and then click **Advanced options**. You can then reconfigure the VPN settings as needed.





Note: Your VPN connection appears on the list of available networks when you click the network icon in the notification area.

Using the Connection Manager Administration Kit

Obviously, you can configure each client computer manually by accessing the Network settings app. However, if you are configuring many computers, this would be too time consuming. The Connection Manager Administration Kit (CMAK) allows you to customize users' remote connection options by creating predefined connections to remote servers and networks. The CMAK wizard creates an executable file, which you can then distribute in many ways, or include during deployment activities as part of the operating system image.

Connection Manager is a client network connection tool that allows a user to connect to a remote network, such as an ISP or a corporate network protected by a VPN server. CMAK is a tool that you can use to customize the remote connection experience for users on your network by creating predefined connections to remote servers and networks. You use the CMAK wizard to create and customize a connection for your users.



Note: CMAK is an optional component that is not installed by default. You must install CMAK to create connection profiles that your users can install to access remote networks. To do this, in Windows 10, from **Control Panel**, choose **Programs and Features**, and then click **Turn Windows features on or off**. Select the **RAS Connection Manager Administration Kit (CMAK)** check box and click **OK**.

Distributing the connection profile

The CMAK wizard compiles the connection profile into a single executable file with an .exe file name extension. You can deliver this file to users through any method that is available to you. Some methods to consider are:

- Include the connection profile as part of the image that is included with new computers. You can install your connection profile as part of the client computer images that are installed on your organization's new computers.
- Deliver the connection profile on removable media for the user to install manually. You can deliver the connection profile installation program on a CD/DVD, USB flash drive, or any other removable media that you permit your users to access. Some removable media support autorun capabilities, which allow you to start the installation automatically, when the user inserts the media into the client computer.
- Deliver the connection profile with automated software distribution tools.
 - Many organizations use a desktop management and software deployment tool such as Microsoft System Center Configuration Manager. Configuration Manager provides the ability to package and deploy software that is intended for your client computers. The installation can be invisible to your users, and you can configure it to report back to the management console whether the installation was successful or not.
 - You can also choose to distribute the connection profile by using GPOs.

Advanced VPN features

When you create and configure VPN connections, you can implement a number of advanced features. These features include:

- Always on. You can configure the VPN profile so that Windows initiates the VPN when:
 - The user signs in.
 - There is a change in the network state. For example, when the connection is no longer connected to the corporate network infrastructure directly.
- App-triggered VPN. You can configure the VPN profile to respond to a specific set of apps. For example, when a user loads a defined app, Windows initiates the VPN connection.
- Traffic filters. With traffic filters, you can configure your VPN profiles to initiate a connection only when certain criteria are met. You define these criteria in policies. For example, you can create traffic-based rules that filter based on protocol, address, and port.
- LockDown VPN. You can configure LockDown to secure your user's device so that only the VPN can be used for network communications, preventing the use of Wi-Fi or other network connections.

• When you create and configure VPN connections, you can implement a number of advanced features

- The advanced features include:
 - Always on
 - App-triggered VPN
 - Traffic filters
 - LockDown VPN

Demonstration: Configuring VPNs

In this demonstration, you will see how to:

- Configure a VPN server.
- Configure a VPN client.
- Test a VPN connection.

Demonstration Steps

Configure a VPN server

1. On **LON-RTR**, in **Server Manager**, open the **Remote Access Management** console, and then enable VPN by clicking **Deploy VPN Only** in the **Getting Started Wizard**.
2. When the Routing and Remote Access console opens, right-click **LON-RTR**, and then configure the VPN server with the following options:
 - **Remote access (Dial-up or VPN)**
 - Remote Access: **VPN**
 - Network interface that connects this server to the Internet: **131.107.0.200**
 - Accept all defaults until finished, and then in the warning windows, click **OK**.
3. Open the **Network Policy Server** console, and then enable the **Connections to Microsoft Routing and Remote Access server** network policy.
4. Close both consoles, and restart **LON-RTR**.

Configure a VPN client

1. On **LON-CL1**, open **Network Connections**.
2. Disable **Ethernet**.
3. Enable **Ethernet 2**.
4. Open the **Internet Protocol Version 4 (TCP/IPv4)** properties of **Ethernet 2**, and then ensure the following settings:
 - IP address: **131.107.0.2**
 - Subnet mask: **255.255.0.0**
5. Close Network Connections.
6. On **LON-CL1**, open the **Settings** app, and then browse to **VPN** in **Network & Internet**.
7. In **Add a VPN connection**, provide the following values, and then click **Save**:
 - VPN provider: **Windows (built-in)**
 - Connection name: **Adatum HQ VPN**
 - Server name or address: **131.107.0.200**

Test a VPN connection

1. Connect to the **Adatum HQ** VPN. Sign in as **Adatum\Administrator**, with the password **Pa\$\$w0rd**.
2. The **Adatum HQ** VPN should show a status of **Connected**.
3. Disconnect the VPN.

Verify the correctness of the statement by placing a mark in the column to the right.

Statement	Answer
The SSTP VPN tunneling protocol supports VPN Reconnect.	

Lab: Implementing DirectAccess

Scenario

A. Datum wants to implement a remote access solution for its employees so they can connect to the corporate network while away from the office. While a VPN solution provides a high level of security, business management is concerned about the complexity of the environment for end users, and IT management is concerned that they are not able to manage the remote clients effectively. To address these issues, A. Datum has decided to implement DirectAccess.

You will configure the DirectAccess environment and validate that client computers can connect to the internal network when operating remotely.



Note: When you verify the DirectAccess deployment, you might not be able to connect to the internal file share or the internal Web site since DirectAccess is not fully functional in Windows Server 2016 TP5. This issue will be resolved in Windows Server 2016 RTM.

Objectives

After completing this lab, you will be able to:

- Configure DirectAccess using the Getting Started Wizard.
- Test DirectAccess client's connectivity to internal resources.

Lab Setup

Estimated Time: 45 minutes

Virtual machines: **20743A-LON-DC1**, **20743A-LON-SVR1**, **20743A-INET1**, **20743A-LON-CL1**, and **20743A-LON-RTR**.

User name: **Adatum\Administrator** or **Administrator**

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

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

1. On the host computer, start **Hyper-V Manager**.
2. In Hyper-V Manager, click **20743A-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 using the following credentials:
 - o User name: **Administrator**
 - o Password: **Pa\$\$w0rd**
 - o Domain: **Adatum**
5. Repeat steps 2 through 4 for **20743A-LON-SVR1**, **20743A-LON-CL1**, and **20743A-LON-RTR**.
6. Repeat steps 2 and 3 for **20743A-INET1**.

Exercise 1: Configuring DirectAccess using the Getting Started Wizard

Scenario

You decided to implement DirectAccess as a solution for remote clients that are not able to connect via VPN. Also you want to address management problems, such as GPO application for remote clients. At the beginning, you will configure prerequisite components, and will configure DirectAccess Server.

The main tasks for this exercise are as follows:

1. Run the **Getting Started Wizard** on **LON-RTR**.
2. Verify that the client is configured.
3. Move **LON-CL1** to the external network.

► Task 1: Run the Getting Started Wizard on LON-RTR

1. On **LON-RTR**, open **Server Manager**, and then from the **Tools** menu, select **Remote Access Management**. Complete the **Getting Started Wizard** with the following settings:
 - a. On the **Configure Remote Access** page, click **Deploy DirectAccess only**.
 - b. Verify that **Edge** is selected, and then in the **Type the public name or IPv4 address used by clients to connect to Remote Access server** text box, type **131.107.0.200**.
 - c. On the **Remote Access Review** page, remove the **Domain Computers** group, and then add the **DA_Clients** group.
 - d. On the **Remote Access Review** page, clear the **Enable DirectAccess for mobile computers only** check box.
 - e. On the **DirectAccess Client Setup** page, set **Windows 10 Workplace Connection** as the **DirectAccess connection name**.
 - f. Provide the defaults for all other pages, and on the **Configure Remote Access** page, click **Finish**, and then in the **Applying Getting Started Wizard Settings** dialog box, click **Close**.
2. Restart **LON-RTR**.
3. Sign in to **LON-RTR** as **Adatum\Administrator** with the password of **Pa\$\$w0rd**.

► Task 2: Verify that the client is configured

1. Switch to **LON-CL1** and then restart the computer. After the computer restarts, sign in as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. Open **Command Prompt**, and then to force apply Group Policy, type **gpupdate /force**.
3. At the command prompt, type **gpresult /R** to verify that the DirectAccess Client Settings GPO is applied to the computer settings.



Note: If the DirectAccess Client Settings GPO is not applied, restart **LON-CL1**, and then repeat steps 2 and 3 on **LON-CL1**.

4. Type the following command at the command prompt, and then press Enter:

```
netsh name show effectivepolicy
```

5. Verify that the "DNS Effective Name Resolution Policy Table Settings" message appears.
6. On **LON-CL1**, open **Windows Internet Explorer**.

7. In the Internet Explorer Address bar, type **http://lon-svr1.adatum.com**, and then press Enter. The default IIS web page appears.
8. Leave the **Internet Explorer** window open.
9. Open **File Explorer**, type **\\LON-SVR1\DA**, and then press Enter. Note that you are able to access the folder.

► **Task 3: Move LON-CL1 to the external network**

1. Simulate moving the **LON-CL1** client computer out of the corporate network and to the Internet by disabling the **Ethernet** network adapter and enabling the **Ethernet 2** network adapter.
2. Configure the Ethernet 2 network adapter with following values:
 - IP address: **131.107.0.2**
 - Subnet mask: **255.255.0.0**
 - Preferred DNS server: **131.107.0.100**

Results: After completing this exercise, you will have successfully deployed DirectAccess.

Exercise 2: Testing DirectAccess

Scenario

You decide to test a client connected to the Internet.

The main tasks for this exercise are as follows:


1. Verify connectivity to the internal network resources.
2. Verify connectivity to the DirectAccess server.
3. Verify client connectivity to the DirectAccess server.
4. Prepare for the next module.

► **Task 1: Verify connectivity to the internal network resources**

1. On **LON-CL1**, run the following command at the command prompt, and then press Enter:


```
netsh name show effectivepolicy
```

2. Verify that DNS Effective Name Resolution Policy Table Settings present two entries, one for **.Adatum.com** and one for **Directaccess-NLS.Adatum.com**.
3. On **LON-CL1**, switch to **Internet Explorer**.
4. In the Internet Explorer address bar, type **http://lon-SVR1.adatum.com**, and then press Enter. Verify that the default IIS web page appears.
5. Leave the **Internet Explorer** window open.
6. On the Start screen, type **\\LON-SVR1\DA**, and then press Enter. Note that you are able to access the folder content.

 **Note:** When you verify the DirectAccess deployment, you might not be able to connect to the internal file share or the internal Web site since DirectAccess is not fully functional in Windows Server 2016 TP5. This issue will be resolved in Windows Server 2016 RTM.

7. At the command prompt, run the following command, and then press Enter:

```
ipconfig
```

 **Note:** Notice the IP address for Tunnel adapter is **IPHTTPSInterface** starting with **2002**. This is an IP-HTTPS address.

► **Task 2: Verify connectivity to the DirectAccess server**


1. On **LON-CL1**, open **Windows PowerShell**.
2. At the Windows PowerShell prompt, type the following command, and then press Enter:

```
Get-DAClientExperienceConfiguration
```

 **Note:** Review the DirectAccess client settings.

► **Task 3: Verify client connectivity to the DirectAccess server**

1. Switch to **LON-RTR**.
2. Start the **Remote Access Management Console**, and then click **Remote Client Status**.

 **Note:** Notice that a **Client** is connected via **IPHttps**. In the **Connection Details** pane, in the lower right of the screen, note the use of Kerberos authentication for the machine and the user.

3. Close all open windows.

► Task 4: Prepare for the next module

When you have finished the lab, revert the virtual machines to their initial state.

1. On the host computer, start **Hyper-V Manager**.
2. In the **Virtual Machines** list, right-click **20743A-LON-CL1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20743A-LON-DC1**, **20743A-LON-RTR**, **20743A-INET1**, and **20743A-LON-SVR1**.

Results: After completing this exercise, you will have successfully verified your DirectAccess deployment.

Question: Your organization requires only selected computers to be able to connect from the Internet to the corporate network resources by using DirectAccess. How will you configure the DirectAccess settings to meet the organization's requirements?

Question: In the lab, you used the Getting Started Wizard to configure DirectAccess. In what situations is using the wizard inappropriate?

Module Review and Takeaways

Review Questions

Question: Users are complaining that they are unable to connect to the corporate network by using VPNs following recent firewall configuration changes. The team responsible for implementing security policies has determined that only TCP port 443 is allowed through into the internal network. Which tunneling protocol supports this restriction?

Question: What are the main benefits of using DirectAccess for providing remote connectivity?

Question: What are the main benefits of using DirectAccess for providing remote connectivity?

MCT USE ONLY. STUDENT USE PROHIBITED

Module 10

Deploying and managing Windows and Hyper-V containers

Contents:

Module Overview	10-1
Lesson 1: Overview of containers in Windows Server 2016	10-2
Lesson 2: Deploying Windows Server and Hyper-V containers	10-8
Lesson 3: Installing, configuring, and managing containers by using Docker	10-16
Module Review and Takeaways	10-33

Module Overview

One of the important new features in Windows Server 2016 is the option to deploy containers. By deploying containers, you can provide an isolated environment for applications. You can deploy multiple containers on a single physical server or virtual server, and each container provides a complete operating environment for installed applications. This module introduces you to Windows and Hyper-V containers in Windows Server 2016, and it explains how to deploy and manage these containers.

Objectives

After completing this module, you will be able to:

- Describe containers in Windows Server 2016.
- Deploy Windows Server and Microsoft Hyper-V containers.
- Install, configure, and manage containers by using Docker.

Lesson 1

Overview of containers in Windows Server 2016

After completing this lesson, students will be able to explain the purpose of Windows Server and Hyper-V containers.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe Windows Server containers.
- Describe Hyper-V containers.
- Describe scenarios for using containers.
- Describe the installation requirements for containers.

Overview of Windows Server containers

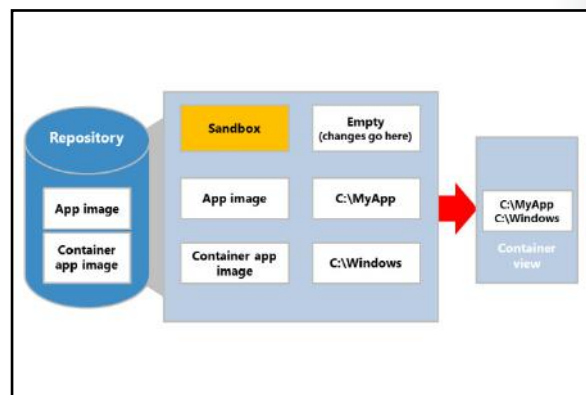
Containers are an isolated operating environment that you can use to provide a controlled and portable space for an app. The container space provides an ideal environment for an app to run without affecting the rest of the operating system (OS) and without the OS affecting the app. Containers enable you to isolate apps from the OS environment.

In many ways, containers are the next evolution in virtualization. Containers are also referred to as *container-based OS virtualization*. Although containers run on the host OS, containers are isolated from one another. Isolated containers improve the security and the reliability of the apps that run within the containers. Containers provide a simulated environment for apps. For example, the local disk appears as a new copy of the OS files, while the memory appears to hold only files and data of the OS that recently started, and the only running component is the OS.

Windows Server 2016 supports two different types of containers, or *runtimes*, each offering different degrees of isolation with different requirements:

- Windows Server containers. These containers provide app isolation through process and namespace isolation technology. Windows Server containers share the OS kernel with the container host and with all other containers that run on the host. While this provides a faster startup experience, it does not provide complete isolation of the containers.
- Hyper-V containers. These containers expand on the isolation that Windows Server containers provide by running each container in a highly optimized virtual machine (VM). However, in this configuration, the OS kernel of the container host does not share with the Hyper-V containers.

Containers appear like a complete OS to an app. Therefore, in many respects, containers are similar to VMs because they run an OS, they support a file system, and you can access them across a network other physical machines or VMs. However, the technology and concepts behind containers are very different from that of VMs.



Container definitions

As you begin creating and working with containers in Windows Server 2016, it is helpful to learn the key concepts that make up the container architecture:

- *Container host.* This element consists of the physical or virtual computer that is configured with the Windows containers feature. The container host can run one or more Windows containers.
- *Container image.* As modifications are made to a containers file system or registry, these changes are captured in the container's sandbox. In many cases, you might want to capture the container image state so that the new containers that you create can inherit the container changes. After you stop the container, you can discard the sandbox, or you can convert it into a new container image. For example, you can install an app into a container and then capture the post-installation state. From this state, you can create a new container image that contains the app. The image will only contain the changes that the installation of the app made, with a layer on top of the container OS image.
- *Container OS image.* While containers are made from images, the container OS image is the first layer in potentially multiple image layers that make up a container. The container OS image provides the OS environment, and it is immutable.
- *Sandbox.* This layer consists of all the changes made to the container, including file system modifications, registry modifications, or software installations. You can keep or discard these changes as required.
- *Container repository.* Each time you make a container image, the container image and its dependencies are stored in a local repository. This allows you to reuse the image many times on the container host.

Finally, it is important to understand that you can manage containers by using the Windows PowerShell command-line interface or by using the open source Docker platform.


Windows Server containers

When you deploy a physical or virtual computer, the computer must have a single user mode that runs on top of a single kernel mode. Computers provide a boundary to allow multiple user modes so you can deploy multiple isolated apps. For example, Hyper-V offers child partitions, or VMs, which can each have their own Windows Server OS with the requisite kernel and user modes, and each app installs in each user mode, or each VM. Containers allow you to have more than one user mode per kernel mode, and they only require one computer per kernel mode.

As noted earlier, a computer deploys with the Windows Server OS with a kernel mode and a user mode. The user mode of the OS manages the container host, or the computer that is hosting the containers. A special stripped down version of the Windows OS, which is stored in a container repository as a container OS image, is used to create a container. This container only features a user mode—this is the distinction between Hyper-V and containers because a VM runs a guest OS with a user mode and a kernel mode. The Windows Server container's user mode allows Windows processes and app processes to run in the container, isolated from the user mode of other containers. When you virtualize the user mode of the OS, Windows Server containers allow multiple apps to run in an isolated state on the same computer, but they do not offer secure isolation.

Overview of Hyper-V containers

It is important to discuss VMs to help you understand Hyper-V containers. VMs also provide an isolated environment for running apps and services. However, a VM provides a full guest OS with kernel and user modes. For example, a computer with the Hyper-V role enabled includes a parent partition, or a management OS, isolated kernel and user modes, and it is responsible for managing the host. Each child partition, or hosted VM, runs an OS with a kernel mode and a user mode.



Similar to VMs, Hyper-V containers are the child partitions that are deployed. On the other hand, the guest OS in Hyper-V containers is not the normal, full Windows OS that we know; it is an optimized, stripped-down version of the Windows Server OS—this is not the same as Nano Server. The boundary provided by the Hyper-V child partition provides secure isolation between the Hyper-V container, other Hyper-V containers on the host, the hypervisor, and the host's parent partition.

Hyper-V containers use the base container image that is defined for the app, and they automatically create a Hyper-V VM by using that base image. When deployed, the Hyper-V container starts in seconds, which is much faster than a VM with a full Windows OS and is even faster than the speed of a Nano Server. The Hyper-V container features an isolated kernel mode, a user mode for core system processes, and a container user mode, which is the same thing that runs in a Windows Server container. In fact, Hyper-V containers use the Windows containers within the VM to store the binaries, libraries, and the app.

Now that the Windows container is running inside a Hyper-V VM, this provides the app with kernel isolation and separation of the host patch and the version level. Because the app is containerized by using Windows containers, you can choose the level of isolation that is required during deployment by selecting a Windows or Hyper-V container. With multiple Hyper-V containers, you can use a common base image that does not require manual management of VMs; the VMs create and delete automatically.

Usage scenarios

Windows Server and Hyper-V container have several practical applications for enterprise environments.

Windows Server containers

While many similarities exist between Windows Server containers and Hyper-V containers, the differences in these virtualization technologies make one more suitable than the other based on your requirements. For example, Windows Server containers are preferred in scenarios where the OS trusts the apps that it hosts, and all the apps must trust each other. In other words, the host OS and apps are within the same trust boundary. This is true for many multiple-container apps, apps that compose a shared service of a larger app, and sometimes apps from the same organization.

Some common usage scenarios for Windows containers include:

- Windows Server containers for:
 - Hosting stateless apps
 - Rapid test deployment
- Hyper-V containers for:
 - Multiple tenants
 - Single tenants
 - Independent lifecycle management

You should ensure that the apps that you deploy in a container on a Windows Server 2016 host are stateless. This type of app does not store any state data in its container. Additionally, keep in mind that containers do not have a GUI. Based on the characteristics of a container, you will probably not run your accounting package in a container. On the other hand, some apps such as games and websites render on local systems, not servers, so they make great examples of apps that are well suited for containers. To summarize, stateless web apps, which do not have a GUI and similar code are the most likely candidates for using Windows container technologies in Windows Server 2016.

Windows Server containers for rapid test deployment

Containers can be used to package and deliver distributed apps quickly. A custom app might require multiple deployments, either weekly or sometimes daily, to keep up with the changes.

Windows Server containers are an ideal way to deploy these apps because you can create packages by using a layered approach to building a deployable app. For example, you can create an image that hosts web sites that includes installed Microsoft Internet Information Services (IIS) and the Microsoft ASP.NET software. Developers can then use that image multiple times to deploy apps without changing the underlying layers. Because Windows Server containers provide greater efficiency in startup times, faster runtime performance, and greater density than Hyper-V containers, developers can spend more time developing apps while requiring fewer resources.



Note: While not unique to Windows Server containers, you can deploy the same package in your test environment to your production environment—it runs the same way it did for the developers and testers. As an added bonus, you also can deploy this container in Microsoft Azure without changing it.

Hyper-V containers

Hyper-V containers each have their own copy of the Windows OS kernel, and have memory assigned directly to them, which is a key requirement of strong isolation. Similar to VMs, you would use Hyper-V containers in scenarios that require central processing unit (CPU), memory, and I/O isolation, for example, a network and storage. The host OS only exposes a small, constrained interface to the container for communication and sharing of host resources. This very limited sharing means that Hyper-V containers are a bit less efficient in startup times and density than Windows Server containers, but they provide the isolation required to allow untrusted apps to run on the same host.

The trust boundary in Hyper-V containers provides secure isolation between the Hyper-V containers on the host, the hypervisor, and the host's other processes. For this reason, Hyper-V containers are the preferred virtualization model in multitenant environments.

Hyper-V containers for multiple tenants

In some situations, you might want to run apps that require different trust boundaries on the same host. For example, you might be deploying a multitenant platform as a service (PaaS) or SaaS offering where you allow your customers to supply their own code to extend the functionality of your service offering. However, you need to ensure that one customer's code does not interfere with your service or gain access to your other customers' data. Hyper-V containers provide the required components for the tenants, but they also ensure that one tenant's application cannot interfere with other applications.

In a typical usage scenario for cloud service providers that host Hyper-V containers, the service provider would have a cluster of Hyper-V hosts that run Windows Server 2016 for a portion of their cloud. This Hyper-V host cluster would host a group of VMs, and each VM would have Windows Server 2016 installed as its guest OS. When you use Windows container technology, each of these VMs would assume the role of a container host. Each container host, or VM, would then be assigned to a different tenant, and the tenant could then create as many containers as it needs on its dedicated container host. If malware or

malicious attack compromised one container host, or VM, the other VMs that belong to other customers would be unaffected.

Hyper-V containers for single tenants

Hyper-V containers might be useful even in a single-tenant environment. One common scenario is where one or more of the apps that you want to host in a Windows container have a dependency on the OS version level or the patch level of the underlying container host. In this scenario, you might consider provisioning a single Hyper-V host, or host cluster, and using Hyper-V containers instead of provisioning and configuring several systems as container hosts and using Windows Server containers.

Hyper-V containers for independent lifecycle management

The other scenario where isolation is very helpful is if you want to run a container with a different version of Windows Server. One of the challenges with Windows Server containers is that they share a significant portion of the OS between the base image and the container image. Consequently, if you upgrade the OS in the base image, then you also need to upgrade the container.

Alternatively, a Hyper-V container enables you to have different versions of base images that allow you to simultaneously host an OS in the container image. This feature is helpful for enterprises that want to have independent lifecycle management for patching, updating, and compliance reasons.

Installation requirements

When planning for Windows containers, you should be aware of the requirements for Windows Server 2016. You should also be familiar with the supported scenarios for Windows Server containers and Hyper-V containers in Windows Server 2016.

Windows container host requirements

When you are planning your deployment, the Windows container host has the following requirements:

- The Windows container role is only available on:
 - Windows Server 2016 Technical Preview 5 (TP5) (Full or Server Core) and newer.
 - Nano Server.
 - Windows 10 (build 14352, and newer).
- If Hyper-V containers are deployed, the Hyper-V role needs to be installed.
- Windows Server container hosts must have the Windows OS installed to **C:**. This restriction does not apply if only Hyper-V containers will deploy.

You should consider the following when planning for Windows containers:

- Windows container host requirements
- Virtualized container host requirements
- Supported scenarios

Host OS	Windows Server container	Hyper-V container
Windows Server 2016 Full UI	Server Core image	Nano Server image
Windows Server 2016 Core	Server Core image	Nano Server image
Windows Server 2016 Nano Server	Nano Server image	Nano Server image
Windows 10 Insider releases	Not available	Nano Server image

Virtualized container host requirements

If you deploy a Windows container host on a Hyper-V VM that is hosting Hyper-V containers, you need to enable nested virtualization. Nested virtualization has the following requirements:

- At least 4 gigabytes (GB) of memory available for the virtualized Hyper-V host.
- On the host system, you will need:
 - Windows Server 2016 TP5 and newer.
 - Windows 10 (build 10565 and newer).
- On the container host VM, you will need:
 - Windows Server TP5 (Full or Server Core).
 - Nano Server.
- A processor with Intel VT-x (this feature is currently only available for Intel processors).
- The container host VM requires at least two virtual processors.

Supported scenarios

Windows Server TP5 is offered with two container OS images: Windows Server Core and Nano Server. Not all configurations support both OS images, however. The following table lists the supported scenario.

Host OS	Windows Server container	Hyper-V container
Windows Server 2016 Full UI	Server Core image	Nano Server image
Windows Server 2016 Core	Server Core image	Nano Server image
Windows Server 2016 Nano Server	Nano Server image	Nano Server image
Windows 10 Insider releases	Not available	Nano Server image

Check Your Knowledge

Question	
In Windows Server 2016 containers, which of the following statements best describes a sandbox?	
Select the correct answer.	
<input type="checkbox"/>	A sandbox is a computer that is configured with containers. This can be a physical computer or a virtual computer.
<input type="checkbox"/>	A sandbox is the first layer of the container hierarchy.
<input type="checkbox"/>	All changes that are made to a running container are stored in the sandbox.
<input type="checkbox"/>	A sandbox is a management tool that you can use instead of the Windows PowerShell command-line interface to manage your containers.

Lesson 2

Deploying Windows Server and Hyper-V containers

Containers provide an isolated and portable operating environment for apps. From the app's perspective, a container appears as an isolated Windows OS with its own file system, devices, and configuration. Windows Server supports two types of containers: Windows Server containers and Hyper-V containers. Windows Server containers achieve isolation through namespace and process isolation, whereas Hyper-V containers encapsulate each container in a lightweight VM. To support your organization's app requirements, you should understand the fundamentals of how to enable and configure Windows Server to support containers.

Lesson Objectives

After completing this lesson, you will be able to:

- Explain how to deploy Windows Server containers.
- Explain how to deploy Hyper-V containers.
- Explain how to manage Windows Server and Hyper-V containers by using Windows PowerShell.
- Deploy Windows Server containers by using Windows PowerShell.

Deploying Windows Server containers

Before you can use containers in Windows Server 2016, you need to deploy a container host. You can choose to deploy containers on a physical host computer or within a VM. You can also choose to use Windows Server 2016, with or without Desktop Experience, or Nano Server. The procedure for deploying container hosts varies depending on the type of container.

Preparing a Nano Server

If you choose to deploy Windows Server containers on a Nano Server, use the high-level steps in the following table to prepare the server for Windows Server containers.

- To prepare a Nano Server:
 - Create a Nano Server virtual hard disk (VHD) file for containers
 - Proceed with the steps below
 - Use the following steps to prepare your Windows Server host for containers:
 1. Install the container feature*
 2. Create a virtual switch
 3. Configure NAT settings
 4. Configure MAC address spoofing
 5. Install container operating system images
- * Step not required if you deploy to a Nano Server

Deployment action	Details
Create a Nano Server virtual hard disk file for containers	<p>Prepare a Nano Server virtual hard disk file with the container and Hyper-V capabilities. This requires that you build a Nano Server image by using the -Compute and -Containers switches. For example, you can type the following code, and then press Enter:</p> <pre>New-NanoServerImage -MediaPath \$WindowsMedia -BasePath c:\nano -TargetPath C:\nano\NanoContainer.vhdx -GuestDrivers -ReverseForwarders -Compute -Containers</pre>

After you deploy the Nano Server, you must then complete the steps listed in the following section.

Preparing the Windows Server host

Use the steps in the following table to prepare your Windows Server host for containers.

Deployment action	Details
Install the container feature*	<p>This step enables the use of Windows Server and Hyper-V containers. You can install this feature by using Server Manager, or you can type the following Windows PowerShell cmdlet, and then press Enter:</p> <pre>Install-WindowsFeature Containers</pre>
Create a virtual switch	<p>All containers connect to a virtual switch for network communications. The switch type can be Private, Internal, External, or NAT. Type the following Windows PowerShell cmdlet to complete this task, and then press Enter:</p> <pre>New-VMSwitch -Name <i>Virtual Switch Name</i> -SwitchType <i>Type</i></pre> <p>If the type is NAT, you must also use the -NATSubnetAddress 172.16.0.0/12 switch, substituting an appropriate subnet address for 172.16.0.0/12.</p>
Configure network address translation (NAT)	<p>If you want to use a virtual switch configured with NAT, you must configure the NAT settings. For example, you can type the following Windows PowerShell command, and then press Enter:</p> <pre>New-NetNat -Name ContainerNat -InternalIPInterfaceAddressPrefix "172.16.0.0/12"</pre> <p>Replace the subnet address with something appropriate for your network.</p>
Configure media access control (MAC) address spoofing	<p>If your container host is virtualized, you must enable MAC address spoofing. For example you can type the following Windows PowerShell command, and then press Enter:</p> <pre>Get-VMNetworkAdapter -VMName <i>Container Host VM</i> Set-VMNetworkAdapter -MacAddressSpoofing On</pre>

* This step is not required if you choose to deploy Windows Server containers to a Nano Server.

Deploying Windows Server containers

Use the high-level steps in the following table to deploy Windows Server containers.

Deployment action	Details
Install container operating system images.	<p>Use the following Windows PowerShell commands to provide the base images for your container deployments:</p> <ol style="list-style-type: none"> This installs the required Windows PowerShell module: <pre>Install-PackageProvider ContainerProvider -Force</pre> This lists the available images by name, version number, and description: <pre>Find-ContainerImage</pre>

Deployment action	Details
	3. This installs the named image: <pre>Install-ContainerImage -Name ImageName -Version Number</pre>

Deploying Hyper-V containers

Before you can use containers in Windows Server 2016, you need to deploy a container host. You can choose to deploy containers either on a physical host computer or within a VM. You can also choose to use Windows Server 2016, with or without Desktop Experience, or Nano Server. The procedure for deploying container hosts varies depending on the type of container.

Preparing a Nano Server

If you choose to deploy Hyper-V containers to a Nano Server, use the high-level steps in the following table to prepare the server for Hyper-V containers.

- To prepare a Nano Server:
 - Create a Nano Server virtual hard disk (VHD) file for containers
 - Proceed with the steps below
 - Use the following steps to prepare your Windows Server host for containers:
 1. Install the container feature*
 2. Enable the Hyper-V role*
 3. Enable nested virtualization
 4. Configure virtual processors
 5. Create a virtual switch
 6. Configure NAT settings
 7. Configure MAC address spoofing
 8. Install container operating system images
- * Step not required if you deploy to a Nano Server

Deployment action	Details
Create a Nano Server virtual hard disk file for containers	Prepare a Nano Server virtual hard disk file with the container and Hyper-V capabilities. This requires that you build a Nano Server image by using the -Compute and -Containers switches. For example, type the following command, and then press Enter: <pre>New-NanoServerImage -MediaPath \$WindowsMedia -BasePath c:\nano -TargetPath C:\nano\NanoContainer.vhdx -GuestDrivers -ReverseForwarders -Compute -Containers</pre>

You must then complete the steps listed in the next section.

Preparing the Windows Server host

Use the steps in the following table to prepare your Windows Server host for containers.

Deployment action	Details
Install the container feature*	You can install this feature by using Server Manager or by using the Install-WindowsFeature Containers Windows PowerShell cmdlet. You can then enable the use of Windows Server and Hyper-V containers.
Enable the Hyper-V role*	You can install this feature by using Server Manager or by using the Install-WindowsFeature Hyper-V Windows PowerShell cmdlet. This is required only if you deploy Hyper-V containers.

Deployment action	Details
Enable nested virtualization	<p>If your container host is a Hyper-V VM, you must enable nested virtualization. Type the following Windows PowerShell command, and then press Enter:</p> <pre>Set-VMProcessor -VMName <i>Container Host VM</i> -ExposeVirtualizationExtensions \$true</pre>
Configure virtual processors	<p>If the container host is a Hyper-V VM, you must configure at least two virtual processors. You can type the following Windows PowerShell command, and then press Enter:</p> <pre>Set-VMProcessor -VMName <i>Container Host VM</i> -Count 2</pre>
Create a virtual switch	<p>All containers connect to a virtual switch for network communications. The switch type can be Private, Internal, External, or NAT. Type the following Windows PowerShell command, and then press Enter:</p> <pre>New-VMSwitch -Name <i>Virtual Switch Name</i> -SwitchType <i>Type</i></pre> <p>If the type is NAT, you must also use the -NATSubnetAddress 172.16.0.0/12 switch, substituting an appropriate subnet address for 172.16.0.0/12.</p>
Configure NAT	<p>If you want to use a virtual switch configured with NAT, you must configure the NAT settings. For example, type the following Windows PowerShell command, and then press Enter:</p> <pre>New-NetNat -Name <i>ContainerNat</i> -InternalIPInterfaceAddressPrefix "172.16.0.0/12"</pre> <p>Replace the subnet address with something appropriate for your network.</p>
Configure MAC address spoofing	<p>If your container host is virtualized, you must enable MAC address spoofing. For this task, type the following Windows PowerShell command, and then press Enter:</p> <pre>Get-VMNetworkAdapter -VMName <i>Container Host VM</i> Set-VMNetworkAdapter -MacAddressSpoofing On</pre>

* These steps are not required if you choose to deploy Hyper-V containers to a Nano Server.

Deploying Hyper-V containers

Use the high-level steps in the following table to deploy Hyper-V containers.

Deployment action	Details
Install container operating system images	<p>Use the following Windows PowerShell commands to provide the base images for your container deployments:</p> <ol style="list-style-type: none"> This installs the required Windows PowerShell module: <pre>Install-PackageProvider ContainerProvider -Force</pre> This lists the available images by name, version number, and description: <pre>Find-ContainerImage</pre> This installs the named image: <pre>Install-ContainerImage -Name ImageName -Version Number</pre>

Managing Windows Server and Hyper-V containers by using Windows PowerShell

After you have deployed a physical or virtual container host, you must create and configure your containers. You can use Windows PowerShell or Docker to administer your containers. Typical tasks include:

- Creating containers.
- Starting containers.
- Connecting to containers.
- Stopping containers.
- Removing containers.
- Configuring containers.

- You can use Windows PowerShell or Docker to administer your containers
- Typical tasks are:
 - Creating containers
 - Starting containers
 - Connecting to containers
 - Stopping containers
 - Removing containers



Note: This module covers how to use Docker to manage Windows containers later.

Creating containers

The first administrative task is to create containers. Start by determining the appropriate container base image. Open an elevated **Windows PowerShell Command Prompt** window, and then use the **Get-ContainerImage** cmdlet to perform this task; the cmdlet returns a list of available base images. Make a note of the container image's name and use the **New-Container** cmdlet to create an image. For example, type the following command to create a new container named **IIS** based on the **WindowsServerCore** container base image, and then press Enter:

```
New-Container -Name IIS -ContainerImageName WindowsServerCore
```

The previous command creates a Windows Server container by default. Alternatively, type the following command to create a new Hyper-V container named **IIS** based on the **WindowsServerCore** container base image, and then press Enter:

```
New-Container -Name IIS -ContainerImageName WindowsServerCore -RunTimeType HyperV
```



Note: Similar to the default method to create a Windows Server container, you can also use the **-RunTimeType** parameter with a value of **Default** to create a Windows Server container. For example type the following code, and then press Enter:

```
New-Container -Name IIS -ContainerImageName WindowsServerCore -RunTimeType Default
```

To start the container, you must complete the creation of the container. Additionally, you must enable a network adapter in your image. Use the following procedure to complete this task:

1. This cmdlet adds a network adapter to your **IIS** image:

```
Add-ContainerNetworkAdapter -ContainerName IIS
```

2. This cmdlet returns a list of available virtual switches; note the name of an appropriate VM switch:

```
Get-VMswitch
```

3. This cmdlet binds the network adapter to your preselected VM switch; substitute **SwitchName** with the value of your preselected switch:

```
Connect-ContainerNetworkAdapter -ContainerName IIS -SwitchName SwitchName
```

Starting containers

You can use the following procedure to start your container. The first cmdlet stores details about the **IIS** container in the **\$container** variable, and the second cmdlet starts the **IIS** container. In Windows PowerShell, run the following cmdlets, and then press Enter:

```
$container = Get-Container -Name IIS
Start-container $container
```

Connecting to containers

You can use PowerShell Direct to connect to a container. This is useful when you must perform a task such as installing software or configuring or troubleshooting a container. Type the following cmdlet to connect to your **IIS** container, and then press Enter:

```
Enter-PSSession -ContainerName IIS -RunAsAdministrator
```

If you are successful in connecting, the Windows PowerShell command prompt changes to include the name of the container—in this case, **IIS**. You can now use any Windows PowerShell cmdlet to add or remove roles and features, to invoke scripts, or to install apps in the **IIS** container.

Stopping containers

You can use the following procedure to stop a container. The first cmdlet stores details about the **IIS** container in the **\$container** variable, and the second cmdlet stops the **IIS** container. In Windows PowerShell, run the following cmdlets, and then press Enter:

```
$container = Get-Container -Name IIS  
Stop-container $container
```

Removing containers

You can use the following procedure to remove a container. The first cmdlet stores details about the **IIS** container in the **\$container** variable, and the second cmdlet removes the **IIS** container. In Windows PowerShell, run the following cmdlets, and then press Enter:

```
$container = Get-Container -Name IIS  
Remove-container $container -Force
```

Configuring containers

As noted earlier, when you create a container in Windows PowerShell, the default container type that is created is a Windows Server container. If your intention is to create a Hyper-V container, one method to resolve this is to remove the container and create it again with the appropriate **-RunTimeType** parameter.

As an alternative to recreating the container, you can use the following procedure to reconfigure the container as the appropriate container type.

The following cmdlet stores details about the **IIS** container in the **\$container** variable; you should run this prior to either of the next two cmdlets:

```
$container = Get-Container -Name IIS
```

The following cmdlet configures the **IIS** container from a Windows Server container to a Hyper-V container:

```
Set-Container $container -RunTimeType HyperV
```

The following cmdlet configures the **IIS** container from a Hyper-V container to a Windows Server container:

```
Set-Container $container -RunTimeType Default
```

Demonstration: Installing the Containers feature and preparing for Docker

In this demonstration, you will see how to:

- Install the Containers feature.
- Prepare for managing containers using Docker.

Demonstration Steps

1. Use **Add roles and features** to add the **Containers** feature.
2. Restart the host, and then sign in as **Adatum\Administrator** with the password **Pa\$\$w0rd**.

3. Open an elevated Windows PowerShell Command Prompt window, by clicking **Run as administrator**.
4. Run the following commands:

```
Invoke-WebRequest "https://get.docker.com/builds/Windows/x86_64/docker-1.12.0.zip" -  
OutFile "$env:TEMP\docker-1.12.0.zip" -UseBasicParsing  
Expand-Archive -Path "$env:TEMP\docker-1.12.0.zip" -DestinationPath $env:ProgramFiles  
[Environment]::SetEnvironmentVariable("Path", $env:Path + ";C:\Program Files\Docker",  
[EnvironmentVariableTarget]::Machine)
```

5. Close and then reopen the a command prompt.
6. In the Windows PowerShell Command Prompt window, run the following commands:

```
dockerd.exe --register-service  
Start-Service docker  
docker pull microsoft/windowsservercore
```



Note: This will take some time to download the image.

Lesson 3

Installing, configuring, and managing containers by using Docker

After completing this lesson, students will be able to install, configure, and manage containers by using Docker.

Lesson Objectives

After completing this lesson, you will be able to:

- Define Docker.
- Describe support for Docker on Windows Server 2016.
- Describe usage scenarios for Docker.
- Explain how to install and configure Docker.
- Describe management with Docker.
- Describe the Docker Hub.
- Describe how Docker integrates with Azure.
- Deploy Hyper-V containers by using Docker.

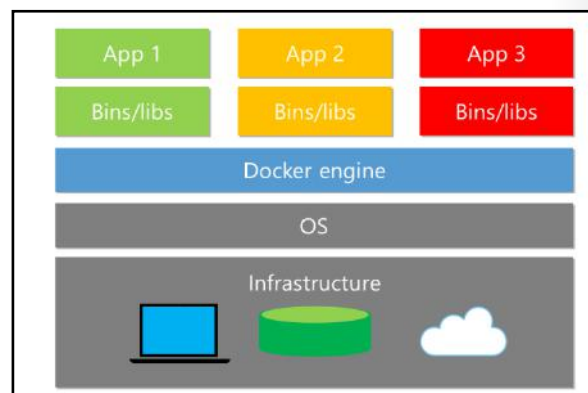
What is Docker?

Docker is a collection of open source tools, solutions, and cloud-based services that provide a common model for packaging, or containerizing, app code into a standardized unit for software development. This standardized unit, also called a *Docker container*, is software wrapped in a complete file system that includes everything it needs to run, which includes code, runtime, system tools, and system libraries, or anything you can install on a server.

Supporting the Docker container is the core of the Docker platform, or the *Docker Engine*. This in-host daemon is a lightweight runtime environment that communicates with the *Docker client* to run commands to build, ship, and run Docker containers. This guarantees that the app always runs the same, regardless of the environment from which it is running.

Docker containers have similar resource isolation and allocation benefits as VMs, but they use a different architectural approach that allows them to be more portable and efficient. These containers include the app and all of its dependencies, but share the kernel with other containers. Each Docker container runs as an isolated process in the user space on the host OS.

Docker containers are based on open standards that allow containers to run on all major Linux distributions and Windows Server 2016 OSs with support for every infrastructure. Because they are not tied to any specific infrastructure, Docker containers can run on any computer, on any infrastructure, and in any cloud.



To guarantee the packaging format remains universal, Docker recently organized the Open Container Initiative, aiming to ensure that container packaging remains an open industry standard with Microsoft as one of the founding members. Microsoft is collaborating with Docker with the aim of enabling developers to create, manage, and deploy both Windows Server and Linux containers by using the same Docker tool set. Developers who target Windows Server will no longer have to make a choice between using the vast range of Windows Server technologies and building containerized apps. By contributing to the Docker project, Microsoft supports using the Docker tool set to manage Windows Server containers and Hyper-V containers.

Docker support on Windows Server 2016

Until recently, it was not possible to use the Windows Server platform to host the Docker Engine without adding an additional layer of virtualization. This changed with the release of Windows Server 2016, which provides a built-in, native Docker daemon for Windows Server hosts. With this component in Windows Server 2016, you can use Docker containers, tools, and workflows in production Windows environments.

The Docker Engine for Windows Server requires Windows Server 2016, and it includes the following key points:

The Docker Engine for Windows Server requires Windows Server 2016, and it includes the following key points:

- No cross-platform containerization
- Two ways to manage containers in Windows OS

- No cross-platform containerization. There is currently no method to present the appropriate kernel to a container from another platform. In other words, Windows containers require a Windows Docker host, and Linux containers require a Linux Docker host.
- Two ways to manage containers in the Windows OS. You can create and manage Windows containers by using the Docker tool set or Windows PowerShell. However, containers that are created with Windows PowerShell cannot be managed with Docker, and containers that are created with the Docker tool set cannot be managed with Windows PowerShell.

Docker components

It is important to understand how Docker works and some basic Docker terminology; some of these are redefined from earlier to provide clarity that is specific to Docker:

- Image. A stateless collection of root file system changes in the form of layered file systems that are stacked on one another.
- Container. A runtime instance of an image, consisting of the image, its operating environment, and a standard set of instructions.

Docker terminology:

- Image, container, Dockerfile, Build

Docker toolbox

- Docker Engine, Docker Compose, Docker machine, Docker client, Kitematic, Docker Registry, Docker Swarm

Docker solutions

- Docker Hub, Docker Trusted Registry, Universal Control Panel, Docker Cloud, Docker Datacenter

- **Dockerfile.** A text file that contains the commands that need to run to build a Docker image.
- **Build.** The process of building Docker images from a **Dockerfile** and any other files in the directory where the image is being built.

Docker toolbox

The Docker toolbox is a collection of Docker platform tools that make building, testing, deploying, and running Docker containers possible. These tools include:

- *Docker Engine.* This is a lightweight runtime environment for building and running Docker containers. The Docker Engine includes an in-host daemon (Linux service) that you communicate with by using the Docker client for building, deploying, and running containers.
- *Docker Compose.* This enables you to define a multiple-container app together with any dependencies so that you can run it with a single command. Docker Compose lets you specify the images your app will use with any necessary volumes or networks.
- *Docker Machine.* This enables you to provision Docker hosts by installing the Docker Engine on a computer in your datacenter or at a cloud provider. Docker Machine also installs and configures the Docker client so that it can talk with the Docker Engine.
- *Docker client.* This includes a command shell that is preconfigured as a Docker command-line environment.
- *Kitematic.* This GUI can help you to quickly build and run Docker containers and to find and pull images from the Docker Hub.
- *Docker Registry.* This open source app forms the basis for the Docker Hub and Docker Trusted Registry.
- *Docker Swarm.* This native clustering capability allows you to combine multiple Docker Engines into a single virtual Docker Engine.

You can download and install Docker software on various platforms, including Windows, Linux, and Mac OS X. After you install Docker software on your computer, you can then proceed to build images and tags and push or pull them to the Docker Hub.

Docker solutions

The software in the Docker toolbox is not all the Docker platform has to offer. The following Docker solutions are also key parts of what makes Docker so powerful for DevOps:

- **Docker Hub.** This is a cloud-hosted service where you can register your Docker images and share them with others.
- **Docker Trusted Registry.** This private, dedicated image registry lets you store and manage images on-premises or in a virtual private cloud.
- **Universal Control Panel.** You can use this to manage Docker apps regardless of whether they are running on-premises or within a virtual private cloud.
- **Docker Cloud.** With this cloud-hosted service, you can directly deploy and manage your Docker apps.
- **Docker Datacenter.** The latest addition to the stable of Docker solutions, Docker Datacenter is an integrated, end-to-end platform for deploying Containers as a Service on-premises or in a virtual private cloud. The self-service capabilities of Docker Datacenter make it easy for developers to build, test, deploy, and manage agile apps.

Usage scenarios

As organizations adopt containers, they will discover the challenge of deploying dozens, hundreds, or thousands of containers that make up an app. Tracking and managing deployment requires advanced management and orchestration.

Some common usage scenarios for Docker include:

- Container orchestration
- DevOps
- Microservices

DevOps

The Docker platform provides developers with tools and services that they can use to:

- Build and share images through a central repository of images.
- Collaborate on developing containerized apps by using version control.
- Manage infrastructure for apps.

Docker helps developer teams to rapidly build, test, deploy, and run distributed apps and services at any scale. Because containerizing apps eliminates the problem of troubleshooting issues with software dependencies and differences between host environments, Docker increases developer productivity and lets you quickly move apps from development to test to production, and you can easily roll apps back if further remediation is necessary.

Another significant achievement is that Docker for Windows supports volume mounting. This means that the container can see your code on your local device. To achieve this, the tool builds a connection between the container and the host. Effectively, this enables *edit and refresh* scenarios for development.

With Docker for Windows, you can now use Docker tools for Microsoft Visual Studio in the following scenarios:

- Docker assets for Debug and Release configurations are added to the project.
- A Windows PowerShell script is added to the project to coordinate the build and composition of containers, enabling you to extend them while keeping the Visual Studio designer experiences.
- With volume mapping configured, pressing the F5 key in Debug config starts the Windows PowerShell script to build and run your **docker-compose.debug.yml** file.
- F5 in Release config starts the Windows PowerShell script to build and run your **docker-compose.release.yml** file, which allows you to verify and push an image to your Docker Registry for deployment to another environment.

Microservices

Another benefit of using Docker containers is that they can be individually scaled and updated.

Microservices are an approach to app development where every part of an app deploys as a fully self-contained component. For example, the subsystem of an app that receives requests from the Internet might be separate from the subsystem that places the request into a queue for a backend subsystem that drops them into a database.

When an app is constructed by using microservices, each subsystem is a microservice. In a development or test environment on a single machine, microservices might each have one instance, but when the app runs in a production environment, each microservice can scale out to multiple instances, spanning across a cluster of servers based on resource demands.

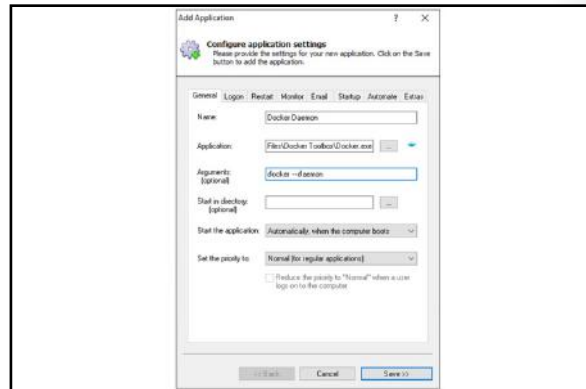
Some of the benefits of using Docker containers in this scenario include:


- Each microservice can quickly scale out to meet increased load.
- The namespace and resource isolation of containers also prevents one microservice instance from interfering with others.
- The Docker packaging format and app programming interfaces (APIs) unlock the Docker ecosystem for the microservice developer and app operator.

With a good microservice architecture, you can solve the management, deployment, orchestration, and patching needs of a container-based service with reduced risk of availability loss while maintaining high agility.

Installing and configuring Docker

Windows Server 2016 does not include the Docker Engine, and you to install and configure it separately. The steps to run the Docker Engine on the Windows OS vary from those on Linux. These instructions detail how to install and configure the Docker Engine on Windows Server 2016 and on Nano Server, in addition to instructions on how to install the Docker client. As a side note, the Docker Engine and the Docker client are now separate installations.



 **Note:** You must enable the Windows container feature before Docker can create and manage Windows containers. For information on enabling this feature, refer to the instructions to deploy Windows containers earlier in this module.

Installing Docker on Windows Server 2016

Because Docker is not included with Windows Server 2016, you must download it from the Internet to deploy the software. At an elevated Windows PowerShell command prompt, use the following procedure to install Docker:

1. Download **Dockerd.exe** from <https://aka.ms/tp5/dockerd>, and place it in the **%SystemRoot%\System32** directory on your container host. For example type the following, and then press Enter:

```
Invoke-WebRequest -Uri "https://aka.ms/tp5/dockerd" -OutFile
"$env:SystemRoot\System32\dockerd.exe"
```

2. Create a directory named **%ProgramData%\Docker**, and then create a file named **runDockerDaemon.cmd** within this directory. For example, type the following, and then press Enter:

```
New-Item -ItemType File -Path %ProgramData%\Docker\runDockerDaemon.cmd -Force
```

3. Using a text editor, copy the following text into the **runDockerDaemon.cmd** file, and then press Enter:

```
@echo off
set certs=%ProgramData%\docker\certs.d
if exist %ProgramData%\docker (goto :run)
mkdir %ProgramData%\docker
:run
if exist %certs%\server-cert.pem (if exist %ProgramData%\docker\tag.txt (goto :secure))
if not exist %systemroot%\system32\dockerd.exe (goto :legacy)
dockerd -H npipe://
goto :eof
:legacy
docker daemon -H npipe://
goto :eof
:secure
if not exist %systemroot%\system32\dockerd.exe (goto :legacysecure)
dockerd -H npipe:// -H 0.0.0.0:2376 --tlsverify --tlscacert=%certs%\ca.pem --tlscert=%certs%\server-cert.pem --tlskey=%certs%\server-key.pem
goto :eof
:legacysecure
docker daemon -H npipe:// -H 0.0.0.0:2376 --tlsverify --tlscacert=%certs%\ca.pem --tlscert=%certs%\server-cert.pem --tlskey=%certs%\server-key.pem
```

4. The NSSM installer creates and configures the Docker service. Download it from <https://nssm.cc/release/nssm-2.24.zip>. For example, type the following, and then press Enter:

```
Invoke-WebRequest -Uri "https://nssm.cc/release/nssm-2.24.zip" -OutFile "$env:ALLUSERSPROFILE\nssm.zip"
```

5. Extract **Nssm.exe** from the contents of the compressed package. For example, type the following, and then press Enter:

```
Expand-Archive -Path $env:ALLUSERSPROFILE\nssm.zip $env:ALLUSERSPROFILE
```

6. Copy **nssm-2.24\win64\nssm.exe** from the extracted location into the **%SystemRoot%\System32** directory. For example, type the following, and then press Enter:

```
Copy-Item $env:ALLUSERSPROFILE\nssm-2.24\win64\nssm.exe $env:SystemRoot\system32
```

7. Run **Nssm.exe install** to configure the Docker service. For example, type the following, and then press Enter:

```
Start-Process nssm install
```

8. In the **NSSM service installer** dialog window, on the **App** tab, type the following data in the appropriate field:

- Path: **C:\Windows\System32\cmd.exe**
- Startup Directory: **C:\Windows\System32**
- Arguments: **/s /c C:\ProgramData\docker\runDockerDaemon.cmd**
- Service Name: **Docker**

9. On the **Details** tab, type the following data in the appropriate field:
 - o Display name: **Docker**
 - o Description: The Docker daemon provides management capabilities of containers for Docker clients.
10. On the **IO** tab, type the following data in the appropriate field:
 - o Output (stdout): **C:\ProgramData\docker\daemon.log**
 - o Error (stderr): **C:\ProgramData\docker\daemon.log**
11. Click **Install service**.



Note: The Docker daemon is now registered as a Windows service and will start when Windows Server starts. If you wish to enable remote Docker management, you also need to open the inbound TCP port 2376.

Installing Docker on Nano Server

Similar to Windows Server 2016, Docker is not included with Nano Server, and you must download it from the Internet to deploy the software. At an elevated Windows PowerShell command prompt, use the following procedure to install Docker:

1. Download **Docker.exe** from <https://aka.ms/tp5/dockerd>, and then place it in the **%SystemRoot%\System32** directory on your container host that is running Nano Server. Because Nano Server does not support the **Invoke-WebRequest** cmdlet, you will need to download the file to another computer, and then copy it to your container host.
2. Create a directory named **%ProgramData%\Docker**, and then create a file named **runDockerDaemon.cmd** within this directory. For example, type the following, and then press Enter:

```
New-Item -ItemType File -Path %ProgramData%\Docker\runDockerDaemon.cmd -Force
```

3. Using a text editor from another computer, copy the following text into the **runDockerDaemon.cmd** file, and then press Enter:

```
@echo off
set certs=%ProgramData%\docker\certs.d
if exist %ProgramData%\docker (goto :run)
mkdir %ProgramData%\docker
:run
if exist %certs%\server-cert.pem (if exist %ProgramData%\docker\tag.txt (goto :secure))
if not exist %systemroot%\system32\dockerd.exe (goto :legacy)
dockerd -H npipe://
goto :eof
:legacy
docker daemon -H npipe://
goto :eof
:secure
if not exist %systemroot%\system32\dockerd.exe (goto :legacysecure)
dockerd -H npipe:// -H 0.0.0.0:2376 --tlsverify --tlscacert=%certs%\ca.pem --
tlscert=%certs%\server-cert.pem --tlskey=%certs%\server-key.pem
goto :eof
:legacysecure
docker daemon -H npipe:// -H 0.0.0.0:2376 --tlsverify --tlscacert=%certs%\ca.pem--
tlscert=%certs%\server-cert.pem--tlskey=%certs%\server-key.pem
```

4. Create a scheduled task, which will start the Docker daemon when the Nano Server starts. For example, type the following, and then press Enter:

```
# Creates a scheduled task to start docker.exe at computer start up.
$dockerData = "$($env:ProgramData)\docker"
$dockerDaemonScript = "$dockerData\runDockerDaemon.cmd"
$dockerLog = "$dockerData\daemon.log"
$action = New-ScheduledTaskAction -Execute "cmd.exe" -Argument "/c
$dockerDaemonScript > $dockerLog 2>&1" -WorkingDirectory $dockerData
$trigger = New-ScheduledTaskTrigger -AtStartup
$settings = New-ScheduledTaskSettingsSet -Priority 5
Register-ScheduledTask -TaskName Docker -Action $action -Trigger $trigger -Settings
$settings -User SYSTEM -RunLevel Highest | Out-Null
Start-ScheduledTask -TaskName Docker
```



Note: If you wish to enable remote Docker management, you also need to open the inbound TCP port 2376. You can use the **netsh** command to enable this port on Nano Server.

Installing the Docker client

The Docker client includes a command shell that is preconfigured as a Docker command-line environment (CLI), and it should install on the container host or any other system where you will run Docker CLI commands. At an elevated Windows PowerShell command prompt, use the following procedure to install the Docker client:

- Download **Docker.exe** from <https://aka.ms/tp5/docker>, and then place it in the **%SystemRoot%\System32** directory on your container host or another system for running the CLI. For example, type the following, and then press Enter:

```
Invoke-WebRequest -Uri "https://aka.ms/tp5/docker" -OutFile
"$env:SystemRoot\system32\docker.exe"
```



Note: While Nano Server does not support the **Invoke-WebRequest** cmdlet, you will need to download the file to another computer, and then copy it to your Nano Server container host.

Overview of management with Docker

As mentioned earlier in the module, you can choose to manage Windows containers by using Windows PowerShell or Docker. The advantage of using Docker on Windows Server is that Docker is an industry-standard container deployment and management tool that enables administrators with skills in container management on other OSs to manage your Windows containers.

With Docker, you can create containers, remove containers, manage containers, and browse the Docker Hub to access and download prebuilt images.

- With Docker, you can:
 - Create containers
 - Remove containers
 - Manage containers
 - Browse the Docker Hub to access and download prebuilt images
- In most organizations, the most common Docker management tasks include:
 - Automating the creation of container images by using Dockerfile on Windows OS
 - Managing containers by using Docker
 - Using **docker run**

In most organizations, the most common management tasks that use Docker include:

- Automating the creation of container images by using **Dockerfile** on Windows OSs.
- Managing containers by using Docker.
- Using **docker run**.

Automating the creation of container images by using Dockerfile on Windows

The Docker Engine includes tools for automating the creation of container images. While you can create container images manually, adopting an automated image creation process provides many benefits, including:

- Storing container images as code.
- Rapid and exact recreation of container images for maintenance and upgrade purposes.
- Continuous integration between container images and the development cycle.

The Docker components that are part of this automation are the **Dockerfile** and the **docker build** command.

- **Dockerfile.** This text file contains the instructions needed to create a new container image. The instructions include the name of an existing image to use as a base, the commands to run during the image creation process, and a command that will run when new instances of the container image are deployed.
- **Docker build.** This Docker Engine command consumes a **Dockerfile**, and then triggers the image creation process.

In its most basic form, a **Dockerfile** can be very simple.

The following example creates a new image, which includes IIS, and a "Hello world" site.

Dockerfile sample


```
# Indicates that the windowsservercore image will be used as the base image
FROM windowsservercore

# Metadata indicating an image maintainer.
MAINTAINER LJackman@adatum.com


# Uses dism.exe to install the IIS role.
RUN dism.exe /online /enable-feature /all /featurename:iis-webserver /NoRestart

# Creates an html file and adds content to this file.
RUN echo "Hello World - Dockerfile" > c:\inetpub\wwwroot\index.html

# Sets a command or process that runs each time a container is run from the new image.
CMD [ "cmd" ]
```

 **Additional Reading:** For more information on other examples of Dockerfiles for Windows, refer to the Dockerfile for Windows Repository at: <http://aka.ms/sqic5n>

Dockerfile instructions provide the Docker Engine with the steps necessary to create a container image. These instructions perform in order, one by one.

 **Additional Reading:** For more information on the complete list of **Dockerfile** instructions, refer to **Dockerfile** at: <http://aka.ms/Ncu9g5>

You can also specify Windows PowerShell commands to run in a **Dockerfile** by using the **RUN** operation. The following is a list of options for using Windows PowerShell commands in a **Dockerfile**:

- You can use Windows PowerShell and the **Invoke-WebRequest** command to gather information or files from a web service. For example, you could download the Python programming language from the vendor's website for installation in a new image.
- You can use Windows PowerShell to download files during the image creation process by using the .Net WebClient library. This method is shown to increase download performance.



Note: Nano Server does not currently support the .Net WebClient.

- You might consider it helpful to copy a Windows PowerShell script into the containers being used during the image creation process, and then run the script from within the container. For example, the following command copies a Windows PowerShell script from the build machine into the container by using the **ADD** instruction, and it then runs the script by using the **RUN** instruction. Type the following, and then press Enter:

```
FROM windowsservercore
ADD script.ps1 /windows/temp/script.ps1
RUN powershell.exe -executionpolicy bypass c:\windows\temp\script.ps1
```

After you create a **Dockerfile** and save it to disk, you can use **docker build** to create the new image. The **docker build** command takes several optional parameters and a path to the **Dockerfile**. For example, the following command creates an image named **IIS**:

```
docker build -t iis .
```



Additional Reading: For more information on **docker build**, including a list of all build options, refer to: "Docker Build" at Docker.com, at: <http://aka.ms/Qmswzf>



Additional Reading: You can use several methods to optimize the Docker build process and the resulting Docker images. For more information on how the Docker build process operates and the tactics that you can use for optimal image creation with Windows containers, refer to: "Optimize Windows Dockerfiles" at: <http://aka.ms/f29xln>

Managing containers by using Docker

You can use Docker to support a container environment. After you install Docker, use the following commands to manage your containers:

- **Docker images.** This lists the installed images on your container host. As you might recall, you use container images as a base for new containers. The Windows PowerShell cmdlet equivalent to this command is **Get-ContainerImage**.
- **Docker run.** This creates a container by using a container image. For example, the following command creates a container named **IIS** based on the Windows Server Core container image:

```
docker run --name IIS -it windowsservercore
```

The Windows PowerShell cmdlet equivalent to this command is **New-Container**.

- **Docker commit.** This commits the changes you made to the container and creates a new container image. For example, the following command creates a new container image named **WinSvrCoreIIS** based on the **IISBase** base image.

```
docker commit iisbase WINSVRCOREIIS
```
- **Docker stop.** This stops a container. The Windows PowerShell cmdlet equivalent to this command is **Stop-Container**.
- **Docker rm.** This removes a container. The Windows PowerShell cmdlet equivalent to this command is **Remove-Container**.



Additional Reading: For more information about administering containers on Windows Server by using Docker, refer to: "Windows Containers Quick Start" at: <http://aka.ms/xl3mdn>

Using docker run

The **docker run** command is the most commonly used Docker command. As part of creating a container, you can use this command to define the container's resources at runtime.

Docker runs processes in isolated containers. These containers are simply a process that is running on a host. The host might be local or remote. When you run **docker run**, the container process that runs is isolated; for example, a separate file system, networking, and process tree from the host. During execution, the **docker run** command must specify an image from which to derive the container. When developing an image, you can define image defaults that relate to:

- Detached or foreground running.
- Container identification.
- Network settings.
- Runtime constraints on CPU and memory.

With **docker run**, you can add to or override the image defaults that were configured during image development. Additionally, you can override nearly all the defaults that the Docker runtime itself set. The ability to override image and Docker runtime defaults is why **docker run** has more options than any other Docker command.

The optional **cmd** switch opens an interactive session with the container to include a default command or other options. While you might have provided a default **COMMAND** by using the **Dockerfile CMD** instruction during image creation, you can override that **CMD** instruction when running a container from the image by specifying a new **COMMAND**. The Windows PowerShell cmdlet equivalent to this command is **Start-Container**. To end an interactive session with the container, type **exit**.



Note: You can now use Windows PowerShell commands to install the required roles and features within a container. For example, **powershell.exe Install-WindowsFeature web-server** installs the IIS components within your container.


The Docker platform simplifies the experience of working across container options. An app that was developed by using Windows Server containers can deploy as a Hyper-V container without change. In fact, managing Hyper-V containers with Docker is almost identical to managing Windows Server containers with Docker. When you create a Hyper-V container by using **docker run**, you would include the **--isolation=hyperv** parameter.

For example, the following command starts a Windows Server container and hosts a long-running ping process:

```
docker run -d windowsservercore ping localhost -t
```

In contrast, this example also starts a Hyper-V container and hosts a long-running ping process:

```
docker run -d --isolation=hyperv nanoserver ping -t localhost
```

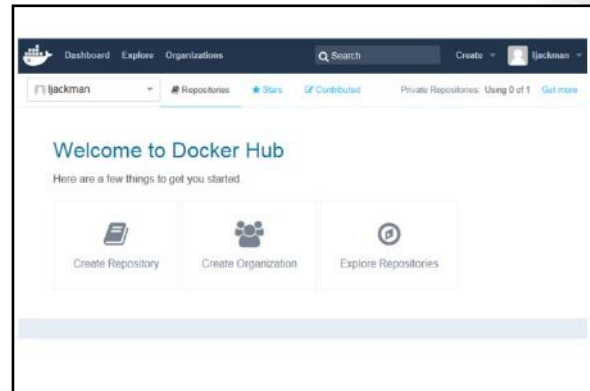
 **Additional Reading:** For more information on using the **docker run** command to define a container's resources at runtime, refer to the Docker run reference at: <http://aka.ms/S2oofi>

Overview of Docker Hub

Docker Hub is a cloud-based public registry service that the Docker company maintains for building and shipping app or service containers. It provides a centralized resource for container image discovery, distribution and change management, user and team collaboration, and workflow automation throughout the development pipeline.

Docker Hub features

Docker Hub provides the following major features and functions:



- **Image repositories.** Docker Hub contains images stored in repositories that you can find, manage, and push and pull images from community, official, and private image libraries to build containers. Specifically, repositories contain images, layers, and metadata about those images. The main concept of containerization is that you can build your own images based on existing images—this is known as *layers*.
- **Organizations and teams.** One useful aspect of private repositories is that you can share them with members of your organization or team only. Docker Hub lets you create organizations, or work groups that require user access, where you can collaborate with colleagues and manage private repositories.
- **Automated builds.** Automated builds automates the building and updating of images from GitHub or Bitbucket, directly on Docker Hub. It works by adding a commit hook to your selected GitHub or Bitbucket repository, triggering a build and update when you push a commit or when you make changes to the source repository.
- **Webhooks.** A webhook is a feature of automated builds that attaches to your repositories and allows you to trigger an event or action when an image or updated image successfully pushes to the repository. With a webhook, for example, you can specify a target URL and a JavaScript Object Notation (JSON) payload to deliver when you push the image.
- **GitHub and Bitbucket integration.** This allows you to add the Docker Hub and your Docker images to current workflows.

Working with image repositories

Docker Hub repositories provide a place to build and ship Docker images. These repositories enable you to share images with coworkers, customers, or the Docker community at large. You can configure Docker Hub repositories in two ways:

- **Repositories.** Repositories allow you to push images at will from your local Docker daemon to the Docker Hub. If you are building images internally, either on your own Docker daemon or by using your own continuous integration services, you can push the images to a Docker Hub repository, which you add to your Docker Hub user or organizational account.
- **Automated builds.** Automated builds allow you to configure GitHub or Bitbucket to trigger the Docker Hub to rebuild repositories when changes to the repository occur. If the source code for your Docker image is on GitHub or Bitbucket, you can use an Automated Build repository, which Docker Hub services build.



Note: You can create public repositories that any other Docker Hub user can access, or you can create private repositories with limited access that you can manage.

Docker provides access to Docker Hub services by using four primary Docker Engine CLI commands: **docker login**, **docker search**, **docker pull**, and **docker push**.

Creating an account for sign-in

If you have not already done so, you will need to create a Docker ID to interact with your repositories on Docker Hub. To do this through Docker Hub, refer to <http://aka.ms/q7bz4b>. After you have a Docker ID, sign in to your account from the CLI. Type the following command, and then press Enter:

```
$ docker login
```

After signing in from the command line, you can use the other Docker Engine subcommands.



Additional Reading: For more information on registering a Docker ID, refer to: "Your Docker ID" at: <http://aka.ms/mudoqr>

Searching for images

You can search for public repositories and images that are available on the Docker Hub in two ways. You can use the **Search** feature on the Docker Hub website, or you can use the **docker search** command from the CLI. Both of these methods will show you a list of the currently available images that match the provided keywords from the public repositories on the Docker Hub.



Note: Images from a private repository do not appear in the repository search results list. To see all the repositories that you can access and their status, view the dashboard on the Docker Hub website.

Image search results are based on criteria such as image name, user name, or description. Using the search criteria "CentOS" returns the following search results from all the repositories and images:

NAME	DESCRIPTION	STARS	OFFICIAL	AUTOMATED
centos	The official build of CentOS.	1034	[OK]	
ansible/centos7-ansible	Ansible on Centos7	43		[OK]
tianon/centos	CentOS 5 and 6, created wi...	13		[OK]
...				

In the preceding example, your search returned three results: **centos**, **ansible/centos7-ansible**, and **tianon/centos**. The second and third results signal that they originate from public repositories named **ansible** and **tianon** respectively; the forward slash (/) character separates a user's repository name from the image name. On the other hand, the first result, **centos**, does not explicitly list a repository. In this latter scenario, this means that the image originates from the trusted, top-level namespace for Official Repositories from Docker Hub.

Docker Hub contains a number of *Official Repositories*. These are public, certified repositories from vendors and contributors to Docker that you can use to build apps and services. With Official Repositories, you know that you are using an optimized and up-to-date image that experts built to power your apps.



Additional Reading: For more information on the Docker repositories that the Docker Hub supports and promotes, refer to: "Official Repositories" on Docker Hub at: <http://aka.ms/wyoip4>

After you locate the image you want, you can download it with the **docker pull** command from the CLI. In the following example, you download the latest image from which you can run your containers:

```
$ docker pull centos
Using default tag: latest
latest: Pulling from library/centos
f1b10cd84249: Pull complete
c852f6d61e65: Pull complete
7322fbe74aa5: Pull complete
Digest: sha256:90305c9112250c7e3746425477f1c4ef112b03b4abe78c612e092037bfec3b7
Status: Downloaded newer image for centos:latest
```

As noted earlier, image repositories contain images, layers, and metadata about those images. Part of this metadata is a tag that you can use to label an image. For example, you could use the following command to download version 5 of centos (**centos5** is the tag labeling an image in the centos repository for a version of CentOS):

```
docker pull centos:centos5
```

Contributing to Docker Hub

While anyone can download public images from the Docker Hub registry, you must register if you want to share your own images.



Additional Reading: For more information on pushing a repository to the Docker Hub registry, refer to: "Build your own images" at: <http://aka.ms/F9rae2>

For example, you can push this repository and upload your image so that it is available for your teammates and the Docker Hub community to use:

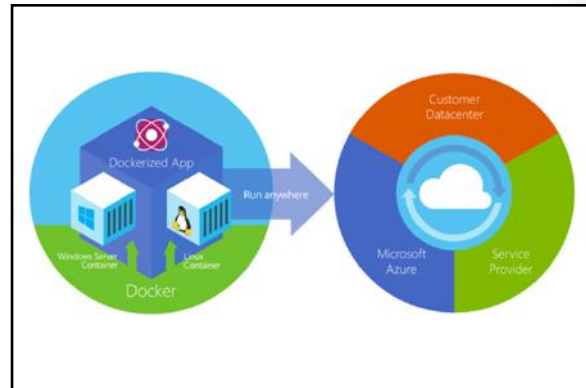
```
$ docker push Docker ID/Image Name
```



Additional Reading: For more information on creating organizations and teams so that you can delegate access to colleagues for shared image repositories, refer to: "Organizations and teams" at: <http://aka.ms/Vr7xie>

Docker with Azure

As an open source engine, Docker automates the deployment of any app as a portable, self-sufficient container that runs almost anywhere—including Azure. Typical VM images contain everything necessary to run, including the app, any dependencies, and the OS. In contrast, Docker containers include the app and some libraries, but the OS and common dependencies remain shared assets. Consequently, Docker containers are extremely lightweight compared with VM images. By making Docker containers significantly smaller than traditional VMs, more of them can run on a single host, they can start more quickly, and they are considerably more portable—these characteristics are ideal for PaaS, such as Azure.



With Azure, you have the flexibility to deploy Docker in a few different scenarios based on your needs:

- Docker Machine Azure driver to deploy Docker hosts within Azure
- Azure Docker VM extension for template deployments
- Deploy an Azure container service cluster

Using Docker Machine Azure driver to deploy Docker hosts on Azure

You can use the Docker Machine Azure driver to deploy Docker hosts on Azure. Docker is one of the most popular virtualization approaches that use Linux containers rather than VMs as a way of isolating app data and computing on shared resources. One common scenario that uses this approach is when you need to prototype an app quickly.

You can create Docker host VMs on Azure by using the **docker-machine create** command with the **-d azure** Azure driver option. For example, you can use the following command for testing a web app; the command creates a new VM named **DockerVM**, opens port 80 to the Internet on the VM, and it enables ops as the sign-in user for Secure Shell (SSH):

```
docker-machine create -d azure \
  --azure-ssh-user ops \
  --azure-subscription-id Azure_Subscription_ID \
  --azure-open-port 80 \
  machine
```



Additional Reading: For more information on using Docker Machine to create new Docker host VMs in Azure for your Linux containers, refer to: "Use Docker Machine with the Azure Driver" at: <http://aka.ms/Qdpt98>

Using the Azure Docker VM extension for template deployments

For a template-based deployment, you can use the Docker VM extension for Azure VMs. This approach allows you to integrate with Azure Resource Manager template deployments and includes all of the related benefits such as role base access, diagnostics, and post-deployment configuration. The Azure Docker VM extension installs and configures the Docker daemon, the Docker client, and Docker Compose in your Linux VM.

You can also use the extension to define and deploy container apps by using Docker Compose. Azure Resource Manager templates enable you to deploy a solution throughout the development lifecycle and have confidence that your resources deployed in a consistent state. Using the Azure Docker VM extension is well suited for robust development or production environments because you have some additional control compared with simply using Docker Machine or manually creating a Docker host.

You can use Azure Resource Manager to create and deploy templates that define the entire structure of your environment, such as the Docker hosts, storage, role-based access controls (RBACs), and diagnostics. The advantage of using Resource Manager templates over simply using Docker Machine is that you can define additional Docker hosts, storage, and access controls, and you can reproduce deployments as needed.



Additional Reading: For more information, refer to: "Azure Resource Manager overview" at: <http://aka.ms/ltdpqg>

Deploying an Azure Container Service cluster

The Azure Container Service provides rapid deployment of popular open source container clustering and orchestration solutions. With Azure Container Service, you can deploy clusters such as a Docker Swarm cluster by using Azure Resource Manager templates or the Azure portal. Docker Swarm clusters are ideal for production-ready, scalable deployments that take advantage of the additional scheduling and management tools that Docker Swarm provides.

Docker Swarm uses the native Docker API to provide an environment for deploying containerized workloads across a pooled set of Docker hosts. During the deployment of Docker Swarm clusters, you will use the Azure compute resource and Azure VM Scale Sets to manage a collection of VMs as a set.



Additional Reading: For more information on using the Azure Container Service to deploy Docker Swarm clusters, refer to: "Deploy an Azure Container Service cluster" at: <http://aka.ms/Yz210g>

Question: Verify the correctness of the statement by placing a mark in the column to the right.

Statement	Answer
Docker is a graphical management tool that you can use to manage Hyper-V containers in Windows Server 2016.	

Demonstration: Deploying containers by using Docker

20743A-LON-HOST1 and **20743A-LON-DC1-B** are required to complete this demonstration. These should be running from the last demonstration. In order to successfully complete this demonstration, you need to complete the previous demonstration.

Demonstration Steps

Install base OS images

- Open an elevated **Command Prompt** window, type the following commands, and then press Enter:

```
Docker images
```

Download a prebuilt Microsoft and IIS Docker image

1. In the **Command Prompt** window, type the following command, and then press Enter:

```
Docker search Microsoft
```

2. In the **Command Prompt** window, type the following command, and then press Enter:

```
docker pull microsoft/iis
```

3. In the **Command Prompt** window, type the following commands, and then press Enter:

```
docker images
```

Deploy a new container with the prebuilt image

- In the **Command Prompt** window, run the following command:

```
docker run -d -p 80:80 microsoft/iis ping -t localhost
```



Note: This command runs the IIS image as a background service (**-d**). It also configures networking such that port 80 of the container host maps to port 80 of the container.

Manage the container

- In the **Command Prompt** window, run the following command to view the running containers:

```
docker ps
```


Module Review and Takeaways

Review Questions

Question: When creating a virtual hard disk for Nano Server by using the **New-NanoServerImage** Windows PowerShell cmdlet, when do you use the **-Guestdrivers** switch?

Question: When using the Nano Server Recovery Console, which two fundamental components can you configure?

Question: When configuring Windows Server containers, what Windows PowerShell cmdlet do you use to create a container and what is the equivalent Docker command?

MCT USE ONLY. STUDENT USE PROHIBITED

Module 11

Implementing failover clustering

Contents:

Module Overview	11-1
Lesson 1: Overview of failover clustering	11-2
Lesson 2: Implementing a failover cluster	11-18
Lesson 3: Configuring highly available applications and services on a failover cluster	11-24
Lesson 4: Maintaining a failover cluster	11-30
Lesson 5: Implementing a stretch cluster	11-36
Lab: Implementing failover clustering	11-43
Module Review and Takeaways	11-49

Module Overview

Providing high availability is very important for any organization that wants to provide continuous services to its users. Failover clustering is one of the main technologies in Windows Server 2016 that can provide high availability for various applications and services. In this module, you will learn about failover clustering, failover clustering components, and implementation techniques.

Objectives

After completing this module, you will be able to:

- Describe failover clustering.
- Implement a failover cluster.
- Configure highly available applications and services on a failover cluster.
- Maintain a failover cluster.
- Implement a stretch cluster.

Lesson 1

Overview of failover clustering

Failover clusters in Windows Server 2016 provide a high-availability solution for many server roles and applications. By implementing failover clusters, you can maintain application or service availability if one or more computers in the failover cluster fail. Before you implement failover clustering, you should be familiar with general high-availability concepts. You must be familiar with clustering terminology and understand how failover clusters work. It is also important to be familiar with the new clustering features in Windows Server 2016.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe availability.
- Describe the failover clustering improvements in Windows Server 2012 R2.
- Describe the failover clustering improvements in Windows Server 2016.
- Describe the failover cluster components.
- Define failover and failback.
- Describe failover cluster networks.
- Describe failover cluster storage.
- Describe a quorum.
- Describe the quorum modes in Windows Server 2016 failover clustering.
- Describe how a quorum works in Windows Server 2016 failover clustering.
- Describe Cluster Shared Volumes (CSVs).
- Describe the CSV improvements in Windows Server 2016.

What is availability?

Availability refers to a level of service that applications, services, or systems provide and is described as the percentage of time that a service or system is available. Highly available systems have minimal downtime—whether planned or unplanned—depending on the needs and the budget of the organization. For example, a system that is unavailable for 8.75 hours per year would have a 99.9 percent availability rating.

To improve availability, you must implement mechanisms that mask or minimize how failures of the service's components and dependencies affect the system. You can achieve fault tolerance by implementing redundancy to single points of failure.

- Availability is a level of service expressed as a percentage of time
- Highly available services or systems are available more than 99 percent of the time
- High-availability requirements differ based on how availability is measured
- Planned outages typically are not included when calculating availability

The information technology (IT) department must disclose the availability requirements to the customer so that no misunderstanding occurs about the implications. When you need higher availability, the infrastructure requires more compute, memory, and storage. This might equate to infrastructure and continued yearly support costs. Miscommunication about service level expectations between the customer and the IT organization might result in poor business decisions, such as unsuitable investment levels and customer dissatisfaction.

The availability measurement period might also have a significant effect on the definition of high availability. For example, a requirement for 99.9 percent availability over a one-year period allows for 8.75 hours of downtime for the whole year, whereas a requirement for 99.9 percent availability over a rolling four-week period allows for only 40 minutes of downtime every four weeks.

You also have to identify and negotiate planned outages, maintenance activities, service pack updates, and software updates. For the purpose of defining service level agreements (SLAs), these are scheduled outages and typically are not included as downtime when calculating the system's availability. You typically calculate availability based only on unplanned outages. However, you have to negotiate exactly which planned outages you consider downtime.

Failover clustering improvements in Windows Server 2012 R2

Failover clustering in Windows Server 2012 R2 is enhanced with many new features, and existing technologies have been updated for better functionality. The quorum model is significantly changed, and you now have many more options and flexibility to maintain the quorum and cluster. In addition, the **Failover Cluster Manager** console in Windows Server 2012 R2 has a cluster dashboard where you can quickly view the health status of all managed failover clusters. In the **Failover Cluster Manager** console, next to each failover cluster that you manage are icons that indicate whether a cluster is running, the number and status of clustered roles, the node status, and the event status.

Significant new features of failover clustering in Windows Server 2012 R2:

- Quorum changes and dynamic witness
- Force quorum resiliency
- Tie breaker for 50% node split
- Global Update Manager mode
- Cluster node health detection
- Active Directory-detached cluster

The most important new features in the failover clustering quorum in Windows Server 2012 R2 are the following:

- **Dynamic quorum.** This feature enables a cluster to recalculate the quorum in the event of node failure and still maintain working clustered roles, even when the number of voting nodes remaining in the cluster is less than 50 percent.
- **Dynamic witness.** This feature dynamically decides if the witness has a vote to maintain the quorum in the cluster.
- **Force quorum resiliency.** This feature provides additional support and flexibility to manage split-brain syndrome cluster scenarios. *Split-brain syndrome* refers to 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 automatically adjust the running node's vote status to keep the total number of votes in the cluster at an odd number.

These new quorum options and modes of work are discussed in more detail later in this lesson.

In addition to updating the quorum, Microsoft has made other valuable changes to failover clustering. The remainder of this topic details the most important changes in Windows Server 2012 R2 failover clustering.

Global Update Manager mode

The Global Update Manager is responsible for updating the cluster database. In Windows Server 2012, it is not possible to configure how these updates work. Windows Server 2012 R2 allows you to configure Global Update Manager updates.

Every time the state of a cluster changes, such as when a cluster resource is offline, all the 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, the Global Update Manager works in Majority (read and write) mode. In this mode, when a change happens to a cluster, a majority of the cluster nodes must receive and process the update before it is committed to the database. When the cluster node wants to read the database, the cluster compares the latest time stamp from a majority of the running nodes and uses the data with the latest time stamp.

In Windows Server 2012 R2, the Global Update Manager can also work in the All (write) and Local (read) mode. When working in this mode, all the nodes in the cluster must receive and process an update before it is committed to the database. However, when the cluster receives the database read request, the cluster reads the data from the database copy that is locally stored. Because all roles receive and process the update, the Global Update Manager considers the local cluster database copy a relevant source of information.

Windows Server 2012 R2 also supports a third mode for the Global Update Manager. This mode is Majority (write) and Local (read). In this mode, a majority of the cluster nodes must receive and process an update before it is committed to the database. When the database read request is received, the cluster reads the data from the database copy that is locally stored.

In Windows Server 2012 R2, the default setting for Microsoft Hyper-V Server 2012 R2 failover clusters is Majority (read and write). All other workloads in the clusters use All (write) and Local (read) mode. Majority (write) and Local (read) is not used by default for any workload.

Changing the working mode for the Global Update Manager improves cluster database performance and increases the performance of cluster workloads, because a cluster database will no longer need 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 5 seconds. In Windows Server 2012 R2, specifically for Hyper-V failover clusters, the default threshold value increases from 5 seconds to 10 seconds if nodes are in the same subnet, and it increases to 20 seconds if nodes are in different subnets. This provides increased resiliency for temporary network failures for virtual machines (VMs) that run on a Hyper-V cluster, and this delays cluster recovery actions in cases of short network interruptions.

AD DS–detached cluster

You integrate failover clusters with Active Directory Domain Services (AD DS), and you cannot deploy a cluster if nodes are not members of same domain. When you create a cluster, appropriate computer objects for the cluster name and the clustered role name are created in AD DS.

In Windows Server 2012 R2, you can deploy an *AD DS–detached cluster*, which is a cluster that does not have dependencies in AD DS for network names. When you deploy clusters in detached mode, you register the cluster network name and the network names for clustered roles in a local Domain Name

System (DNS), but corresponding computer objects for the cluster and the clustered roles are not created in AD DS.

Cluster nodes still have to be joined to the same Active Directory domain, but the person that creates a cluster does not need to have permission to create new objects in AD DS. In addition, you do not need later management of these computer objects.

The deployment of AD DS–detached clusters also has side effects. Because you do not create computer objects, you cannot use Kerberos authentication when accessing cluster resources. Although you use Kerberos authentication between cluster nodes, because they have their computer accounts and objects created outside the cluster, NTLM authentication is used. Therefore, we recommend that you do not deploy Active Directory–detached clusters for any scenario that requires Kerberos authentication.

To create an AD DS–detached cluster, you must run Windows Server 2012 R2 or later on all cluster nodes. You cannot configure these features by using the Failover Cluster Manager, so you must use Windows PowerShell.

Failover clustering improvements in Windows Server 2016

Failover clustering in Windows Server 2016 continues to build on earlier versions of Windows Server. The most significant changes in Windows Server 2016, including Microsoft Azure Cloud Witness options, storage replication, and site-aware failover clusters, relate to stretch cluster configurations. Microsoft continued to make the challenges introduced by stretch clusters easier to handle by creating features such as the dynamic quorum, introduced in Windows Server 2012 R2. Windows Server 2016 provides some improvements to the overall stability and versatility of the cluster upgrade processes and new features in implementing the Hyper-V role in Windows Server 2016 on a failover cluster.

Failover clustering improvements in Windows Server 2016:

- Cluster operating system rolling upgrades
- Storage replica
- Azure Cloud Witness
- VM resiliency
- Site-aware failover clusters
- Work group and multidomain clusters

Cluster operating system rolling upgrades

When you upgrade from Windows Server 2012 R2 to Windows Server 2016, the upgrade path is different for Scale-Out File Servers and Hyper-V clusters than with earlier version upgrades. In earlier Windows Server versions, you had to create a new cluster with the new operating system on all nodes, and then you had to move the clustered services over to the new cluster. This created downtime and made it difficult to quickly complete upgrades. With Windows Server 2016, rather than taking downtime, you can take one node at a time offline, upgrade the operating system, and bring the node back into the cluster. This changes the cluster to a mixed operating system mode. After you upgrade all nodes, you can change the functional level of the entire cluster from the Windows Server 2012 R2 version to Windows Server 2016 version.

Storage Replica

The Storage Replica feature is an expansion on the Windows Server 2012 R2 and Windows Server 2012 storage software offerings in Windows Server. Using Server Message Block (SMB) 3.1.1, Storage Replica provides block-level replication for any type of data in complete volumes. This allows the disaster recovery of stretch cluster, cluster-to-cluster, or server-to-server situations. Even though Storage Replica can technically replicate the volume data for all applications, we recommend that you use the best practices

for those applications that provide replication functionality, especially applications such as AD DS, Microsoft Exchange Server, Hyper-V, or Microsoft SQL Server.

Azure Cloud Witness

In Windows Server 2016, you can now use an Azure Cloud Witness share to create a quorum. Previously when creating a stretch cluster, a third offsite quorum was required. Now with Windows Server 2016, you can create an Azure Cloud Witness instead.

VM resiliency

Windows Server 2016 introduces three new states to increase the resilience of the failover cluster during a transient failure. These new states are Unmonitored, Isolated, and Quarantined. This provides some new logic around how Hyper-V can handle transient failures, such as storage and specific node disconnects from the cluster.

Site-aware failover clusters

In Windows Server 2016, it is important to consider also building out site-aware failover clusters. This allows more control over stretch configurations. This includes what nodes the roles fail over to and failover affinity, including handling the CSV loads more accurately for site-specific configurations. It also gives you the ability to configure different heartbeat thresholds for the servers at different sites. Identifying the sites for each cluster node adds a level of control that alleviates some of the concerns for stretch clusters.

Workgroup and multiple-domain clusters


Windows Server 2016 supports failover clusters with different domains and clusters where not all nodes are in a domain. At this time, this configuration comes with many caveats, including requiring a local administrator on all servers with the same user name and password.

Failover cluster components

A *failover cluster* is a group of computers that work together to increase the availability of applications and services. Physical cables and software connect the clustered servers, known as *nodes*. If one of the cluster nodes fails, another node begins to provide service. We refer to this process as *failover*. With failover, users should experience a minimum amount of service disruptions.

A failover clustering solution consists of several components, which include:

- **Nodes.** *Nodes* are computers that are members of a failover cluster. These computers run the Cluster service and any resources and applications associated to the cluster.
- **A network.** Cluster nodes can communicate with one another and with clients across this network. You can use three types of networks in a cluster. The "Failover cluster networks" topic discusses these networks in more detail.
- **A resource.** A *resource* is an entity that a node hosts. The Cluster service manages the resource and can start, stop, and move to another node.



- Cluster storage. *Cluster storage* is a storage system that cluster nodes share. In some scenarios, such as clusters of servers that run Exchange Server, shared storage is not required.
- Clients. *Clients* are computers (or users) that use the Cluster service.
- A service or application. This is a software entity that Microsoft presents to clients and that clients use.
- A witness. A *witness* can be a file share or a disk, which you use to maintain the quorum. Ideally, the witness is located on a network that is both logically and physically separate from those used by the failover cluster. However, the witness must remain accessible by all cluster node members. Later lessons discuss the concepts of the quorum and the witness in more detail.

In a failover cluster, each node in the cluster:

- Has full connectivity and communication with the other nodes in the cluster.
- Is aware when another node joins or leaves the cluster.
- Connects to a network through which client computers can access the cluster.
- Connects through a shared bus or Internet Small Computer System Interface (iSCSI) connection to shared storage.
- Is aware of the services or applications that run locally and the resources that run on all other cluster nodes.

Cluster storage usually refers to logical devices—typically drives or logical unit numbers (LUNs)—that all the cluster nodes attach to through a shared bus. This bus is separate from the bus that contains the system and boot disks. The shared boot disks store resources such as applications and file shares that the cluster will manage.

A failover cluster typically defines at least two data communication networks: one network allows the cluster to communicate with clients, and the second, isolated network allows the cluster node members to communicate directly with one another. If directly connected shared storage is not being used, a third network segment (for iSCSI or Fibre Channel) can exist between the cluster nodes and a data storage network.

Most clustered applications and their associated resources are assigned to one cluster node at a time. The node that provides access to those cluster resources is the active node. If the nodes detect the failure of the active node for a clustered application, or if the active node is offline for maintenance, the clustered application starts on another cluster node. To minimize the impact of the failure, client requests automatically redirect to an alternative node in the cluster as quickly as possible.


What are failover and failback?

A failover transfers the responsibility of providing access to resources in a cluster from one node to another. A failover can occur when unplanned downtime of one node happens because of hardware failure. In addition, a service failure on an active node can initiate a failover to another node. A failover also occurs when an administrator intentionally moves resources to another node for maintenance.

A failover attempt consists of the following steps:

1. The Cluster service takes all the resources in the instance offline in an order that the instance's dependency hierarchy determines. That is, dependent resources go first, followed by the resources on which they depend. For example, if an application depends on a physical disk resource, the Cluster service takes the application offline first, which enables the application to write changes to the disk before the disk is offline.
2. After all the resources are offline, the Cluster service attempts to transfer the instance to the node that is listed next on the instance's list of preferred owners.
3. If the Cluster service successfully moves the instance to another node, it attempts to bring all the resources online. This time, it starts in the reverse order of the dependency hierarchy. In the preceding example, the Cluster service attempts to bring the disk back online first, followed by the application. The failover is complete when all the resources are online on the new node.

- During failover, the clustered instance and all associated resources are moved from one node to another
- A failover occurs when:
 - The node that hosts the instance becomes inactive for any reason
 - One of the resources within the instance fails
 - An administrator makes a switchover
- The cluster service can fail back after the offline node becomes active again
- Both planned and unplanned failovers can occur

 **Note:** Exceptions to this rule exist. For instance, after 2012 R2 when VM roles fail over, the role does not go offline but instead writes to both the source location and the destination of the resource owner until the failover is complete. It then moves the I/O to the new failover cluster node. Module 12, "Implementing failover clustering with Windows Server 2016 Hyper-V," covers this in more depth.

For instances that were originally hosted on the offline node, you can preconfigure a Cluster service to use failback after the offline node becomes active again. When the Cluster service uses failback for an instance, it uses the same procedure that it performs during failover. That is, the Cluster service takes all the resources in the instance offline, moves the instance, and then brings all the resources in the instance back online.

Planned vs. unplanned failover

A failover cluster completes the preceding steps in a planned failover. In an unplanned failover, the steps for failback are the same as for a planned failover. However, an unplanned failover usually occurs when one node goes offline without any planning—therefore, the services abruptly shut down on the node that owns them during the failure. This causes the Failover Cluster Manager to skip to step 3, and other nodes attempt to bring the offline services back online as quickly as possible.

Failover cluster networks

Networks and network adapters are important parts of each cluster implementation. You cannot configure a cluster without configuring the networks that the cluster will use. A network can perform one of the following roles in a cluster:

- **Private network.** A private network carries internal cluster communication. By using this network, cluster nodes exchange heartbeats and check for other nodes. The failover cluster authenticates all internal communication. However, administrators who are especially concerned about security might want to restrict internal communication to physically security-enhanced networks.
- **Public network.** A public network provides client systems with access to cluster application services. You create IP address resources on networks that provide clients with access to the Cluster service.
- **Public-and-private network.** A public-and-private network (also known as a *mixed network*) carries internal cluster communication and connects clients to cluster application services.

Network	Description
Public network	Clients use this network to connect to the clustered service
Private network	Nodes use this network to communicate with each other
Public-and-private network	Required to communicate with external storage systems

- One network can support both client and node communications
- Multiple network cards are recommended to provide enhanced performance and redundancy
- iSCSI storage should have a dedicated network

When you configure networks in failover clusters, you might also need to dedicate a network to connect to the shared storage. If you use iSCSI for the shared storage connection, the network will use an IP-based Ethernet communication network. However, do not use a storage network for node or client communication. Sharing the storage network in this manner might result in contention and latency issues for both users and the resource that the cluster provides.

Though not a best practice, you can use the private and public networks for both client and node communication. Preferably, you should dedicate an isolated network for private node communication. The reasoning for this is similar to that for using a separate Ethernet network for iSCSI—namely, to avoid resource congestion and contention issues. The public network is configured to allow client connections to the failover cluster. Although the public network can provide a backup for the private network, a better design practice is to define alternative networks for the primary private and public networks or, at least, to use bandwidth provisioning when teaming the network interfaces.

The networking features in clusters that are based on Windows Server 2016 include the following:

- The nodes transmit and receive heartbeats by using User Datagram Protocol (UDP) unicast instead of UDP broadcast (which is used in clusters based on earlier operating systems). The messages are sent on port 3343.
- You can include clustered servers on different IP subnets, which reduces the complexity of setting up stretch clusters.
- The Failover Cluster Virtual Adapter is a hidden device that is added to each node when you install the Failover Clustering feature. The adapter is assigned a media access control (MAC) address based on the MAC address that is associated with the first enumerated physical network adapter in the node.
- Failover clusters fully support IPv6 for both node-to-node and node-to-client communications.

- You can use Dynamic Host Configuration Protocol (DHCP) to assign IP addresses or static IP addresses to all nodes in the cluster. However, if some nodes have static IP addresses and you configure others to use DHCP, the **Validate a Configuration Wizard** will display an error. The cluster IP address resources are obtained based on the configuration of the network interface supporting that cluster network.

Failover cluster storage

Most failover clustering scenarios require shared storage to provide consistent data to a highly available service or application after a failover. The following are five shared-storage options for a failover cluster:

- Shared serial attached SCSI (SAS). SAS provides the lowest cost option. However, it is not very flexible for deployment because the two cluster nodes must be physically close together. In addition, the shared storage devices that support SAS 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 you use 1 gigabit per second (Gbps) or 10 Gbps with Ethernet as the physical medium for data transmission. This type of SAN is inexpensive to implement because no specialized networking hardware is required. In Windows Server 2016, you can implement iSCSI target software on any server and present local storage over iSCSI to clients.
- Fibre Channel. Fibre Channel SANs typically have better performance than iSCSI SANs but are much more expensive. Specialized knowledge and hardware are required to implement a Fibre Channel SAN.
- Shared virtual hard disk. In Windows Server 2012 R2 and later, you can use a shared virtual hard disk as storage for VM guest clustering. A shared virtual hard disk should be located on a CSV or Scale-Out File Server cluster, or you can add it to two or more VMs that are participating in a guest cluster by connecting to a SCSI or guest Fibre Channel interface.
- Scale-Out File Server. In Windows Server 2016, you can use shared SMB storage as the shared location for some failover cluster roles—specifically, for SQL Server and Hyper-V. This removes the need to have storage locally on the nodes hosting the SQL Server or Hyper-V roles, and all storage is handled over SMB 3.0 at the Scale-Out File Server.

- Failover clusters require shared storage to provide consistent data to a virtual server after a failover
- Shared storage options include:
 - SAS
 - iSCSI
 - Fibre Channel
 - Shared .vhdx
 - Scale-Out File Server
- You can also implement clustered storage spaces to achieve high availability at the storage level



Note: The Microsoft iSCSI Software Target is an integrated feature in Windows Server 2016. It can provide storage from a server over a TCP/IP network, including shared storage for applications that a failover cluster hosts. In addition, in Windows Server 2016, you can configure a highly available iSCSI Target Server as a clustered role by using either Failover Cluster Manager or Windows PowerShell.

In Windows Server 2016, in addition to using storage as a cluster component, you also can use failover clustering to provide high availability for the storage. This occurs when you implement clustered storage spaces. When you implement clustered storage spaces, you help to protect your environment from risks such as:

- Physical disk failures
- Data access failures
- Data corruption
- Volume unavailability
- Server node failures



Reference Links: For more information, refer to: "Deploy Clustered Storage Spaces" at: <http://aka.ms/b5cjdj>

Storage requirements and considerations

You also should be aware of the following shared storage requirements and considerations:

- To use the native disk support included in failover clustering, use basic disks and not dynamic disks.
- We recommend that you format the partitions with NTFS or Resilient File System (ReFS). For the disk witness, the partition must be NTFS or ReFS because a file allocation table (FAT) file system is not supported. Scale-Out File Servers do not support ReFS at this time.
- For the partition style of the disk, you can use either master boot record (MBR) or GUID partition table (GPT).
- Because improvements in failover clusters require that the storage respond correctly to specific SCSI commands, the storage must follow the SCSI Primary Commands - 3 (SPC-3) standard. In particular, the storage must support Persistent Reservations, as specified in the SPC-3 standard.
- The miniport driver used for the storage must work with the Microsoft Storport storage driver, when storage is located on a storage area network (SAN).
- You must isolate storage devices—that is, use one cluster per device. Servers from different clusters must be unable to access the same storage devices. In most cases, a LUN that one set of cluster servers use should be isolated from all other servers through LUN masking or zoning.
- Consider using Multipath I/O (MPIO) software. In a highly available storage fabric, you can deploy failover clusters with multiple host bus adapters by using MPIO software. This provides the highest level of redundancy and availability. For Windows Server 2016, you must base your multipath solution on MPIO. Your hardware vendor usually supplies an MPIO device-specific module (DSM) for your hardware, although Windows Server 2016 includes one or more DSMs as part of the operating system.
- If you use a shared virtual hard disk, you must have a CSV or a file server cluster to store the virtual hard disk.



Reference Links: For more information, refer to: "Failover Clustering Hardware Requirements and Storage Options" at: <http://aka.ms/kr8ahr>

What is quorum?

Quorum is the number of elements that must be online for a cluster to continue to run. In effect, each element can cast one vote to determine whether the cluster continues to run. Each cluster node is an element that has one vote. In case an even number of nodes exists, an additional element referred to as a *witness* is assigned to the cluster. The witness element can be a disk, a file share, or an Azure Cloud Witness if you run Windows Server 2016 nodes. Each voting element contains a copy of the cluster configuration, and the Cluster service works to keep all the copies synchronized at all times.

- In failover clusters, quorum defines the consensus that enough cluster members are available to provide services
- Quorum:
 - Is based on votes in Windows Server
 - Enables nodes, file shares, or a shared disk to have a vote, depending on the quorum mode
 - Enables the failover cluster to remain online when sufficient votes are available

The cluster stops providing failover protection if most of the nodes fail, or if a problem occurs with communication between the cluster nodes. Without a quorum mechanism, each set of nodes can continue to operate as a failover cluster. This results in a partition within the cluster.

Quorum prevents two or more nodes from concurrently operating a failover cluster resource. If you do not achieve a clear majority among the node members, the vote of the witness becomes crucial to maintain the validity of the cluster. Concurrent operation can occur when network problems prevent one set of nodes from communicating with another set of nodes. That is, a situation might occur in which more than one node tries to control access to a resource. If that resource is, for example, a database application, damage might result. Imagine the consequence if two or more instances of the same database were made available on the network, or if data was accessed and written to a target from more than one source at a time. If no damage to the application occurs, the data can easily become corrupted.

Because a specific 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 will stop running. That is, it will not provide failover protection if a node failure occurs. Nodes will still listen for the presence of other nodes on port 3343, in case another node appears again on the network, but the nodes will not function as a cluster until either a majority consensus occurs or you achieve a quorum.



Note: The full functioning of a cluster depends not only on a quorum but also on the capacity of each node to support the services and applications that fail over to that node. For example, a cluster that has five nodes can still have a quorum after two nodes fail, but each remaining cluster node will continue serving clients only if it has enough capacity—such as disk space, processing power, random access memory (RAM), or network bandwidth—to support the services and applications that failed over to it. An important part of the design process is planning each node's failover capacity. A failover node must run its own load and the load of additional resources that might fail over to it.

Achieving quorum

A cluster must complete several phases to achieve quorum. As a particular node comes up, it determines whether other cluster members exist, with which it can communicate. This process might simultaneously occur on multiple nodes. After establishing communication with other members, the members compare their membership views of the cluster until they agree on one view (based on time stamps and other information). A determination is made about whether this collection of members has quorum or has enough members to create sufficient votes so that a split scenario cannot exist. A split scenario means that

another set of nodes that are in this cluster run on a part of the network inaccessible to these nodes. Therefore, more than one node might be actively trying to provide access to the same clustered resource. If not enough votes exist to achieve quorum, the voters (the currently recognized members of the cluster) wait for more members to appear. After at least the minimum vote total is attained, the Cluster service begins to bring cluster resources and applications into service. With quorum attained, the cluster becomes fully functional.

Quorum modes in Windows Server 2016 failover clustering

The same quorum modes from Windows Server 2008 are present in Windows Server 2016. However, the process and recommendations for configuring quorum have changed. As before, a majority of votes determines whether a cluster achieves quorum. Nodes can vote and, where appropriate, so can either a disk in cluster storage (known as a *disk witness*), a file share (known as a *file share witness*), or an Azure Cloud Witness.

Before Windows Server 2012, only four quorum modes existed:

- Use dynamic quorum mode with:
 - Disk witness
 - File share witness
 - Azure Cloud Witness
- Use all other quorum modes only in specific use cases—the default and recommended best practice is to always use dynamic quorum

- **Node Majority.** Each node that is available and is in communication can vote. The cluster functions only with a majority, or more than half of the votes. This model is preferred when the cluster consists of an odd number of server nodes (no witness is needed to maintain or achieve quorum).
- **Node and Disk Majority.** Each node plus a designated disk in the cluster storage (the disk witness) can vote when they are available and in communication. The cluster functions only with a majority (more than half) of the votes. This model is based on an even number of server nodes being able to communicate with one another in the cluster in addition to the disk witness.
- **Node and File Share Majority.** Each node plus a designated file share created by the administrator, which is the file share witness, can vote when they are available and in communication. The cluster functions only with a majority of the votes. This model is based on an even number of server nodes being able to communicate with one another in the cluster in addition to the file share witness.
- **No Majority: Disk Only.** The cluster has quorum if one node is available and in communication with a specific disk in the cluster storage. Only the nodes that are also in communication with that disk can join the cluster.

Dynamic quorum

In Windows Server 2012, a new mode was introduced called *dynamic quorum*. Dynamic quorum dynamically adjusts the quorum votes based on the number of servers online. For example, assume you have a five-node cluster, you place two of the nodes in a paused state, and then one of the remaining nodes crashes. In any of the earlier configurations, your cluster would fail to achieve quorum and go offline. However, dynamic quorum adjusts the voting of the cluster when the first two servers are offline, making the number of votes for a quorum of the cluster two instead of three. The cluster with dynamic quorum stays online.

Building on the dynamic quorum mode, Windows Server 2012 R2 introduced the dynamic witness. A dynamic witness is a witness that dynamically has a vote depending on the number of nodes in the cluster. If an even number of nodes exists, the witness has a vote. If an odd number of nodes exist, the

witness does not have a vote. The recommended configuration for a cluster is to create a witness only when you have an even number of nodes. However, with the ability of a dynamic witness to remove voting to always have an odd number of votes in the cluster, you should always configure a witness for all clusters. This configuration is now the default mode for any configuration and is a best practice in most scenarios for Windows Server 2016 and Windows Server 2012 R2.

In Windows Server 2016, you can choose whether to use a disk witness, file share witness, or Azure Cloud Witness:

- **Disk witness.** A disk witness is still the primary witness in most scenarios, especially for local clustered scenarios. In this configuration, all the nodes have access to a shared disk. One of the greatest benefits of this configuration is that the cluster stores a copy of the cluster database on the disk witness.
- **File share witness.** A file share witness is ideal when shared storage is not available or when the cluster spans geographical locations. This option does not store a copy of the cluster database.
- **Azure Cloud Witness.** New to Windows Server 2016, the Azure Cloud Witness is the ideal option when you run Internet-connected stretch clusters. This removes the need to set up either a file share witness at a third datacenter location or a VM in the cloud. Instead, this option is built in to a failover cluster. This does not store a copy of the cluster database.

You should also consider the capacity of the nodes in your cluster and their ability to support the services and applications that might fail over to that node. For example, a cluster that has four nodes and a disk witness still has quorum after two nodes fail. However, if you have several applications or services deployed on the cluster, each remaining cluster node might not have the capacity to provide services.

What are CSVs?

In a classic failover cluster deployment, only a single node at a time controls a LUN in the shared storage. This means that the other nodes cannot see the shared storage until each node becomes an active node. CSVs are a technology introduced in Windows Server 2008 R2 that enables multiple nodes to concurrently share a single LUN. Each node obtains exclusive access to individual files on the LUN instead of to the whole LUN. In other words, CSVs provide a distributed file access solution so that multiple nodes in the cluster can access the same NTFS file system simultaneously.

• **The benefits of CSVs include:**

- Fewer required LUNs
- Better use of disk space
- Resources in a single logical location
- No special required hardware
- Increased resiliency

• **To implement a CSV:**

1. Create and format volumes on shared storage
2. Add the disks to failover cluster storage
3. Add the storage to the CSV

In Windows Server 2008 R2, CSVs are designed only for hosting VMs that run on a Hyper-V server in a failover cluster. This enabled administrators to have a single LUN that hosts multiple VMs in a failover cluster. Multiple cluster nodes have access to the LUN, but each VM runs on only one node at a time. If the node on which the VM runs fails, a CSV allows the VM to restart on a different node in the failover cluster. Additionally, this provides simplified disk management for hosting VMs compared to each VM requiring a separate LUN.

In Windows Server 2012, CSVs are enhanced. You can use CSVs with VMs for other roles and not only Hyper-V. For example, in a Scale-Out File Server scenario, you can configure the File Server role in a failover cluster. The Scale-Out File Server is designed to provide Scale-Out File Server shares that are continuously available for file-based server application storage. Scale-Out File Server shares provide the

ability to share the same folder from multiple nodes of the same cluster. In this context, CSVs in Windows Server 2012 introduces support for a read cache, which can significantly improve performance in certain scenarios. In addition, a CSV File System (CSVFS) can perform **chkdsk** without affecting applications with open handles on the file system.

Other important improvements in CSVs in Windows Server 2012 include:

- CSVFS benefits. In disk management, CSVs now display as **CSVFS**. However, this is not a new file system. The underlying technology is still the NTFS or the ReFS file system, and you still format CSVFS volumes with NTFS or ReFS. However, because volumes display as **CSVFS**, applications can discover that they run on CSVs, which helps to improve compatibility. In addition, because of a single file namespace, all files have the same name and path on any node in a cluster.
- Multiple subnet support for CSVs. CSVs are enhanced to integrate with SMB Multichannel to help achieve faster throughput for CSVs.
- Support for BitLocker. Windows Server 2012 supports BitLocker volume encryption for both traditional clustered disks and CSVs. Each node performs decryption by using the computer account for the cluster itself.
- Support for SMB 3.0 storage. CSVs in Windows Server 2012 provide support for SMB 3.0 storage for Hyper-V and for applications such as SQL Server.
- Integration with SMB Multichannel and SMB Direct. This integration allows CSV traffic to stream across multiple networks in the cluster and to fully utilize network adapters that support Remote Direct Memory Access (RDMA).
- Integration with the Storage Spaces feature in Windows Server 2012. This integration can provide virtualized storage on clusters of inexpensive disks.
- The ability to scan and repair volumes. CSVs in Windows Server 2012 support the ability to scan and repair volumes with virtually zero offline time.

Implementing CSVs

Before you can add storage to the CSV, the LUN must be available as shared storage to the cluster. When you create a failover cluster, all of the shared disks that are configured are added to the cluster, and you can add them to a CSV. If you add more LUNs to the shared storage, you must first create volumes on the LUN, add the storage to the cluster, and then add the storage to the CSV.

As a best practice, you should configure the CSV before you make any VMs highly available. However, you can convert from regular disk access to CSV after deployment. The following considerations apply:

- When you convert from regular disk access to CSV, the LUN's drive letter or mount point is removed. This means that you must re-create all VMs that are stored on the shared storage. If you must retain the same VM settings, consider exporting the VMs, switching to CSVs, and then importing the VMs in Hyper-V.
- You cannot add shared storage to a CSV if it is in use. If you have a running VM that is using a cluster disk, you must shut down the VM, and then add the disk to the CSV.



Additional Reading:

- For more information, refer to: "Server Message Block Overview" at: <http://aka.ms/rep0rf>
- For more information, refer to: "Storage Spaces Overview" at: <http://aka.ms/txzql4>

CSV improvements

Windows Server 2012 R2 provides improvements to CSVs. These improvements include optimized CSV placement policies, increased CSV resiliency, and CSV cache allocation. In addition, with Windows Server 2012 R2 and later, you now have both the ability to diagnose CSVs and enhanced interoperability between CSVs and other technologies.

The enhancements and new functionalities to CSVs Windows Server 2012 R2 include:

- Optimized CSV placement policies
- Increased CSV resiliency
- CSV cache allocation
- The ability to diagnose CSV
- CSV interoperability



Note: There have not been any major changes in CSVs for Windows Server 2016. However, this topic is still important because of the number of changes in Windows Server 2012 R2 from earlier versions of Windows Server.

Optimized CSV placement policies

In a failover cluster for Windows Server 2012, one node in the cluster is designated as the coordinator for a CSV, and no automatic rebalance occurs for this designation. The coordinator node for the CSV owns the physical disk resource that is associated with a LUN. All I/O operations that are specific to the file system are performed through the coordinator node.

In Windows Server 2012 R2 and later, CSV ownership is distributed evenly among the cluster nodes. This distribution is performed based on the number of CSVs that each node owns. The failover cluster automatically performs a rebalance in certain scenarios, such as when a node rejoins a cluster, when you add a new cluster, or when you restart a cluster node.

Increased CSV resiliency

CSVs since Windows Server 2012 use SMB as the transport mechanism for I/O forwarding between the nodes in a cluster. SMB uses the Server service on cluster nodes, and if this service becomes unavailable, it might result in decreases in performance or in the ability to access storage. Windows Server 2012 R2 implements multiple instances of the Server service, which improves the resilience and scalability of internode SMB traffic. The default instance of the Server service now accepts clients that access regular file shares, and a second Server service instance manages only internode CSV traffic. In addition, if the Server service becomes unhealthy on one cluster node, CSV ownership automatically transitions to another node to help ensure greater resiliency.

CSV cache allocation

To improve performance, the CSV cache enables the server to use RAM as a cache for write-through operations. In Windows Server 2012, the CSV cache is disabled by default, and when it is enabled, you can allocate up to 20 percent of the total RAM for the cache.

In Windows Server 2012 R2 and later, you can allocate up to 80 percent of memory for the CSV cache, which allows you to achieve performance gains for the clustered server role. This is especially useful for Scale-Out File Server clusters. In deployments where a Hyper-V cluster runs on a Scale-Out File Server cluster, we recommend that you enable and use the CSV cache but with a greater allocation for a Scale-Out File Server deployment to achieve the maximum performance of VMs stored on file servers.



Note: In Windows Server 2012 R2 and later, the name of the private property of the cluster physical disk resource has changed from **CsvEnableBlockCache** to **EnableBlockCache**.

Ability to diagnose CSV

In Windows Server 2012 R2 and later, you can view the state of the CSV on a per-node basis. For example, you can see whether I/O is direct or redirected or whether the CSV is unavailable. If a CSV is in I/O redirected mode, you can also view the reason for this. You can retrieve this information by using the Windows PowerShell **Get-ClusterSharedVolumeState** cmdlet with the **StateInfo**, **FileSystemRedirectedIOReason**, or **BlockRedirectedIOReason** parameter. This cmdlet provides you with a better view of how CSV works across cluster nodes.

CSV interoperability

CSVs in Windows Server 2012 R2 and later also support interoperability with the following technologies:

- ReFS
- Data deduplication
- Parity storage spaces
- Tiered storage spaces
- Storage Spaces write-back caching

This added support expands the number of scenarios in which you can use CSVs and allows you to use the efficiencies that these features introduce.

Question: What are some of the improvements to failover clustering in Windows Server 2016?

Question: What quorum configuration is a best practice for Windows Server 2016 failover clusters?

Lesson 2

Implementing a failover cluster

Failover clusters that you create in Windows Server 2016 have specific, recommended hardware and software configurations that allow Microsoft to support the cluster. The intent of failover clusters is to provide a higher level of service than standalone servers. Therefore, cluster hardware requirements are frequently stricter than the requirements for standalone servers.

This lesson describes how to prepare for cluster implementation. It also discusses the hardware, network, storage, infrastructure, and software requirements for Windows Server 2016 failover clusters. Finally, this lesson outlines the steps for using the **Validate a Configuration Wizard** to help ensure the correct cluster configuration.

Lesson Objectives

After completing this lesson, you will be able to:

- Prepare for implementing failover clustering.
- Describe the hardware requirements for a failover cluster implementation.
- Describe the network requirements for failover clustering.
- Describe the AD DS and infrastructure requirements for a failover cluster.
- Describe the software requirements for a failover cluster implementation.
- Validate and configure a failover cluster.

Preparing for implementing failover clustering

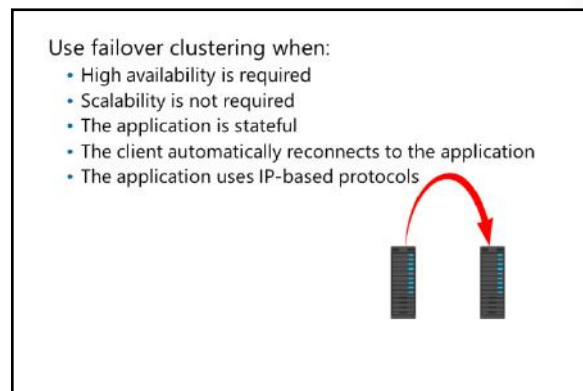
Before you implement failover clustering technology, you must identify the services and applications that you want to make highly available. You cannot apply failover clustering to all applications.

Failover clustering is best suited for stateful applications that are restricted to a single set of data. A database is one example of such an application. The data is stored in a single location and is used by only one database instance. You can also use failover clustering for Hyper-V VMs and for stateful applications that Hyper-V VMs implement.

The best results for failover clustering occur when the client can automatically reconnect to the application after a failover. If the client does not automatically reconnect, the user must restart the client application.

Consider the following guidelines when planning node capacity in a failover cluster:

- Distribute the highly available applications from a failed node. When all the nodes in a failover cluster are active, the highly available services or applications from a failed node should distribute among the remaining nodes to prevent a single node from overloading.



- Ensure that each node has sufficient capacity to service the highly available services or applications that you allocate to it when another node fails. This capacity should provide a sufficient buffer to avoid nodes that run at near capacity after a failure event. The failure to adequately plan resource utilization can result in a decrease in performance following a node failure.
- Use hardware with similar capacity for all the nodes in a cluster. This simplifies the planning process for failover because the failover load will distribute evenly among the surviving nodes.
- Use standby servers to simplify capacity planning. When a passive node is included in the cluster, all the highly available services or applications from a failed node can fail over to the passive node. This avoids the need for complex capacity planning. If you select this configuration, it is important that the standby server have sufficient capacity to run the load from more than one node failure.

You should also examine all the 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 multipathing software. These solutions reduce the probability that a single device failure will cause a cluster failure. Typically, server-class computer hardware has options for multiple power supplies for power redundancy and for creating Redundant Array of Independent Disks (RAID) sets for disk data redundancy.

Hardware requirements for failover cluster implementation

When you select hardware for cluster nodes, you must understand the hardware requirements. Failover clusters must satisfy the following hardware criteria to meet availability and support requirements:

- You should certify hardware for Windows Server.
- You should install the same or similar hardware on each failover cluster node. For example, if you choose a specific model of network adapter, you should install this adapter on each of the cluster nodes to avoid compatibility and capacity issues.
- If you use SAS or Fibre Channel storage connections, the mass-storage device controllers that are dedicated to the cluster storage should be identical in all clustered servers. They also should use the same firmware version.
- If you are using iSCSI storage connections, each clustered server should have one or more network adapters or host bus adapters dedicated to the cluster storage. You should not use the network that you use for iSCSI storage connections for nonstorage 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 10-gigabit Ethernet or more.
- After you configure the servers with the hardware, all the tests provided in the **Validate a Configuration Wizard** must pass before the cluster is considered a configuration that is supported by Microsoft.
- Each node must run the same processor architecture. This means that each node must have the same processor family, which might be the Intel or the Advanced Micro Devices (AMD) family.

The hardware requirements for a failover implementation include:

- The server hardware must be certified for Windows Server
- The server nodes should all have the same configuration and contain the same or similar components
- All the tests in the Validate a Configuration Wizard pass


Network requirements for failover cluster implementation

In addition to the hardware, the components should meet certain requirements and pass the tests in the **Validate a Configuration Wizard**. The suggested requirements include the following:

- The network adapters in each node should be identical and have the same IP protocol version, speed, duplex, and flow control capabilities.
- The networks and network equipment to which you connect the nodes should be redundant so that even a single failure allows the nodes to continue communicating with one another. You can team network devices to provide single network redundancy. We recommend multiple networks to provide multiple paths among nodes for internode communication; otherwise, a warning will appear during the validation process.
- The network adapters in a cluster network should have the same IP address assignment method, which means that they all use either static IP addresses or DHCP.
- Network settings and IP addresses. When you use identical network adapters for a network, also use identical communication settings on those adapters, such as those for speed, duplex mode, flow control, and media type. In addition, compare the settings between the network adapter and the switch to which it connects, and ensure that no settings are in conflict. Otherwise, network congestion or frame loss might occur, which can adversely affect how the cluster nodes communicate among themselves, with clients, or with storage systems.
- Unique subnets. If you have private networks that are not routed to the rest of the network infrastructure, ensure that each of those private networks uses a unique subnet. This is necessary even if you give each network adapter a unique IP address. For example, if you have a cluster node in a central office that uses one physical network and another node in a branch office that uses a separate physical network, do not specify 10.0.0.0/24 for both networks—even if you give each adapter a unique IP address. This helps to avoid routing loops and other network communication problems if, for example, the segments are accidentally configured into the same collision domain because of incorrect virtual local area network (VLAN) assignments.

The network requirements for a failover implementation include:

- The server should be connected to multiple networks for communication redundancy or to a single network with redundant hardware to remove single points of failure
- The network adapters should be identical and have the same IP protocol versions, speed, duplex, and flow control capabilities

 **Note:** If you connect cluster nodes with a single network, the network will pass the redundancy requirement in the **Validate a Configuration Wizard**. However, the report from the wizard will include a warning that the network should not have single points of failure.

AD DS and infrastructure requirements for failover clusters

Failover clusters depend on infrastructure services. In Windows Server 2012, each server node must be in the same Active Directory domain, and if you use DNS, the nodes should use the same DNS servers for name resolution. However, in Windows Server 2016, multiple-domain clusters and workgroup clusters are supported.



Note: Though you can deploy multiple-domain clusters and workgroup clusters, we recommend that you do not use this configuration for Hyper-V or file server clusters.

- The infrastructure requirements for a failover cluster implementation include the following:
 - The nodes in the cluster typically use DNS for name resolution
 - All servers in the cluster should be in the same Active Directory domain for Windows Server 2012, however this is not required for Windows Server 2016
 - The user account that creates the cluster must have administrator rights and permissions on all servers, and the Create Computer Objects permission in the domain
- Failover cluster infrastructure recommendations include:
 - Do not install the AD DS role on any of the cluster nodes

We recommend that you install the same Windows Server 2016 features and roles on each node. Inconsistent configurations on cluster nodes might cause instability and performance issues. In addition, you should not install the AD DS role on any of the cluster nodes because AD DS has its own fault-tolerance mechanism.

You must have the following network infrastructure elements for a failover cluster:

- DNS. The servers in the cluster typically use DNS for name resolution. DNS dynamic update is a supported configuration.
- A domain role. In Windows Server 2012 and Windows Server 2012 R2, it is required that all servers in the cluster be in the same Active Directory domain. We recommend that all clustered servers have the same domain role (either member server or domain controller). The recommended role is member server because AD DS inherently includes its own failover protection mechanism. In Windows Server 2016 failover clusters, the cluster nodes do not need to be members of the same domain.
- An account for administering the cluster. When you first create a cluster or add servers to it, you must sign in to the domain with an account that has administrator rights and permissions on all servers in that cluster. The account does not have to be a Domain Admins account but can be a Domain Users account that is in the Administrators group on each clustered server. In addition, if the account is not a Domain Admins account, the account (or the group in which the account is a member) must be given the Create Computer Objects permission in the domain. The permission to create computer objects is not required when you create detached clusters in AD DS.

In Windows Server 2016, no Cluster service account exists. Instead, the Cluster service automatically runs in a special context that provides the specific permissions and credentials that are necessary for the service (similarly to the local system context but with reduced credentials). When a failover cluster is created and a corresponding computer object is created in AD DS, that object is configured to help prevent accidental deletion. In addition, the cluster Network Name resource has additional health check logic, which periodically checks the health and properties of the computer object that represents the Network Name resource.


Software requirements for a failover cluster implementation


As a best practice, we recommend that each cluster node run the same edition of Windows Server 2016. The edition can be either Windows Server 2016 Standard or Windows Server 2016 Datacenter. The nodes should also have the same software updates and service packs. Depending on the role that is clustered, a Server Core or Nano Server installation of Windows Server 2016 might also meet the software requirements.

The same versions of service packs or any operating system updates should exist on all the nodes that are part of a cluster.

The software best practices for a failover cluster implementation include:

- All nodes should have the same edition of Windows Server 2016, which can be any of the following:
 - Windows Server 2016 Standard, Desktop Experience, Server Core, or Nano Server installation
 - Windows Server 2016 Datacenter, Desktop Experience, Server Core, or Nano Server installation
- All nodes should have the same service pack and updates

 **Note:** With Windows Server 2016 and Cluster Operating System Rolling Upgrade, you might have different operating systems in a cluster. As stated, it is a best practice to have a cluster with the same operating system, edition, and updates. However, the cluster can run without this configuration, especially during an upgrade process.

 **Note:** Windows Server 2016 provides Cluster-Aware Updating (CAU) technology that can help you to maintain updates on cluster nodes. The “Maintaining a failover cluster” lesson discusses this feature in more detail.

Demonstration: Validating and configuring a failover cluster

The **Validate a Configuration Wizard** runs tests that determine if the hardware and hardware settings are compatible with failover clustering. By using this wizard, you can run the complete set of configuration tests or a subset of the tests. We recommend that you run the tests on servers and storage devices before you configure the failover cluster and again after any major changes are made to the cluster. You can access the test results in the **%windir%\cluster\Reports** directory.

In this demonstration, you will see how to validate and configure a failover cluster.

Demonstration Steps

1. Start the Failover Cluster Manager on the **LON-SVR2** VM.
2. Start the **Validate Configuration Wizard**.
3. Add **LON-SVR2** and **LON-SVR3** as cluster nodes.
4. Review the report.
5. Create a new cluster.
6. Add **LON-SVR2** and **LON-SVR3** as cluster nodes.

7. Name the cluster **Cluster1**.
8. Use **172.16.0.125** as the **IP address**.

Question: Does Windows Server 2016 require all nodes to be in the same domain?

Question: Can a node that runs Windows Server 2016 and one that runs Windows Server 2012 R2 both run in the same cluster?

Lesson 3

Configuring highly available applications and services on a failover cluster

After you configure a clustering infrastructure, you should configure specific roles or services to be highly available. Not all roles can be clustered. Therefore, you should first identify the resource that you want to put in a cluster and then verify whether it is supported. In this lesson, you will learn about configuring roles and applications in clusters and about configuring cluster settings.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe and identify cluster resources and services.
- Describe the process for clustering server roles.
- Cluster a file server role.
- Describe failover cluster management tasks.
- Manage cluster nodes.
- Configure application failover settings.
- Sequence an activity.

Identifying cluster resources and services

A clustered service that contains an IP address resource and a network name resource (and other resources) publishes to a client on the network under a unique server name. Because this group of resources displays as a single logical server to clients, it is called a *cluster instance*.

Users access applications or services on a cluster instance in the same manner they would if the applications or services were on a nonclustered server. Usually, applications or users do not know that they are connecting to a cluster or to which node they are connecting.

A *cluster resource* is a physical or logical entity (such as a file share, disk, or IP address) that the failover cluster manages. A resource might provide a service to clients or be an important part of the cluster. A resource is the most basic and smallest configurable unit. At any time, a resource can run only on a single node in a cluster, and it is online on a node when it provides its service to that specific node.

- **Clustered services:**
 - Are services or applications that are made highly available by installing them on a failover cluster
 - Are active on one node but can be moved to another node
- **Resources:**
 - Are the components that make up a clustered service
 - Are moved to another node when one node fails
 - Can run on only one node at a time
 - Include components such as shared disks, names, and IP addresses

Server cluster resources

A *server cluster resource* is any physical or logical component that has the following characteristics:

- It can be brought online and taken offline.
- It can be managed in a server cluster.
- It can be hosted (owned) by only one node at a time.

To manage resources, the Cluster service communicates to a resource dynamic-link library (DLL) through a resource monitor. When the Cluster service makes a request of a resource, the resource monitor calls the appropriate entry point function in the resource DLL to check and control the resource state.

Dependent resources

A dependent resource is one that requires another resource to operate. For example, a network name must be associated with an IP address. Because of this requirement, a network name resource depends on an IP address resource.

Dependent resources are taken offline before the resources on which they depend are taken offline. Similarly, they are brought online after the resources on which they depend are brought online. A resource can specify one or more resources on which it depends. Resource dependencies also determine bindings. For example, clients will be bound to the particular IP address on which a network name resource depends.

When you create resource dependencies, consider the fact that although some dependencies are strictly required, others are not required but are recommended. For example, a file share that is not a Distributed File System (DFS) root has no required dependencies. However, if the disk resource that contains the file share fails, the file share will be inaccessible to users. Therefore, it is logical to make the file share dependent on the disk resource.

A resource can also specify a list of nodes on which it can run. Possible nodes and dependencies are important considerations when administrators organize resources into groups.

Clustering server roles process

Failover clustering supports the clustering of several Windows Server roles, such as File Services, DHCP, and Hyper-V. To implement clustering for a server role or for an external application such as SQL Server or Exchange Server, perform the following procedure:

1. Install the Failover Clustering feature. Use Server Manager or Windows PowerShell to install the Failover Clustering feature on all the computers that will be cluster members.
2. Verify the configuration, and create a cluster with the appropriate nodes. Use the Failover Cluster Management snap-in to first validate a configuration and to then create a cluster with selected nodes.
3. Install the role on all cluster nodes. Use Server Manager to install the server role that you want to use in the cluster.
4. Create a clustered role by using the Failover Clustering Management snap-in.

1. Install the failover clustering feature
2. Verify the configuration, and create a cluster
3. Install the role on all cluster nodes, by using Server Manager
4. Create a clustered application by using the Failover Clustering Management snap-in
5. Configure the application
6. Test the failover

5. Configure the cluster role. Configure the options on the application that is used in the cluster.
6. Test failover. Use the Failover Cluster Management snap-in to test failover by intentionally moving the service from one node to another.



Note: If the cluster is a mixed domain or workgroup cluster, you must use only Windows PowerShell to configure the cluster. Beginning with Windows Server 2016 Technical Preview 5, the Failover Clustering Management snap-in is not supported.

After you create the cluster, you can use the **Failover Cluster Management** console to monitor its status and manage the available options.

Demonstration: Clustering a file server role

In this demonstration, you will see how to cluster a file server role.

Demonstration Steps

1. Open the **Failover Cluster Manager**, and then verify that three cluster disks are available.
2. Start the **Configure Role Wizard**, and then configure **File Server** as a clustered role.
3. For the client access point, use the name **AdatumFS**.

Failover cluster management tasks

You can perform several failover cluster management tasks. These tasks range from adding and removing cluster nodes to modifying the quorum settings. Some of the most frequently used configuration tasks include:

- Managing cluster nodes. For each node in a cluster, you can temporarily stop the Cluster service, pause it and initiate Remote Desktop to the node, or evict the node from the cluster. You can also choose to drain nodes in the cluster—for example, if you want to perform maintenance or install updates. This functionality is part of the infrastructure that enables CAU for patching nodes in a cluster.
- Managing cluster networks. You can add or remove cluster networks, and you also can configure networks that are dedicated only for intercluster communication.
- Managing permissions. By managing permissions, you delegate rights to administer clusters.
- Configuring cluster quorum settings. By configuring quorum settings, you determine how to achieve a quorum in addition to who can have a vote in a cluster.
- Migrating services and applications to a cluster. You can implement existing services in the cluster and make them highly available.

Common management tasks include:

- Managing cluster nodes
- Managing cluster networks
- Managing permissions
- Configuring cluster quorum settings
- Migrating services and applications to a cluster
- Configuring new services and applications
- Removing clusters
- Upgrading cluster nodes to new operating system versions


- Configuring new services and applications to work in a cluster. You can implement new services in the cluster.
- Removing a cluster.

You can perform most of these administrative tasks by using the **Failover Cluster Management** console.

Upgrading failover clusters for Windows Server 2016

For Windows Server 2016, the process of upgrading a failover cluster has changed. With cluster operating system upgrades, the cluster operating system can be upgraded first before you upgrade the functional level of the cluster. For example, if you have a two-node cluster with Windows Server 2012 R2, you can upgrade it to Windows Server 2016 by draining the roles from one node, taking it offline, and then removing it from the cluster. You can then upgrade that node to Windows Server 2016 and add the node back into the cluster. The cluster will continue to run at the functional level of Windows Server 2012 R2. You can then drain the roles back to the Windows Server 2016 node. Then remove the Windows Server 2012 R2 node from the cluster, upgrade it, and add it back to the cluster. Finally, with both nodes running Windows Server 2016, you can upgrade the functional level by running the following Windows PowerShell command.

```
Update-ClusterFunctionalLevel
```

 **Note:** You can achieve this same upgrade by adding a new Windows Server 2016 node to the cluster, upgrading or removing all the Windows Server 2012 R2 nodes from the cluster, and then upgrading the new node. Many ways to achieve this same cluster upgrade exist, and these are just some examples.

Managing cluster nodes

After you create a cluster and put it into production, you might occasionally have to perform management tasks on the cluster nodes. Cluster node management tasks typically belong to one of the following three categories:

- Adding a node. You can add a node to an established failover cluster by selecting **Add Node** in the **Actions** pane of the **Failover Cluster Management** console. The **Add Node Wizard** prompts you for information about the additional node.
- Pausing a node. You can pause a node to prevent resources from failing over or being moved to the node. You typically pause a node when it is undergoing maintenance or troubleshooting.
- Evicting a node. You can evict a node from the cluster. After you evict the node, you must add it back to the cluster. You evict a node when it is damaged or no longer needed in the cluster. If you evict a damaged node, you can repair or rebuild it and then add it back to the cluster by using the **Add Node Wizard**.

- To manage cluster nodes, you can:
 - Add nodes after you create a cluster
 - Pause nodes, which prevents resources from running on that node
 - Evict nodes from a cluster, which removes the node from the cluster configuration
 - All of these actions are available in the Failover Cluster Management console, Actions pane

You can manage a cluster in the **Actions** pane of the **Failover Cluster Management** console.

Configuring application failover settings

You can adjust the failover settings, including preferred owners and failback settings, to control how the cluster responds when roles or services fail. You can configure these settings on the property sheet for the clustered service or application (on either the **General** or the **Failover** tab).

You can individually set preferred owners in the properties of each role. You can select multiple preferred owners and place them in any order that is preferred. Selecting preferred owners provides more control over what node a particular role fails over to and actively runs on.

- Considerations for using preferred owners:
 - You set preferred owners are set on the clustered role
 - You can set multiple preferred owners can be set in an ordered list
 - Setting preferred owners gives control over:
 - The order in which a role selects a node to run
 - The roles that can be run on the same nodes
- Options to modify failover and failback settings:
 - Setting the number of times the Cluster service restarts a clustered role in a set period
 - Setting or preventing failback of the clustered role to the preferred node when it becomes available

Settings also exist that you can change for each role for failover and failback. Failover settings can control how many times the cluster is allowed to attempt restarting a role in a particular amount of time. The default in Windows Server 2012 R2 is three failures in six hours to restart. In Windows Server 2016, this is changed to allowing only one failure every six hours. You can set the failback settings to **Prevent Failback** or **Allow Failback**. When allowing failback, you can set the role to immediately use failback or to use failback during a certain number of hours. The default settings are **Prevent Failback**. The following table provides examples that show how these settings work.

Setting	Result
Example 1: General tab, Preferred owner: Node1 Failover tab, Failback setting: Allow failback (Immediately)	If the service or application fails over from Node1 to Node2, when Node1 is again available, the service or application fails 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 is restarted or failed over every time. If the application or service fails a third time in a six-hour period, it remains in a 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 among nodes indefinitely.

Question: In the **Failover Cluster Manager** console, what are some of the Microsoft roles you can configure?

Sequencing Activity

The following are the steps for clustering server roles. Arrange them in the correct order by numbering each step.

	Steps
	Install the Failover Clustering feature. Use Server Manager or Windows PowerShell to install the Failover Clustering feature on all computers that will be cluster members.
	Verify the configuration, and create a cluster with the appropriate nodes. Use the Failover Cluster Management snap-in to first validate a configuration and to then create a cluster with selected nodes.
	Install the role on all the cluster nodes. Use Server Manager to install the server role that you want to use in the cluster.
	Create a clustered application by using the Failover Clustering Management snap-in.
	Configure the application. Configure the options on the application that is used in the cluster.
	Test the failover. Use the Failover Cluster Management snap-in to test the failover by intentionally moving the service from one node to another.

Lesson 4

Maintaining a failover cluster

After you have your cluster infrastructure running, you should establish monitoring procedures to prevent possible failures. In addition, you should have backup and restore procedures for cluster configuration. In Windows Server 2016, CAU allows you to update cluster nodes without downtime. In this lesson, you will learn about monitoring, backing up and restoring, and updating cluster nodes.

Lesson Objectives

After completing this lesson, you will be able to:

- Monitor failover clusters.
- Backup and restore failover cluster configurations.
- Troubleshoot failover clusters.
- Describe CAU.
- Configure CAU.

Monitoring failover clusters

Many tools are available to help you monitor failover clusters. You can use standard Windows Server operating system tools, such as the Event Viewer and the Performance and Reliability Monitor snap-in, to review cluster event logs and performance metrics. You also can use the **Tracerpt.exe** tool to export data for analysis. In addition, you can use the Multipurpose Internet Mail Extension Hypertext Markup Language (MHTML)-formatted cluster configuration reports and the **Validate a Configuration Wizard** to troubleshoot problems with the cluster configuration and hardware changes.

Tools you can use to monitor clusters include:

- The Event Viewer
- Tracerpt.exe
- MHTML-formatted cluster configuration reports
- The Performance and Reliability Monitor snap-in

Event Viewer

If problems arise in a cluster, you can use the Event Viewer to view events with a Critical, Error, or Warning severity level. In addition, you can view informational-level events in the Failover Clustering Operations log, which you can access in the Event Viewer in the **Applications and Services Logs\Microsoft \Windows** folder. Informational-level events are usually common cluster operations, such as cluster nodes leaving and joining the cluster or resources going offline or coming online.

In earlier versions of Windows Server, event logs were replicated to each node in the cluster. This simplified cluster troubleshooting because you could review all the event logs on a single cluster node. Windows Server 2012 and later does not replicate the event logs among nodes. However, the Failover Cluster Management snap-in has a **Cluster Events** option that allows you to view and filter events across all cluster nodes. This feature is helpful in correlating events across cluster nodes.

The Failover Cluster Management snap-in also provides a **Recent Cluster Events** option that queries all the Error and Warning events in the last 24 hours from all the cluster nodes.

You can access additional logs, such as the Debug and Analytic logs, in the Event Viewer. To display these logs, modify the view on the menu by selecting the **Show Analytic** and **Debug Logs** options.

Event tracing for Windows

Event tracing for Windows is a kernel component that is available soon after startup and late into shutdown. It is designed to allow the fast tracing and delivery of events to both trace files and consumers. Because it is designed to be fast, it allows only the basic, in-process filtering of events based on event attributes.

The event trace log contains a comprehensive accounting of the failover cluster actions. To view the data, use **Tracerpt.exe** to access the information in the event trace log. **Tracerpt.exe** parses the event trace logs only on the node on which it runs. All of the individual logs are collected in a central location. To transform the XML file into a text file or into an HTML file that you can open in Microsoft Edge or Windows Internet Explorer, you can parse the XML-based file by using the Microsoft Extensible Stylesheet Language (XSL) parsing command prompt utility **Msxsl.exe** and an XSL style sheet.

Performance and Reliability Monitor snap-in

You can also use the Performance and Reliability Monitor snap-in to help monitor failover clusters. The Performance and Reliability Monitor snap-in allows you to:

- See how application performance is trending on each node. To determine how an application is performing, you can view specific information on the system resources that are used on each node and see how that information is trending.
- See how the application failures and stability on each node are trending. You can pinpoint when application failures occur and match the application failures with other events on the node.
- Modify trace log settings. You can start, stop, and adjust trace logs, including their size and location.

Backing up and restoring a failover cluster configuration

Configuring clusters can be a time-consuming and detail-oriented process. To avoid having to repeat the process, you should always back up cluster configurations. You can perform a backup and restore of cluster configurations with Windows Server Backup or with another non-Microsoft backup tool.

When you back up your cluster configuration, be aware of the following:

- You must test your backup and recovery process before you put a cluster into production.
- You must first add the Windows Server Backup feature if you decide to use it. You can do this by using Server Manager.

- When backing up failover clusters, keep in mind that:
 - Windows Server Backup is a Windows Server 2016 feature
 - Non-Microsoft tools are available to perform backup and restore operations
 - You must perform system-state backups
- A nonauthoritative restore operation completely restores a single node in the cluster
- An authoritative restore operation restores the entire cluster configuration to a certain point

Windows Server Backup is the built-in backup and recovery software for Windows Server 2016. To complete a successful backup, consider the following:

- For a backup to succeed in a failover cluster, the cluster must be running and have quorum. In other words, enough nodes must be running and communicating (perhaps with a witness disk or witness file share, depending on the quorum configuration) that the cluster has achieved quorum.

- You must back up all clustered applications. If you cluster a SQL Server database, you must have a backup plan for the databases and configuration outside the cluster configuration.
- If you must back up the application data, the disks on which you store the data must be made available to the backup software. You can achieve this by running the backup software from the cluster node that owns the disk resource or by running a backup against the clustered resource over the network.
- The Cluster service keeps track of which cluster configuration is the most recent, and it replicates that configuration to all cluster nodes. If the cluster has a witness disk, the Cluster service also replicates the configuration to the witness disk.

Restoring a cluster

Two types of restore operations exist for clusters:

- **Nonauthoritative restore operation.** Use a nonauthoritative restore operation when a single node in the cluster is damaged or rebuilt, and the rest of the cluster is operating correctly. Perform a nonauthoritative restore operation by restoring the system recovery (system state) information to the damaged node. When you restart that node, it joins the cluster and automatically receives the latest cluster configuration.
- **Authoritative restore operation.** Use an authoritative restore operation when you must roll back the cluster configuration. For example, you use an authoritative restore operation if an administrator accidentally removed clustered resources or modified other cluster settings, and you need to revert the cluster to a previous point in time. To perform the authoritative restore operation, stop the cluster resource on each node and then perform a system recovery on a single node by using the Windows Server Backup command-line interface. After the restored node restarts the Cluster service, the remaining cluster nodes can also start the Cluster service.

Troubleshooting failover clusters

Although cluster validation in Windows Server 2016 failover clustering prevents misconfigurations and nonworking clusters, in some cases, you must still perform cluster troubleshooting. To troubleshoot a failover cluster, use the following guidelines:

- Use the **Validate a Configuration Wizard** to identify configuration issues that might cause cluster problems.
- Review the cluster events and trace logs to identify application or hardware issues that might cause a cluster to become unstable.
- Review the hardware events and logs to help pinpoint specific hardware components that might cause a cluster to become unstable.
- Review the SAN components, switches, adapters, and storage controllers to help identify any potential problems.

Failover cluster troubleshooting techniques include:

- Using the Validate a Configuration Wizard
- Reviewing the events in logs (cluster, hardware, storage)
- Defining a process for troubleshooting failover clusters
- Reviewing the storage configuration
- Checking for group and resource failures

When troubleshooting failover clusters, you must:

- Identify the perceived problem by collecting and documenting the symptoms of the problem.
- Identify the scope of the problem so that you can understand what the problem affects and what impact the effect has on the application and the clients.
- Collect information so that you can accurately understand and pinpoint the possible problem. After you identify a list of possible issues, you can prioritize them by probability or by the impact of a repair. If you cannot identify the problem, you should attempt to re-create the problem.
- Create a schedule for repairing the problem. For example, if the problem affects only a small subset of users, you can delay the repair to an off-peak time so that you can schedule downtime.
- Complete and test each repair one at a time so that you can identify the fix.

To troubleshoot SAN issues, start by checking the physical connections and by reviewing each of the hardware component logs. Next, run the **Validate a Configuration Wizard** to verify that the current cluster configuration is still supported. When you run the **Validate a Configuration Wizard**, ensure that the storage tests that you select can run on an online failover cluster. Several of the storage tests cause a loss of service on the clustered disk when the tests run.

Troubleshooting group and resource failures

To troubleshoot group and resource failures:

- Use the Dependency Viewer in the Failover Cluster Management snap-in to identify dependent resources.
- Review the Event Viewer and trace logs for errors from the dependent resources.
- Determine whether the problem happens on only a specific node or nodes by trying to re-create the problem on different nodes.

What is CAU?

Applying operating system updates to nodes in a cluster requires special attention. With earlier versions of Windows Server, if you wanted to provide zero downtime for a clustered role, you had to manually update the cluster nodes one after another, and you had to manually move resources from the node being updated to another node. This procedure could be very time consuming. In Windows Server 2012, Microsoft implemented CAU, a feature for automatically updating cluster nodes.

CAU is a feature that allows administrators to automatically update cluster nodes with little or no loss in availability during the update process. During an update procedure, CAU transparently takes each cluster node offline, installs the updates and any dependent updates, performs a restart if necessary, brings the node back online, and then moves to update the next node in a cluster.

CAU is an automated feature in Windows Server 2016 that:

- Updates nodes in a cluster
- Has these benefits:
 - Updating is automatic
 - Updating can be scheduled
 - Updating causes minimal or no downtime

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 2016, such as Hyper-V with live migration or file server with SMB transparent failover, CAU can orchestrate cluster updates with no effect on the service availability.

How CAU works

CAU is based on orchestrating a process of cluster node updating. CAU can orchestrate the complete cluster updating in one of the following two modes:

- Remote-updating mode. In this mode, a computer that runs Windows Server 2012 R2 or Windows 8.1 or later is called and configured as a *CAU orchestrator*. To configure a computer as a CAU orchestrator, you must install the failover clustering administrative tools. The CAU orchestrator is not a member of the cluster that is updated during the procedure. From the CAU orchestrator, the administrator triggers on-demand updating by using a default or custom Updating Run profile. Remote-updating mode is useful for monitoring real-time progress during the updating run and for clusters that run on Server Core installations of Windows Server 2016.
- 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 CAU orchestrator. 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 sequentially on each cluster node.

CAU can work in two modes:

- Remote-updating mode:
 - You configure a separate computer as an orchestrator
 - You must install the failover clustering administrative tools
 - The CAU orchestrator must not be a cluster member
- Self-updating mode:
 - You configure the CAU clustered role as a workload
 - No dedicated orchestrator exists
 - The cluster updates itself

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. 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 Windows PowerShell **Get-CauRun** cmdlet.

To use CAU, you must install the Failover Clustering feature in Windows Server 2016 and create a failover cluster. The components that support CAU functionality then automatically install on each cluster node.

You must also install the CAU tools, which are included in the Failover Clustering Tools that are part of the Remote Server Administration Tools (RSAT). The CAU tools consist of the CAU UI and the CAU Windows PowerShell cmdlets. When you install the Failover Clustering feature, the Failover Clustering Tools install by default on each cluster node. You can also install these tools on a local or remote computer that runs Windows Server 2016 or Windows 10 and that has network connectivity to the failover cluster.

Demonstration: Configuring CAU

In this demonstration, you will see how to configure CAU.

Demonstration Steps

1. Ensure that the cluster is configured and running on **LON-SVR2** and **LON-SVR3**.
2. Add the **Failover Clustering** feature to **LON-DC1**.
3. Run **Cluster-Aware Updating** on **LON-DC1**, and then configure it to connect to **CLUSTER1**.
4. Preview the updates that are available for the nodes **LON-SVR2** and **LON-SVR3**.
5. Review the available options for the **Updating Run** profile.
6. Apply the available updates to **CLUSTER1** from **LON-DC1**.
7. After the updates apply, configure the **Cluster self-updating options** on **LON-SVR2**.

Question: What are some of the troubleshooting techniques for failover clusters?

Question: You have an eight-node cluster running Hyper-V in Windows Server 2016. How would you run Windows updates on each node on a schedule without downtime?

Lesson 5

Implementing a stretch cluster

In some scenarios, you must deploy cluster nodes on different sites. Usually, you do this when you build disaster recovery solutions. In this lesson, you will learn about deploying stretch clusters.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe a stretch cluster.
- Describe synchronous and asynchronous replication.
- Describe site-aware failover clusters.
- Choose a quorum witness.
- Describe the considerations for deploying stretch clusters.
- Describe the considerations for stretch cluster failover and failback.

What is a stretch cluster?


A *stretch cluster* provides highly available services in more than one location. Although stretch clusters can solve several specific problems, they also present specific challenges.

In a stretch cluster, each site usually has a separate storage system with replication among the sites. Stretch cluster storage replication allows each site to be independent and provides fast access to the local disk. With separate storage systems, you cannot share a disk among sites.

Compared to a remote server, a stretch cluster has three main advantages in a failover site:

- When a site fails, a stretch cluster can automatically fail over the clustered service or application to another site.
- Because the cluster configuration automatically replicates to each cluster node in a stretch cluster, less administrative overhead exists than with a standby server, which requires you to manually replicate changes.
- The automated processes in a stretch cluster reduce the possibility of human error, which is inherent in manual processes.

Because of the increased cost and the complexity of a stretch cluster, it might not be an ideal solution for every application or business. When you are considering whether to deploy a stretch cluster, you should evaluate the importance of the applications to the business, the types of applications, and any alternative solutions. Some applications can easily provide stretch cluster redundancy with log shipping or other processes and can still achieve sufficient availability with only a modest increase in cost and complexity.



The complexity of a stretch cluster requires more-detailed architectural and hardware planning than is required for a single-site cluster. It also requires you to develop business processes to routinely test the cluster functionality.

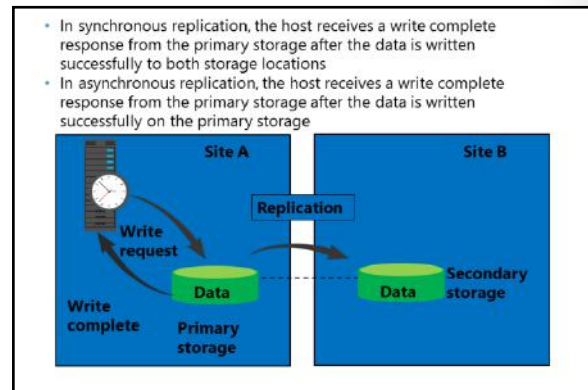


Note: Many of the new features for Windows Server 2016 were created to address the challenges of stretch clusters. Azure Cloud Witness, site-aware failover clusters, and Storage Replica were all implemented to address the challenges of stretch clusters in many of the earlier Windows Server versions.

Synchronous and asynchronous replication

Until recently, a geographically dispersed failover cluster running Windows Server could not use shared storage among the physical locations without a vendor-specific storage replication solution. In Windows Server 2016, this specific hardware is no longer required. Storage Replica uses either synchronous or asynchronous replication, separately from the vendor storage that might exist at the location:

- When you use synchronous replication, after the data writes successfully on both storage systems, the host receives a write complete response from the primary storage. If the data does not write successfully to both storage systems, the application must attempt to write to the disk again. With synchronous replication, the data on both storage systems is identical.
- When you use asynchronous replication, after the data writes successfully on the primary storage, the node receives a write complete response from the 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 stretch cluster. For example, DFS Replication provides file-level asynchronous replication. However, it does not support stretch clustering replication. This is because DFS Replication is designed to replicate smaller documents that are not continuously kept open. As a result, it was not designed for high-speed, open-file replication.

When to use synchronous or asynchronous replication

Use synchronous replication when it is imperative that you avoid data loss. Synchronous replication solutions require low-latency disk write operations, because the application waits for both storage solutions to acknowledge a data write operation. The requirement for low-latency disk write operations also limits the distance between the storage systems, because an increased distance might cause a higher latency. If the latency is high, the performance and even the stability of the application might be affected.

Asynchronous replication overcomes the latency and distance limitations by acknowledging only the local disk write operations and by reproducing a disk write operation on the remote storage system in a separate transaction. However, because asynchronous replication writes to the remote storage system after it writes to the local storage system, the possibility of data loss during a failure increases.

Site-aware failover clusters

Site-aware failover clusters comprise one of the major changes in Windows Server 2016. You can now configure the sites with values that allow you to differentiate among site locations. For example, if you have a four-node cluster with two nodes at one site and two nodes at another site, you can run the following commands to make the nodes site aware.

```
(Get-ClusterNode Node1).Site=1
(Get-ClusterNode Node2).Site=1
(Get-ClusterNode Node3).Site=2
(Get-ClusterNode Node4).Site=2
```

Site-aware failover cluster services:

- Failover affinity
- Cross-site heartbeating
- Preferred site configuration

After these have been defined, multiple new and improved failover cluster services can then be used, including:

- **Failover affinity.** One change is to the way failover affinity works in a stretch cluster. For instance, in the four-node cluster previously mentioned, suppose that you need to take Node2 down for maintenance. With site-aware failover clusters, you can help ensure that the services running on Node2 will fail over to a node on the same site, rather than failing over to a node on the other site.
- **Cross-site heartbeating.** Previously, the default heartbeat settings were configured based on subnets that could cause issues depending on your network configuration. With site-aware failover clusters, you can easily configure this based on location and regardless of network configuration. The subnet heartbeat settings are now relevant for site-aware failover clusters only to determine the heartbeat settings for the same-site nodes.
- **Preferred site configuration.** Previously, with a two-site cluster, you could configure one of the sites with **LowerQuorumPriorityNodeID**. This has been deprecated in Windows Server 2016 and replaced with the preferred site configuration based on the site-aware failover clusters. After configuring the nodes into preferred sites, you can then identify a preferred site by running the following command in Windows PowerShell:

```
(Get-Cluster).PreferredSite = "<Site number that you want to be preferred>"
```

For example, after you run the preceding configuration of the four-node cluster, you can set Site 1 as the preferred site by running the following command:

```
(Get-Cluster).PreferredSite = 1
```

This allows you to identify what nodes the roles should attempt to come up on first. You can also take it a step further and configure preferred sites based on cluster groups so that different groups can have different preferred sites. You can also configure this through Windows PowerShell by using the following command:

```
(Get-ClusterGroupGroupName).PreferredSite = "<Site number that you want to be preferred>"
```


Choosing quorum witness

When creating a stretch cluster across geographically dispersed nodes in Windows Server 2016, we recommend that you use an Azure Cloud Witness whenever possible. However, in some cases, it might be more practical to use a file share witness.

Regardless of the witness selected, you should use dynamic witness mode, which is the default mode in both Windows Server 2016 and Windows Server 2012 R2. A stretch cluster spreads across multiple datacenters, and it is possible that an entire datacenter might go offline. In this situation, the quorum might lose half or more of the cluster nodes at once and have some servers in maintenance mode, so it is important that you use dynamic quorum and dynamic witness to avoid a shutdown of the cluster.

- **File share witness:**
 - Requires three or more datacenter locations
 - Is available in Windows Server 2012 R2 and Windows Server 2016
- **Azure cloud witness:**
 - Requires two datacenter locations
 - Requires an Internet connection for all nodes
 - Is available only in Windows Server 2016
- **No witness:**
 - Not recommended
 - Used for manual failover (disaster recovery site)



Note: In a stretch cluster, shared storage is not accessible to the nodes at different locations. A disk witness is not a suggested witness selection for this scenario.

File share witness

The major issue with a file share witness for most scenarios is the minimum requirement of three datacenters to create the witness. However, if you are working in an environment where three or more datacenters are already in operation, creating a file share witness on a share at one of the locations might be the quickest and easiest witness option. A file share witness does require a file share that all nodes in the cluster can access by using the SMB protocol. A file share witness does not keep a copy of the cluster database. A file share witness is available on both Windows Server 2016 and Windows Server 2012 R2.

Azure Cloud Witness

An Azure Cloud Witness builds on the foundation of the file share witness. An Azure Cloud Witness uses the same basic format as the file share witness regarding its arbitration logic and does not keep a copy of the cluster database. However, rather than requiring a share and writing over SMB, Azure Cloud Witness uses Blob storage and the REST-based API for Azure Storage. This configuration does require an Azure account and Internet connectivity for all the nodes at each site. An Azure Cloud Witness is available only on Windows Server 2016.

No witness

You can also configure a cluster to not use any witness. We recommend that you avoid this solution; however, it is supported to prevent split-brain syndrome. You perform this configuration in Windows Server 2016 by using site-aware clustering. You can also configure no witness for manual failovers—for instance, in disaster recovery scenarios. You can accomplish this by removing the votes for the nodes at the disaster recovery site, manually forcing quorum for the site that you want to bring online, and then preventing quorum at the site that you want to keep offline.

Considerations for deploying a stretch cluster

Stretch clusters are not appropriate for every application or every business. When you design a stretch cluster solution with a hardware vendor, clearly identify the business requirements and expectations. Not every scenario that involves more than one location is appropriate for a stretch cluster.

Stretch clustering is a highly available strategy that primarily focuses on hardware platform availability. However, specific stretch cluster configurations and deployments have availability ramifications, ranging from users' ability to connect to the application to the quality of application performance. Stretch clustering can be a powerful solution for managing planned and unplanned downtime, but you must examine its benefits against all the dimensions of application availability.

Stretch clusters do require more overhead than local clusters. Instead of a local cluster in which each node of the cluster attaches to the mass storage device, each site of a stretch cluster must have comparable storage. In addition, you must consider setting up replication among the cluster sites. Storage Replica in Windows Server 2016 provides storage-agnostic replication. However, you must also consider paying for additional network bandwidth among the sites and developing the management resources with your organization to efficiently administer your stretch cluster. Carefully consider the quorum witness that you use to help ensure that it will maintain functionality in the event of a failure and the location of the available cluster votes.



Note: Application data such as SQL Server, Hyper-V, Exchange Server, and AD DS should use their individual application stretch configurations (Hyper-V Replica, Database Availability Groups, and so on).

When deploying stretch clusters:

- Ensure that the business requirements are met
- Use storage replication among sites:
 - Use hardware vendor (Windows Server 2012 R2 or earlier)
 - Use Storage Replica (Windows Server 2016)
- Choose the correct quorum witness to properly maintain functionality in the event of failures
- Choose the correct storage replication solution to meet these needs

Considerations for stretch cluster failover and failback

When you establish a stretch clustering structure, it is very important that you define a procedure for the tasks that you should perform in the case of a site disaster. In addition, you should define a procedure for the tasks that you should perform for failback.

In most cases, the failover of critical services to another site does not occur automatically but consists of a manual or partly manual procedure. When defining your failover process, you should consider the following factors:

When implementing stretch clusters in disaster recovery scenarios, consider the following:

- Failover time
- The services for failover
- Quorum maintenance
- The storage connection
- Published services and name resolution
- Client connectivity
- The failback procedure

- Failover time. You must decide how long you should wait before you pronounce a disaster and start the failover process to another site.

- The services for failover. You should clearly define the critical services, such as AD DS, DNS, and DHCP, that should fail over to another site. It is not enough to have a cluster designed to fail over to another site. Failover clustering requires that you have Active Directory services run on a second site. You cannot make all the necessary services highly available by using failover clustering, so you must consider other technologies to achieve that result. For example, for AD DS and DNS, you can deploy additional domain controllers and DNS servers or VMs on a second site.
- Quorum maintenance. It is important to design the quorum model in a way that each site has enough votes for maintaining the cluster functionality. If that is not possible, you can use options such as forcing a quorum or dynamic quorum (in Windows Server 2016 and Windows Server 2012 R2) to establish a quorum in case of a disaster.
- The storage connection. A stretch cluster usually requires that you have storage available at each site. Because of this, you should carefully design storage replication and the procedure for how to fail over to secondary storage in the case of a disaster.
- Published services and name resolution. If you have services published to your internal or external users (such as email and webpages), in some cases, failover to another site requires name or IP address changes. If that is the case, you should have a procedure for changing DNS records in the internal or a public DNS. To reduce the downtime, we recommended that you reduce the Time to Live (TTL) on critical DNS records.
- Client connectivity. A failover plan must include a design for client connectivity in the case of a disaster. This includes both internal and external clients. If your primary site fails, you should have a way for your clients to connect to a second site.
- The failback procedure. You should plan and implement a failback process to perform after the primary site comes back online. Failback is as important as a failover, because if you perform it incorrectly, you can cause data loss and service downtime. Because of this, you must clearly define the steps for how to perform failback to a primary site without data loss or corruption. Very rarely is the failback process automated, and it usually happens in a very controlled environment.

Establishing a stretch cluster consists of much more than defining the cluster, cluster role, and quorum options. When you design a stretch cluster, you should consider the much larger picture of failover as part of a disaster recovery strategy. Windows Server 2016 has several technologies that can help with failover and failback, but you should also consider the other technologies that participate in your infrastructure. In addition, each failover and failback procedure greatly depends on the services implemented in a cluster.

Question: What added features does enabling site-aware clustering in a stretch cluster provide?

Check Your Knowledge

Question	
You have only two datacenter locations with a Windows Server 2016 stretch cluster built across both sites. What type of dynamic witness is the best to create?	
Select the correct answer.	
<input type="checkbox"/>	File share witness
<input type="checkbox"/>	Azure Cloud Witness
<input type="checkbox"/>	Disk witness
<input type="checkbox"/>	No witness

Lab: Implementing failover clustering

Scenario

As the business of A. Datum Corporation grows, it is becoming increasingly important that many of the applications and services on the network be available at all times. A. Datum Corporation has many services and applications that must be available to internal and external users who work in different time zones around the world. Many of these applications cannot be made highly available by using Network Load Balancing (NLB). Therefore, you have to use a different technology to make these applications highly available.

As one of the senior network administrators at A. Datum Corporation, you are responsible for implementing failover clustering on the servers running Windows Server 2016 to provide high availability for network services and applications. You are also responsible for planning the failover cluster configuration and deploying applications and services on the failover cluster.

Objectives

After completing this lab, you will be able to:

- Configure a failover cluster.
- Deploy and configure a highly available file server on the failover cluster.
- Validate the deployment of the highly available file server.
- Configure CAU on the failover cluster.

Lab Setup

Estimated Time: 60 minutes

Virtual machines: **20743A-LON-DC1**, **20743A-LON-SVR1**, **20743A-LON-SVR2**, **20743A-LON-SVR3**, and **MSL-TMG1**

User name: **Adatum\Administrator**

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

For this lab, you need to use the available VM environment. Before you begin the lab, you must complete the following steps:

1. On the host computer, start **Hyper-V Manager**.
2. In Hyper-V Manager, click **20743A-LON-DC1**, and then in the **Actions** pane, click **Start**.
3. In the **Actions** pane, click **Connect**. Wait until the VM starts.
4. Sign in by using the following credentials:
 - User name: **Administrator**
 - Password: **Pa\$\$w0rd**
 - Domain: **Adatum**
5. Repeat steps 2 through 4 for **20743A-LON-SVR1**, **20743A-LON-SVR2**, and **20743A-LON-SVR3**.
6. To provide access to the Internet start **MSL-TMG1**.

Exercise 1: Configuring iSCSI storage

Scenario

A. Datum Corporation has important applications and services that they want to make highly available. Some of these services cannot use NLB. Therefore, you have decided to implement failover clustering. You decide to use iSCSI storage for failover clustering. First, you will configure iSCSI storage to support your failover cluster. Note that the iSCSI role service has already been installed on LON-SVR1.

The main task for this exercise is as follows:

1. Configure the iSCSI targets.

► Task 1: Configure the iSCSI targets

1. On **LON-SVR1**, start **Server Manager**.
2. In **Server Manager**, navigate to **File and Storage Services**, and then navigate to **iSCSI**.
3. Create an iSCSI virtual disk with the following values:
 - Storage location: **C:**
 - Disk name: **iSCSIDisk1**
 - Size: **5 GB**
4. Create a new iSCSI target with the following values:
 - Target name: **lon-svr1**
 - Two iSCSI initiators with the following IP addresses:
 - IP Address: **172.16.0.12**
 - IP Address: **172.16.0.13**
5. Repeat step 4 to create two more iSCSI virtual disks with the disk names **iSCSIDisk2** and **iSCSIDisk3**.

Results: After completing this exercise, you should have successfully installed an iSCSI Target Server.

Exercise 2: Configuring a failover cluster

Scenario

You will now configure a failover cluster. First, you will implement the core components for failover clustering. Next, you will validate the cluster, and then you will create the failover cluster.

The main tasks for this exercise are as follows:

1. Connect clients to the iSCSI targets.
2. Install the Failover Clustering feature.
3. Validate and create a failover cluster.

► Task 1: Connect clients to the iSCSI targets

1. On **LON-SVR2**, open **Server Manager**, start the **iSCSI Initiator**, and then configure **Discover Portal** with the IP address **172.16.0.11**.
2. In the **Targets** list, connect to the discovered target.

3. Repeat steps 1 and 2 on **LON-SVR3**.
4. On **LON-SVR2**, open **Disk Management**.
5. Bring online and initialize the three new 5 GB disks.
6. Make a simple volume on each disk, and format them with NTFS. Label the disks **Data**, **Data2**, and **Data3**.
7. On **LON-SVR3**, open **Disk Management**, and then bring the three new disks online.

► **Task 2: Install the Failover Clustering feature**

1. On **LON-SVR2**, install the **Failover Clustering** feature by using **Server Manager**.
2. On **LON-SVR3**, install the **Failover Clustering** feature by using **Server Manager**.

► **Task 3: Validate and create a failover cluster**

1. On **LON-SVR2**, open the **Failover Cluster Manager** console.
2. Start the **Validate a Configuration Wizard**.
3. Use **LON-SVR2** and **LON-SVR3** as nodes for the tests.
4. Run all the tests.
5. Review the results. No errors should appear, but some warnings are expected.
6. Choose to create the cluster now.
7. Specify **Cluster1** as the cluster name.
8. Specify the IP address as **172.16.0.125**
9. Confirm the information, and then start the cluster creation.

Results: After this exercise, you should have installed and configured the Failover Clustering feature.

Exercise 3: Deploying and configuring a highly available file server

Scenario

At A. Datum Corporation, file services are important services that must be made highly available. After you have created a cluster infrastructure, you decide to configure a highly available file server and then implement settings for failover and failback.

The main tasks for this exercise are as follows:

1. Add the file server application to the failover cluster.
2. Add a shared folder to a highly available file server.
3. Configure the failover and failback settings.

► **Task 1: Add the file server application to the failover cluster**

1. On **LON-SVR2**, open the **Failover Cluster Manager** console.
2. In the **Storage** node, click **Disks**, and then verify that three cluster disks are online.
3. Add **File Server** as a cluster role. Select the **File Server for general use** option.

4. Specify **AdatumFS** for the **Client Access Name**, **172.16.0.130** for the address, and **Cluster Disk 2** for the storage.
5. Close the wizard.

► **Task 2: Add a shared folder to a highly available file server**

1. On **LON-SVR3**, from **Server Manager**, open the **Failover Cluster Manager** console.
2. Start the **New Share Wizard**, and then add a new shared folder to the **AdatumFS** cluster role.
3. Specify the file share profile as **SMB Share - Quick**.
4. Accept the default values on the **Select the server and the path for this share** page.
5. Name the shared folder **Docs**.
6. Accept the default values on the **Configure share settings** and **Specify permissions to control access** pages.
7. At the end of the **New Share Wizard**, create the share.

► **Task 3: Configure the failover and failback settings**

1. On **LON-SVR3**, in the **Failover Cluster Manager** console, open the properties for the **AdatumFS** cluster role.
2. Enable failback for between 4 and 5 hours.
3. Select both **LON-SVR2** and **LON-SVR3** as the preferred owners.
4. Move **LON-SVR3** to be first in the preferred owners list.

Results: After this exercise, you should have configured a highly available file server.

Exercise 4: Validating the deployment of the highly available file server

Scenario

In the process of implementing a failover cluster, you want to perform failover and failback tests. In addition, you want to change the disk witness in the quorum.

The main tasks for this exercise are as follows:

1. Validate the highly available file server deployment.
2. Validate the failover and quorum configuration for the File Server role.

► **Task 1: Validate the highly available file server deployment**

1. On **LON-DC1**, open **File Explorer**, and then attempt to access the **\\AdatumFS** location. Verify that you can access the **Docs** folder.
2. Create a test text document inside this folder.
3. On **LON-SVR2**, in the **Failover Cluster Manager** console, move **AdatumFS** to the second node.
4. On **LON-DC1**, in File Explorer, verify that you can still access the **\\AdatumFS** location.

► Task 2: Validate the failover and quorum configuration for the File Server role

1. On **LON-SVR2**, determine the current owner for the **AdatumFS** role.
2. Stop the **Cluster** service on the node that is the current owner of the **AdatumFS** role.
3. Try to access **\\AdatumFS** from **LON-DC1** to verify that **AdatumFS** has moved to another node and that the **\\AdatumFS** location is still available.
4. Start the **Cluster** service on the node on which you stopped it in step 2.
5. Browse to the **Disks** node, and then take the disk marked **Disk Witness in Quorum** offline.
6. Verify that **AdatumFS** is still available by trying to access it from **LON-DC1**.
7. Bring the disk witness online.
8. Open **Cluster Quorum Settings**.
9. Choose to perform advanced configuration.
10. Change the witness disk to **Cluster Disk 3**. Do not make any other changes.

Results: After this exercise, you should have tested the failover scenarios.

Exercise 5: Configuring CAU on the failover cluster**Scenario**

Before Windows Server 2012, implementing updates to servers with critical services caused unwanted downtime. To enable cluster updating with zero downtime, you want to implement the CAU feature and test updates for cluster nodes.

The main tasks for this exercise are as follows:

1. Configure CAU.
2. Update the failover cluster and configure self-updating.
3. Prepare for the next module.

► Task 1: Configure CAU

1. On **LON-DC1**, from **Server Manager**, install the **Failover Clustering** feature.
2. On **LON-SVR2**, open the **Windows Firewall with Advanced Security** window, and then ensure that the following two inbound rules are enabled:
 - **Inbound Rule for Remote Shutdown (RPC-EP-In)**
 - **Inbound Rule for Remote Shutdown (TCP-In)**
3. Repeat step 2 on **LON-SVR3**.
4. On **LON-DC1**, from **Server Manager**, open **Cluster-Aware Updating**.
5. Connect to **Cluster1**.
6. Preview the updates available for the nodes in **Cluster1**. (Note that there may be no updates currently. Updates will show as available.)

► **Task 2: Update the failover cluster and configure self-updating**

1. On **LON-DC1**, start the update process for **CLUSTER1** by clicking **Apply updates to this cluster**.
2. Accept the default values in the update wizard.
3. Wait until the updating process completes.



Note: This updating process might require a restart of both nodes. The process is finished when both display a value of **Succeeded** in the **Last Run status** column.

4. On **LON-SVR2**, open **Cluster-Aware Updating**, and then connect to **CLUSTER1**.
5. Select the **Configure cluster self-updating options** option.
6. Select to add the **CAU** clustered role with the self-updating mode enabled to this cluster.
7. Configure self-updating to be performed **weekly**, on **Sundays**, at **4:00 AM**.
8. Apply the settings and close the wizard.

► **Task 3: Prepare for the next module**

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

1. On the host computer, start **Hyper-V Manager**.
2. In the **Virtual Machines** list, right-click **20743A-LON-DC1**, and then click **Revert**.
3. In the **Revert VM** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20743A-LON-SVR1**, **20743A-LON-SVR2**, **20743A-LON-SVR3**, and **MSL-TMG1**.

Results: After this exercise, you should have configured CAU.

Question: What information do you need for planning a failover cluster implementation?

Question: After running the **Validate a Configuration Wizard**, how can you resolve the network communication single point of failure?

Question: In which situations might it be important to enable failback for a clustered application during a specific time?

Module Review and Takeaways

Review Questions

Question: Why is using a disk-only quorum configuration generally not a good idea?

Question: What is the purpose of CAU?

Question: What is the main difference between synchronous and asynchronous replication in a stretch cluster scenario?

Question: What is an enhanced feature in stretch clusters in Windows Server 2016?

Real-world Issues and Scenarios

Your organization is considering the use of a geographically dispersed cluster that includes an alternate datacenter. Your organization has only a single physical location together with an alternate datacenter. Can you provide an automatic failover in this configuration?

Answer: Yes, you can provide an automatic failover in this configuration. To provide an automatic failover, you must configure an Azure Cloud Witness.

Tools

The following is a list of the tools that this module references:

- Failover Cluster Manager console
- Cluster-Aware Updating console
- Windows PowerShell
- Server Manager
- Internet Small Computer System Interface (iSCSI) initiator
- Disk Management

Best Practices

- Try to avoid using a quorum model that depends only on the disk for Hyper-V high availability or Scale-Out File Server.
- Perform regular backups of the cluster configuration.
- Ensure that in the case of one node failure, other nodes can manage the load.
- Carefully plan stretch clusters.

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
The Cluster Validation Wizard reports an error.	
The Create Cluster Wizard reports that not all nodes support the wanted clustered role.	
You cannot create a print server cluster.	

Module 12

Implementing failover clustering with Windows Server 2016 Hyper-V

Contents:

Module Overview	12-1
Lesson 1: Overview of the integration of Hyper-V Server 2016 with failover clustering	12-2
Lesson 2: Implementing Hyper-V virtual machines on failover clusters	12-7
Lesson 3: Implementing Windows Server 2016 Hyper-V virtual machine migration	12-20
Lesson 4: Implementing Hyper-V Replica	12-24
Lab: Implementing failover clustering with Windows Server 2016 Hyper-V	12-29
Module Review and Takeaways	12-36

Module Overview

One benefit of implementing server virtualization is the opportunity to provide high availability, both for applications or services that have built-in high availability functionality, and for applications or services that do not provide high availability in any other way. With the Windows Server 2016 Hyper-V technology, failover clustering, and Microsoft System Center 2012 R2 Virtual Machine Manager (VMM), you can use several different options to configure high availability.

In this module, you will learn how to implement failover clustering in a Hyper-V scenario to achieve high availability for a virtual environment.



Note: Many of the features that this module describes are also available in Windows Server 2012 R2 and Windows Server 2012. This module explicitly calls out the features that are new to Windows Server 2016.

Objectives

After completing this module, you will be able to:

- Describe how Hyper-V integrates with failover clustering.
- Implement Hyper-V virtual machines in failover clusters.
- Implement Hyper-V virtual machine movement.
- Implement Hyper-V Replica.

Lesson 1

Overview of the integration of Hyper-V Server 2016 with failover clustering

Failover clustering makes applications or services highly available. To make virtual machines highly available in a Hyper-V environment, you should implement failover clustering on the Hyper-V host computers. This lesson summarizes the high availability options for Hyper-V–based virtual machines, focuses on how failover clustering works, and shows how to design and implement failover clustering for Hyper-V.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the options for making applications and services highly available.
- Describe how failover clustering works with Hyper-V nodes.
- Describe the improvements to the functionality of Hyper-V with failover clustering.
- Describe the best practices for implementing high availability in a virtual environment.

Options for making applications and services highly available

Most organizations have some applications that are business critical and must be highly available. To make an application or a service highly available, you must deploy it in an environment that provides redundancy for all components that the application requires. You can choose between several options to provide high availability for virtual machines and the services hosted within virtual machines. You can:

- Implement virtual machines as a clustered role (host clustering).
- Implement clustering inside virtual machines (guest clustering).
- Use Network Load Balancing (NLB) inside virtual machines.

High-availability options	Description
Host clustering	<ul style="list-style-type: none"> • Virtual machines are highly available • Does not require virtual machine operating system or application to be cluster-aware
Guest clustering	<ul style="list-style-type: none"> • Virtual machines are failover cluster nodes • Virtual machine applications must be cluster aware • Requires iSCSI or virtual Fibre Channel interface for shared storage connections
NLB	<ul style="list-style-type: none"> • Virtual machines are NLB cluster nodes • Use for web-based applications

Host clustering

By using host clustering you can configure a failover cluster when you use 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 clustering 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 is still highly available.

Some examples of non-cluster-aware applications are a print server, or a proprietary network-based application such as an accounting application. Should the host node that controls the virtual machine unexpectedly become unavailable, the secondary host node takes control and restarts or resumes the virtual machine as quickly as possible. You also 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 from one node to another while updating the host management Windows Server 2016 operating system.

The applications or services running in the virtual machine do not have to be compatible with failover clustering, and they do not have to be aware that the virtual machine is clustered. Because the failover is at the virtual machine-level, there are no dependencies on software that is installed inside the virtual machine.

Guest clustering

You configure guest failover clustering very similarly to physical-server failover clustering, except that the cluster nodes are virtual machines. In this scenario, you create two or more virtual machines and install and implement failover clustering within the guest operating systems. The application or service is then able to take advantage of high availability between the virtual machines. Because you implement failover clustering within the guest operating system of each virtual machine node, you can locate the virtual machines on a single host. This is a quick and cost-effective configuration in a test or staging environment.

For production environments, however, you can protect the application or service more robustly if you deploy the virtual machines on separate failover clustering-enabled Hyper-V host computers. With failover clustering implemented at both the host and virtual machine levels, you can restart the resource regardless of whether the node that fails is a virtual machine or a host. Such high-availability configurations for virtual machines running mission-critical applications in a production environment are considered optimal.

You should consider several factors when you implement guest clustering:

- The application or service must be failover cluster-aware. This includes any of the Windows Server 2016 services that are cluster-aware, as well as any applications, such as clustered Microsoft SQL Server and Microsoft Exchange Server.
- Hyper-V virtual machines in Windows Server 2016 can use Fibre Channel-based connections to shared storage, or you can implement Internet Small Computer System Interface (iSCSI) connections from the virtual machines to the shared storage. You also can use the shared virtual hard disk feature to provide shared storage for virtual machines.

You should deploy multiple network adapters on the host computers and the virtual machines. Ideally, you should dedicate a network connection to the iSCSI connection (if you are using this method to connect to storage), to the private network between the hosts, and to the network connection that the client computers use.

NLB

NLB works with virtual machines in the same way that it works with physical hosts. It distributes IP traffic to multiple instances of a TCP/IP service, such as a web server that is running on a host within the NLB cluster. NLB transparently distributes client requests among the hosts, and it enables the clients to access the cluster by using a virtual host name or a virtual IP address. From the client computer's perspective, the cluster appears to be a single server that answers these client requests. As enterprise traffic increases, you can add another server to the cluster.

Therefore, NLB is an appropriate solution for resources that do not have to accommodate exclusive read or write requests. Examples of NLB-appropriate applications are web-based front-end virtual machines to database applications or Exchange Server Client Access servers.

When you configure an NLB cluster, you must install and configure the application on all virtual machines that will participate in the NLB cluster. After you configure the application, you install the NLB feature in Windows Server 2016 within each virtual machine's guest operating system (not on the Hyper-V hosts), and then configure an NLB cluster for the application. Earlier versions of Windows Server also support

NLB, so that the guest operating system is not limited to only Windows Server 2016; however, you should use the same operating system versions within one NLB cluster. Similar to a guest cluster across hosts, the NLB resource typically benefits from overall increased input/output (I/O) performance when the virtual machine nodes are located on different Hyper-V hosts.



Note: As with earlier versions of Windows Server, you should not implement NLB and failover clustering within the same operating system because the two technologies conflict with each other.

How does a failover cluster work with Hyper-V nodes?

When you implement failover clustering and configure virtual machines as highly available resources, the failover cluster treats the virtual machines like any other application or service. For example, if there is host failure, failover clustering acts to restore access to the virtual machine as quickly as possible on another host in the cluster. Only one node runs the virtual machine at a time. However, you also can move the virtual machine to any other node in the same cluster as part of a planned migration.

The failover process transfers the responsibility of providing access to resources in a cluster from one node to another. Planned failover (also known as *switchover*) can occur when an administrator intentionally moves resources to another node for maintenance or other reasons, or when unplanned downtime of one node occurs because of hardware failure or other reasons.

The failover process consists of the following steps:

1. The node where the virtual machine is running owns the clustered instance of the virtual machine, controls access to the shared bus or iSCSI connection to the cluster storage, and owns any disks or logical unit numbers (LUNs) that are assigned to the virtual machine. All of the nodes in the cluster use a private network to send regular signals, known as *heartbeat signals*, to one another. The heartbeat indicates that a node is functioning and communicating on the network. The default heartbeat configuration specifies that each node send a heartbeat over TCP/UDP port 3343 each second (or 1,000 milliseconds [ms]).
2. Failover initiates when the node that is hosting the virtual machine does not send regular heartbeat signals over the network to the other nodes. By default, this is five consecutively missed heartbeats (or 5,000 ms elapsed). Failover might occur because of a node failure or network failure. When heartbeat signals stop arriving from the failed node, one of the other nodes in the cluster begins taking over the resources that the virtual machines use.

You define the one or more nodes that could take over by configuring the **Preferred Owner** and the **Possible Owners** properties. The **Preferred Owner** property specifies the hierarchy of ownership if a resource has more than one possible failover node. By default, all nodes are members of Possible Owners. Therefore, removing a node as a Possible Owner excludes it from taking over the resource in a failure situation.

For example, suppose that you implement a failover cluster by using four nodes. However, you configure only two nodes as Possible Owners. In a failover event, the third node might still take over the resource if neither of the Preferred Owners is online. Although you did not configure the fourth node as a Preferred Owner, as long as it remains a member of Possible Owners, the failover cluster uses it to restore access to the resource, if necessary.

Resources come online in order of dependency. For example, if the virtual machine references an iSCSI LUN, access to the appropriate host bus adapters (HBAs), network (or networks), and LUNs are 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 might not notice.

3. You also 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 all of the resources associated with that instance offline, moves the instance, and then brings all of the resources in the instance back online.

Failover clustering with Windows Server 2016 Hyper-V features

There are many improvements to the functionality of Hyper-V with failover clustering since the introduction of Hyper-V in Windows Server 2008. Windows Server 2016 continues to build on Hyper-V with failover clustering with some updated features and improvements in the following areas:

- Maximum nodes and virtual machines supported. Failover clustering supports up to 64 nodes and 8,000 virtual machines per cluster (and 1024 virtual machines per node).
- File share storage. Windows Server 2012 introduced the possibility to store virtual machines on server message block (SMB) file shares in a file server cluster. This is a way to provide shared storage that is accessible by multiple clusters by providing the ability to move virtual machines between clusters without moving the storage. To enable this feature, deploy a file server cluster role and select Scale-Out File Server for application data.
- Shared virtual disk. Windows Server 2012 R2 introduced the ability to use a .vhdx as a shared virtual hard disk for guest clusters. Windows Server 2016 introduced improved features to the shared disks and introduced a new disk format, .vhds (VHD Set).
- Rolling Hyper-V cluster upgrades. In Windows Server 2016, you can upgrade the nodes one at a time when upgrading from Windows Server 2012 R2. After upgrading all nodes in a Hyper-V cluster, you can upgrade the functional level of the entire cluster.
- Virtual machine configuration version. Windows Server 2016 builds on the rolling upgrades by not updating the configuration version of the virtual machines automatically. You can now manually update the virtual machine configuration version. This allows a virtual machine to migrate back and forth from both Windows Server 2016 and Windows Server 2012 R2 until rolling upgrades are completed and you are ready to upgrade to the version for Windows Server 2016 and take advantage of the new features for Windows Server 2016 Hyper-V.

- Failover clustering with Windows Server 2016 Hyper-V features:
 - Maximum nodes and virtual machine support
 - File share storage:
 - vhds (Windows Server 2016 only)
 - .vhdx (Windows Server 2012 R2 and Windows Server 2016 only)
 - Shared virtual disk
 - Rolling Hyper-V cluster upgrades
 - Virtual machine configuration version

Best practices for implementing high availability in a virtual environment

After you determine which applications you will deploy on highly available failover clusters, you plan and deploy the failover clustering environment. Follow these recommendations when you implement the failover cluster:

- Plan for failover scenarios
- Plan the network design for failover clustering
- Plan the shared storage for failover clustering
- Use the default failover cluster quorum mode
- Deploy standardized Hyper-V hosts
- Develop standard management practices

- Plan for failover scenarios. When you design the hardware requirements for the Hyper-V hosts, ensure that you include the hardware capacity that is required when hosts fail. For example, if you deploy a six-node cluster, you must determine the number of host failures that you want to accommodate. If you decide that the cluster must sustain the failure of two nodes, then the four remaining nodes must have the capacity to run all of the virtual machines in the cluster.
- Plan the network design for failover clustering. To optimize the failover cluster performance and failover, you should dedicate a fast network connection for internode communication. Similar to earlier versions, this network should be logically and physically separate from the network segment (or segments) that clients use to communicate with the cluster. You also can use this network connection to transfer virtual machine memory during a live migration. If you are using iSCSI for any virtual machines, ensure that you also dedicate a network connection to the iSCSI network connection. This also applies if you are using SMB 3.0 shares for virtual machines.
- Plan the shared storage for failover clustering. When you implement failover clustering for Hyper-V, the shared storage must be highly available. If the shared storage fails, the virtual machines will all fail, even if the physical nodes are functional. To ensure storage availability, plan for redundant connections to the shared storage and for redundant array of independent disks (RAID) on the storage device. If you decide to use a shared virtual hard disk, ensure that the shared disk is located on a highly available resource such as a Scale-Out File Server.
- Use the recommended failover cluster quorum mode. For failover clustering in Windows Server 2016, the default is Dynamic Quorum mode and Dynamic witness. You should not modify the default configuration unless you understand the implications of doing so.
- Deploy standardized Hyper-V hosts. To simplify the deployment and management of the failover cluster and Hyper-V nodes, develop a standard server hardware and software platform for all nodes.
- Develop standard management practices. When you deploy multiple virtual machines in a failover cluster, you increase the risk that a single mistake might shut down a large part of the server deployment. For example, if an administrator accidentally configures the failover cluster incorrectly and the cluster fails, all virtual machines in the cluster will be offline. To avoid this, develop and thoroughly test standardized instructions for all administrative tasks.

Question: Why is using shared storage a best practice in Windows Server Hyper-V failover clustering?

Question: You have two clusters. One is a Windows Server 2016 cluster (Cluster1), and the other is a mixed mode cluster of Windows Server 2012 R2 and Windows Server 2016 (Cluster2) that is in the process of upgrading, but the upgrade is not completed. You also have two virtual machines called VM1 and VM2. VM1 and VM2 need to migrate back and forth between Cluster1 and Cluster2 occasionally. Should you upgrade the configuration version on VM1?

Lesson 2

Implementing Hyper-V virtual machines on failover clusters

Implementing highly available virtual machines is somewhat different from implementing other roles in a failover cluster. Failover clustering in Windows Server 2016 provides many features for Hyper-V clustering, in addition to tools for managing the high availability of virtual machines. In this lesson, you will learn about how to implement highly available virtual machines.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the components of a Hyper-V cluster.
- Describe the prerequisites for implementing Hyper-V failover clusters.
- Describe how to implement Hyper-V virtual machines on a failover cluster.
- Describe how to configure Clustered Shared Volumes (CSVs).
- Explain how to configure a shared virtual hard disk.
- Explain how to implement Scale-Out File Servers for virtual machine storage.
- Describe the considerations for implementing Hyper-V cluster.
- Explain how to maintain and monitor virtual machines in clusters.

Components of Hyper-V clusters

As a role, Hyper-V has some specific requirements for cluster components. To form a Hyper-V cluster, you must have at least two physical nodes. Whereas other clustered roles (such as Dynamic Host Configuration Protocol [DHCP] or file server) allow nodes to be virtual machines, Hyper-V nodes in most production environments should be physical servers. However, in Windows Server 2016 you can enable nested virtualization, which allows you to use a guest virtual machine for configuring Hyper-V host.

Hyper-V cluster components include:

- Cluster nodes
- Cluster networks
- Virtual networks
- Storage for virtual machines
- Virtual machines

In addition to having nodes, you also must have physical and virtual networks. Failover clustering requires a cluster network interface for internal cluster communication and a network interface for clients. You also can implement a storage network separately, depending on the type of storage you are using. As a reminder, specific to the Hyper-V role, you also should consider virtual networks for clustered virtual machines. Creating the same virtual networks on all physical hosts that participate in one cluster is very important. Failure to do so causes a virtual machine to lose network connectivity when it moves from one host to another.

Storage is an important component of virtual machine clustering. You can use any type of storage that Windows Server 2016 failover clustering supports. We recommend that you configure storage as a Clustered Shared Volume (CSV), as discussed in a following topic.

When using host clustering, virtual machines are also components of a Hyper-V cluster. In Failover Cluster Manager, you can create new highly available virtual machines, or you can make existing virtual machines highly available. In both cases, the virtual machine storage location must be on shared storage that is accessible to both nodes. You might not want to make all virtual machines highly available. In Failover Cluster Manager, you can choose which virtual machines are part of a cluster configuration.

Prerequisites for implementing Hyper-V failover clusters

To deploy a Hyper-V cluster, you must ensure that you meet the hardware, software, account, and network-infrastructure requirements. The following sections detail these requirements.

Hardware requirements for failover clustering with Hyper-V

You must have the following hardware for a two-node failover cluster:

- Server hardware. Hyper-V on Windows Server 2016 requires an x64-based processor, hardware-assisted virtualization, and hardware-enforced Data Execution Prevention (DEP). As a best practice, the servers should have very similar hardware.

- Hardware requirements for cluster nodes and storage:
 - Server hardware
 - Network adapters
 - Storage adapters
 - Storage
- Software recommendations for cluster nodes:
 - Should run either Windows Server 2016 Standard, Datacenter, or Microsoft Hyper-V Server 2016 editions
 - Require the same software updates and service packs
 - Must be either a full installation or a Server Core installation
- Network infrastructure requirements:
 - Network settings and IP addresses
 - Private networks
 - DNS
 - Domain role
 - Account for administering the cluster



Note: Microsoft supports a failover cluster solution only if all of the hardware features are marked as *Certified for Windows Server*. In addition, the complete configuration (servers, network, and storage) must pass all tests in the **Validate This Configuration Wizard**, which is included in the Failover Cluster Manager snap-in.

- Network adapters. The network adapter hardware, like other features in the failover cluster solution, must be marked as *Certified for Windows Server*. To provide network redundancy, you can connect cluster nodes to multiple networks. Alternatively, to remove single points of failure, you can connect the nodes to one network that uses the following hardware:
 - Redundant switches
 - Teamed network adapters
 - Redundant routers
 - Any similar hardware

We recommend that you configure multiple physical network adapters on the host computer that you configure as a cluster node. One network adapter should connect to the private network that the inter-host communications use.

- Storage adapters. If you use Serial Attached SCSI (SAS) or Fibre Channel, the mass-storage device controllers in all clustered servers should be identical and should use the same firmware version. If you are using iSCSI, each clustered server should have one or more network adapters that are dedicated to the cluster storage. The network adapters that you use to connect to the iSCSI storage target should be identical, and you should use a Gigabit Ethernet or faster network adapter.

- Storage. You must use shared storage that is compatible with Windows Server 2016. If you deploy a failover cluster that uses a witness disk, the storage must contain at least two separate volumes (LUNs). One volume functions as the witness disk, and additional volumes contain the virtual machine files that the cluster nodes share. Storage considerations and recommendations include the following:
 - Use basic disks, not dynamic disks. Format the disks with the NTFS file system.
 - Use either master boot record (MBR) or globally unique identifier (GUID) partition table (GPT). Keep in mind that there is a two terabyte (TB) limit on the MBR disks. Most production clusters today use GPT volumes for storing virtual disks.
 - If you are using a storage area network (SAN), the miniport driver that the storage uses must work with the Microsoft Storport storage driver.
 - Consider using Microsoft Multipath I/O (MPIO) software. If your SAN uses a highly available network design with redundant components, deploy failover clusters with multiple host-bus adapters. To do this, use MPIO. This provides the highest level of redundancy and availability.
 - For environments without direct access to SAN or iSCSI storage, consider using shared virtual hard disks.

Software recommendations for using Hyper-V and failover clustering

The following are the software recommendations for using Hyper-V and failover clustering:

- All of the servers in a failover cluster should be running Windows Server 2016 Standard, Datacenter, or Microsoft Hyper-V Server 2016 editions. However, different editions are supported during a rolling upgrade failover cluster.
- All of the servers should have the same software updates and service packs.
- All of the servers should have the same drivers and firmware.

Network infrastructure requirements

The following network infrastructure is required for a failover cluster and an administrative account with the following domain permissions:

- Network settings and IP addresses. Use identical communication settings on all network adapters, including the speed, duplex mode, flow control, and media-type settings. Ensure that all network hardware supports the same settings.
- Private networks. If you use private networks that are not routed to your entire network infrastructure for communication between cluster nodes, ensure that each of these private networks uses a unique subnet.
- Domain Name System (DNS). The servers in the cluster must use DNS for name resolution. You should use the DNS dynamic update protocol.
- Domain role. All servers in the cluster must be in the same Active Directory Domain Services (AD DS) domain. As a best practice, all clustered servers should have the same domain role (either member server or domain controller). The recommended role is member server.
- Account for administering the cluster. When you first create a cluster or add servers to it, you must be signed in to the domain with an account that has administrator rights and permissions on all of the cluster's servers. In addition, if the account is not a Domain Admins account, the account must have the Create Computer Objects permission in the domain.


Implementing Hyper-V virtual machines on a failover cluster

To implement failover clustering for Hyper-V, you must complete the following high-level steps:


1. Install and configure the required versions of Windows Server 2016. After you complete the installation, configure the network settings, join the computers to an Active Directory domain, and configure the connection to the shared storage.
2. Configure the shared storage. You must use Disk Manager to create disk partitions on the shared storage.
3. Install the Hyper-V and Failover Clustering features on the host servers. You can use Server Manager in Microsoft Management Console (MMC) or Microsoft Windows PowerShell to do this.
4. Validate the cluster configuration. The **Validate This Cluster Wizard** checks all of the prerequisite components that are required to create a cluster and provides warnings or errors if any components do not meet the cluster requirements. Before you continue, resolve any issues that the **Validate This Cluster Wizard** identifies.

To implement a Hyper-V virtual machine on a failover cluster:

1. Install and configure Windows Server 2016
2. Configure shared storage
3. Install the Hyper-V and Failover Clustering features
4. Validate the cluster configuration
5. Create the cluster
6. Create a virtual machine in one of the cluster nodes
7. Make the virtual machine highly available (for an existing virtual machine)
8. Test virtual machine failover

 **Note:** Although it is possible to create a cluster without running cluster validation, we strongly recommend that you run the **Validate This Cluster Wizard** and resolve all issues before creating a cluster and putting it into production.

5. Create the cluster. When the components pass validation by the **Validate This Cluster Wizard**, you can create a cluster. When you configure the cluster, assign a cluster name and an IP address. A computer account for the cluster name is created in AD DS, and the IP address is registered in DNS. In Windows Server 2016, you also can create an Active Directory–detached cluster.

 **Note:** You can enable Clustered Shared Storage for the cluster only after you create the cluster and add eligible storage to it. If you want to use CSV, you should configure CSV before you move to the next step.

6. Create a virtual machine in one of the cluster nodes. When you create the virtual machine, ensure that all files associated with the virtual machine—including both the virtual hard disk and virtual machine configuration files—are stored on the shared storage. You can create and manage virtual machines in either Hyper-V Manager or Failover Cluster Manager. We recommend that you use the Failover Cluster Manager console for creating virtual machines. When you create a virtual machine by using Failover Cluster Manager, the virtual machine is made highly available automatically.
7. Make the virtual machine highly available only for existing virtual machines. If you created a virtual machine before implementing failover clustering, you should manually make it highly available. To make the virtual machine highly available, in the Failover Cluster Manager, chose a new service or application to make highly available. Failover Cluster Manager then presents a list of services and applications that can be made highly available. When you select the option to make virtual machines highly available, you can select the virtual machine that you created on shared storage.



Note: When you make a virtual machine highly available, you see a list of all virtual machines that are hosted in all cluster nodes, including virtual machines that are not stored on the shared storage. If you make a virtual machine that is not located on shared storage highly available, you receive a warning, but Hyper-V adds the virtual machine to the services and applications list. However, when you try to migrate the virtual machine to a different host, the migration will fail.

8. Test virtual machine failover. After you make the virtual machine highly available, you can migrate the computer to another node in the cluster. You can choose to perform a quick migration or a live migration. In most cases, you should perform a live migration to reduce downtime. We explain the differences later in this course.

Configuring CSVs

CSVs in a Windows Server 2016 failover cluster allow multiple nodes in the cluster to simultaneously have read-write access to the same disk that is provisioned as an NTFS volume and add them as storage to the cluster. When you use CSVs, clustered roles can fail over from one node to another more quickly and without requiring a change in drive ownership or dismounting and remounting a volume. CSVs also help in simplifying the management of a potentially large number of LUNs in a failover cluster.

CSVs provide a general-purpose, clustered file system, which is layered on NTFS. They are not restricted to specific clustered workloads; they are only supported for Hyper-V clusters and Scale-Out File Server clusters.

Although CSVs provide additional flexibility and reduce downtime, you do not need to configure and use CSV when you implement high availability for virtual machines in Hyper-V. You also can create clusters on Hyper-V by using the regular approach (with disks that are not assigned as CSVs). However, we recommend that you use CSVs because of the following advantages:

- Reduced LUNs for the disks. You can use CSVs to reduce the number of LUNs that your virtual machines require. When you configure a CSV, you can store multiple virtual machines on a single LUN, and multiple host computers can access the same LUN concurrently.
- Better use of disk space. Instead of placing each .vhd file on a separate disk with empty space so that the .vhd file can expand, you can oversubscribe disk space by storing multiple .vhd files on the same LUN.
- Single location for virtual machine files. You can track the paths of .vhd files and other files that virtual machines use. Instead of using drive letters or GUIDs to identify disks, you can specify the path names.

When you implement CSV, all added storage displays in the **\ClusterStorage** folder. The **\ClusterStorage** folder is created on the cluster node's system folder, and you cannot move it. This means that all Hyper-V hosts that are members of the cluster must use the same drive letter as their system drive, or virtual machine failovers fail.

- CSV benefits:
 - Fewer LUNs required
 - Better use of disk space
 - Virtual machine files are in a single logical location
 - No special hardware required
 - Increased resiliency
- To implement CSV:
 1. Create and format volumes on shared storage
 2. Add the disks to failover cluster storage
 3. Add the storage to the CSV

- No specific hardware requirements. Implementation of CSVs requires no specific hardware. You can implement CSVs on any supported disk configuration, and on either Fibre Channel or iSCSI SANs.
- Increased resiliency. CSVs increase resiliency because the cluster can respond correctly even if connectivity between one node and the SAN is interrupted or part of a network is down. The cluster reroutes the CSV traffic through an intact part of the SAN or network.

Implementing CSVs

After you create the failover cluster, you can enable a CSV for the cluster, and then add storage to the CSV.

Before you can add storage to the CSV, the LUN must be available as shared storage to the cluster. When you create a failover cluster, all of the shared disks that you configured in Server Manager are added to the cluster, and you can add them to a CSV. You also have the option to add storage to the cluster after the cluster is created. If you add more LUNs to the shared storage, you must first create volumes on the LUN, add the storage to the cluster, and then add the storage to the CSV.

We recommend that you should configure CSVs before you make any virtual machines highly available. However, after deployment, you can convert a virtual machine from regular disk access to CSV. The following considerations apply for conversion from regular disk access to CSV after deployment:

- Converting from regular disk access to CSV removes the LUN's drive letter (or *mount point*). This means that you must recreate all virtual machines that are stored on the shared storage. If you must keep the same virtual machine settings, consider exporting the virtual machines, switching to CSV, and then importing the virtual machines in Hyper-V.
- You cannot add shared storage to a CSV if it is in use. If you have a running virtual machine that is using a cluster disk, you must shut down the virtual machine and then add the disk to the CSV.

Configuring a shared virtual hard disk

In previous versions of Windows Server, you had to expose shared storage to the virtual machine to implement guest clustering. You could connect to the shared storage by using a virtual Fibre Channel interface or by using iSCSI. In some scenarios, it was a complicated task to perform if you did not have the support of appropriate drivers for the virtual Fibre Channel or if you did not have iSCSI support on the storage. In addition, in some scenarios, such as when the virtual machine was hosted at a hosting provider, administrators did not want to expose a storage layer to the virtual machine users or tenant administrators.

- Failover cluster runs inside virtual machines
- Shared virtual disk used as a shared storage:
 - Virtual machines do not need access to iSCSI or failover clustering SAN
 - Presented as a virtual SAS disk
 - Can be used only for data
- Requirements for shared virtual hard disk:
 - Must be in .vhdx or .vhds format
 - Connected by using virtual SCSI adapter
 - Stored on a Scale-Out File Server or CSV
- Supported operating system in virtual machine:
 - Windows Server 2012 or later

To address these issues, Windows Server 2016 provides an additional layer of abstraction for virtual machine cluster storage. It is possible to share a virtual hard disk (in .vhdx or .vhds format only) between two or more virtual machines and use that virtual hard disk as shared storage when building guest clusters. You can use the shared virtual hard disk as a witness disk or as a data disk in a cluster.

How does a shared virtual hard disk work?

You add shared virtual hard disks as SCSI drives in the virtual machine settings. They appear as virtual SAS disks in the virtual machine. You can add a shared virtual hard disk to any virtual machine with a supported guest operating system running on a Windows Server 2016 Hyper-V platform. When you use this technology, configuration of guest clustering is simplified because you have several options for providing shared storage for guest clusters. These options include shared virtual hard disk, Fibre Channel, SMB, storage spaces, and iSCSI storage. You can use shared virtual disks to provide storage for solutions such as SQL Server databases and file server clusters.

How to configure shared virtual hard disks

Shared virtual disks are used only in guest cluster scenarios. To configure a guest failover cluster that uses shared virtual hard disks, you must meet the followings requirements:

- At least a two-node Hyper-V failover host cluster.
- All servers must be running Windows Server 2012 R2 or later.
- All servers should belong to the same Active Directory domain.
- Configured shared storage resources must be available—for example, CSVs on block storage (such as clustered storage spaces) or a Scale-Out File Server cluster (running Windows Server 2012 R2 or later) with SMB 3.0 (for file-based storage).
- Sufficient memory, disk, and processor capacity within the failover cluster is necessary to support multiple virtual machines implemented as guest failover clusters.

For the guest operating systems, you can use only Windows Server 2012 or later. However, if you are using Windows Server 2012 in virtual machines that are using shared virtual hard disks, you must install Hyper-V integration services from Windows Server 2012 R2 or later. Both Generation 1 and Generation 2 virtual machines are supported.

When you decide to implement shared virtual hard disks as storage for guest clusters, you must first decide where to store the shared virtual hard disk. You can deploy the shared virtual hard disk at the following locations:

- Clustered Shared Volume location. In this scenario, all virtual machine files, including the shared .vhdx or .vhds files, are stored on a CSV configured as shared storage for a Hyper-V failover cluster.
- Scale-Out File Server SMB 3.0 share. This scenario uses SMB file-based storage as the location for the shared .vhdx or .vhds files. You must deploy a Scale-Out File Server and create an SMB file share as the storage location. You also need a separate Hyper-V failover cluster.



Note: You should not deploy a shared virtual hard disk on an ordinary file share or on a local hard drive on the host machine. You must deploy the shared virtual hard disk on a highly available location.

You can configure a shared virtual hard disk in a Windows Server 2016 Hyper-V cluster when you use the Failover Cluster Manager GUI or Windows PowerShell. If you use a .vhdx extra steps are required to create the guest shared virtual disk to let Hyper-V and failover cluster know that the .vhdx is a shared disk. However, with the .vhds format introduced in Windows Server 2016, those steps are not needed and the process is simplified.



Additional Reading: For more information, refer to: "Deploy a Guest Cluster Using a Shared Virtual Hard Disk" at: <http://aka.ms/isec0h>

When you use Hyper-V Manager, you can create a virtual hard disk using the VHD Set (.vhds). We recommend that you always attach virtual hard disks on a separate virtual SCSI adapter rather than the virtual disk with the operating system. However, you are able to connect to the same adapter when running a Generation 2 virtual machine.



Note: Adding virtual SCSI adapters requires the virtual machine to be offline. If the SCSI adapters are already added, you can complete all other steps while the virtual machine is online.

In Windows Server 2016, to add a shared virtual disk to two virtual machines, go to Failover Cluster Manager, select the virtual SCSI controller, and then select **Shared Drive**. Browse to the created disk, and then click **Apply**. Repeat this procedure on all virtual machines that will use this shared virtual hard disk.

To add a shared virtual hard disk by using Windows PowerShell, you should use the **Add-VMHardDiskDrive** cmdlet with the **-ShareVirtualDisk** parameter. You must run this command under administrator privileges on the Hyper-V host for each virtual machine that uses the shared .vhds file.

For example, if you want to create and add a shared virtual hard disk (**Data1.vhds**) that is stored on volume 1 of the CSV to two virtual machines named **VM1** and **VM2**, you would use the following commands in Windows PowerShell:

```
New-VHD -Path C:\ClusterStorage\Volume1\Data1.vhds -Dynamic -SizeBytes 127GB
Add-VMHardDiskDrive -VMName VM1 -Path C:\ClusterStorage\Volume1\Data1.vhds -
ShareVirtualDisk
Add-VMHardDiskDrive -VMName VM2 -Path C:\ClusterStorage\Volume1\Data1.vhds -
ShareVirtualDisk
```

In addition, if you want to add a shared virtual hard disk (**Witness.vhdx**) that is stored on an SMB file share (**\\Server1\Share1**) to a virtual machine that is named **VM2**, you should use the following command in Windows PowerShell:

```
Add-VMHardDiskDrive -VMName VM2 -Path \\Server1\Share1\Witness.vhds -ShareVirtualDisk
```

The following chart lists the different Hyper-V capabilities for each shared storage option when compared to a shared virtual disk.

Capability	Shared .vhdx and .vhds	Virtual Fibre Channel	ISCSI in virtual machine
Supported storage	Storage spaces, SAS, Fibre Channel, iSCSI, SMB	Fibre Channel SAN	iSCSI SAN
Storage is presented to the virtual machine through	Virtual SAS	Virtual Fibre Channel LUN	iSCSI LUN
Data flows through the Hyper-V switch	No	No	Yes
Storage is configured at the Hyper-V host level	Yes	Yes	No
Provides low latency and low CPU use	Yes (remote direct memory access [RDMA] or Fibre Channel)	Yes (Fibre Channel)	No

Capability	Shared .vhdx and .vhds	Virtual Fibre Channel	ISCSI in virtual machine
Requires specific hardware	No	Yes	No
Requires switch to be reconfigured when virtual machine is migrated	No	Yes	No
Exposes storage architecture	No	Yes	Yes

Question: What is the primary benefit of using shared virtual hard disks?

Implementing Scale-Out File Server for virtual machines

It is possible to use one more technique to make virtual machine storage highly available. Instead of using host or guest clustering, now you can store virtual machine files on a highly available SMB 3.0 file share. When you use this approach, you achieve storage high availability not by clustering Hyper-V nodes but by clustering file servers that host virtual machine files on their file shares. With this new capability, Hyper-V can store all virtual machine files, including configuration, files, and checkpoints, on highly available SMB file shares.

- In Windows Server 2016, you can store virtual machine files on a SMB 3.0 file share
- File servers should be running Windows Server 2012 or later
- File server cluster should be configured as a Scale-Out File Server for application data.
- Use Hyper-V Manager to create or move virtual machine files to a SMB file share

What is a Scale-Out File Server?

A Scale-Out File Server, introduced in Windows Server 2012, provides continuously available storage for file-based server applications. You configure Scale-Out File Server by creating a File Server role on a failover cluster and selecting the **Scale-Out File Server for application data** option instead of **File Server for general use**. This requires the use of a CSV volume for storage of data.

Scale-Out File Server is different from the file server clusters that were the only option in previous versions of Windows Server in several ways. An ordinary file-server cluster serves the clients only by using one node at a time; however, a Scale-Out File Server can engage all nodes simultaneously. This is achieved with the new Windows Server failover clustering features and the new capabilities in the new version of Windows file server protocol, SMB 3.0. Therefore, by adding nodes to the failover cluster running the File Server role with the Scale-Out File Server feature, performance of the entire cluster increases. As a result, it is now possible to store resources such as databases or virtual machine hard disks on the folder shares hosted on the Scale-Out File Server.

The key benefits of using a Scale-Out File Server are:

- **Active-active clustering.** When all other failover clusters work in an active-passive mode, a Scale-Out File Server cluster works in a way that all nodes can accept and serve SMB client requests. In Windows Server 2012 R2, SMB 3.0 is upgraded to SMB 3.0.2. This version improves scalability and manageability for Scale-Out File Server. SMB client connections in Windows Server 2012 R2 are tracked per file share (instead of per server), and clients are then redirected to the cluster node with the best access to the volume used by the file share.
- **Increased bandwidth.** In previous versions of Windows Server, bandwidth of the file server cluster was constrained to the bandwidth of a single cluster node. Now, because of the active-active mode in the Scale-Out File Server cluster, you can have much higher bandwidth, which you can additionally 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 system memory (RAM) as a write-through cache. The CSV Cache provides caching of read-only unbuffered I/O. This can improve performance for applications such as Hyper-V, which conducts unbuffered I/O when accessing a .vhd file. With Windows Server 2012, you can allocate up to 20 percent of the total physical RAM for CSV write-through cache, and 80 percent with Windows Server 2012 R2 and Windows Server 2016. The total physical RAM for CSV write-through cache will be consumed from nonpaged pool memory.
- **Abstraction of the storage layer.** When you use a Scale-Out File Server as the storage location for virtual disks you can live-migrate virtual machines from cluster to cluster without needing to migrate the storage, provided the URL location is accessible from the destination cluster.

To implement a Scale-Out File Server, you must meet the following requirements:

- One or more computers running Windows Server 2012 or later with the Hyper-V role installed.
- One or more computers running Windows Server 2012 or later with the File and Storage Services role installed.
- A common Active Directory infrastructure. The servers that are running AD DS do not need to run Windows Server 2016.

Before you implement virtual machines on an SMB file share, you should set up a file server cluster. To do that, you must have at least two cluster nodes with file services and failover clustering installed. In Failover Cluster Manager, you must create a file server and select the Scale-Out File Server for application data configuration. After you configure the cluster, you must deploy the SMB Share – Applications profile, which is designed for Hyper-V and other application data. After creating the share, you can use the Hyper-V Manager console to deploy new virtual machines on the SMB file share, or you can migrate existing virtual machines to the SMB file share when you use the storage migration method.

Question: Have you considered storing virtual machines on the SMB share? Why or why not?

Considerations for implementing Hyper-V clusters

By implementing failover clustering on servers with the Hyper-V feature installed, you can make virtual machines highly available. However, this also adds significant cost and complexity to a Hyper-V deployment. You must invest in additional server hardware to provide redundancy, and you should implement or have access to a shared storage infrastructure.

Use the following recommendations to ensure that the failover clustering strategy meets the organization's requirements:

- Use the following recommended failover clustering requirements:
 - 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
- Windows Server 2016 Hyper-V Live Migration considerations:
 - Verify basic requirements
 - Configure a dedicated network adapter or virtual network adapter
 - Use similar host hardware
 - Verify network configuration

- Identify the applications or services that require high availability. If you were to ask the people who use the organization's applications about their preferences, most of them would probably say that they want all applications to be highly available. However, unless you have the option of making all virtual machines highly available, you must develop priorities for which applications you will make highly available.
- Identify the application components that must be highly available to make the applications highly available. In some cases, the application might run on a single server. If so, you only need to make that server highly available. Other applications might require that several servers and other components (such as storage or the network) be highly available.
- Identify the application characteristics. You must understand several things about an application:
 - Is virtualizing the server that is running the application an option? Some applications are not supported or recommended in a virtual environment.
 - What options are available for making the application highly available? You can make some applications highly available through options other than host clustering. If other options are available, evaluate the benefits and disadvantages of each option.
 - What are the performance requirements for each application? Collect performance information on the servers currently running the applications to gain an understanding of the hardware requirements that you must meet when you virtualize the server.
 - What capacity is required to make the Hyper-V virtual machines highly available? As soon as you identify all of the applications that you must make highly available by using host clustering, you can start to design the actual Hyper-V deployment. By identifying the performance requirements and the network and storage requirements for applications, you can define the hardware that you must implement in a highly available environment.

Live Migration is one of the most important aspects of Hyper-V clustering. This module will discuss this in more detail in a later lesson. However, when implementing live migration, consider the following:

- Verify basic requirements. The basic requirements for live migration in a cluster requires that all hosts must be part of a Windows Server 2008 R2 or later failover cluster and that host processors must be from the same manufacturer. In addition, all hosts in the cluster must have access to shared storage.
- Configure a dedicated network adapter or virtual network adapter for live migration communication. When you implement failover clustering, you should configure a separate virtual LAN (VLAN) live migration network. You use this network to transfer the virtual machine memory during a failover. To optimize this configuration, configure a network adapter for this network that has a capacity of one gigabit per second (Gbps) or higher.

- Use similar host hardware. All failover cluster nodes should use the same hardware for connecting to shared storage, and all cluster nodes must have processors from the same manufacturer. Although you can enable failover for virtual machines on a host with different processor versions by configuring processor compatibility settings, the failover experience and performance is more consistent if all servers have very similar hardware.
- Verify network configuration. As with all failover clusters, the network configurations should be the same for all nodes in the failover cluster. All trunking and VLAN-tagged traffic should be the same on all failover cluster nodes. This ensures network connectivity for the guest virtual machine when taking advantage of Hyper-V virtual networking.

Maintaining and monitoring virtual machines in clusters

Failover clusters provide high availability for the roles configured in the cluster. However, you must monitor the roles and take action when there is an issue with role availability. Virtual Machine is one of the cluster roles, and when this role does not respond to a heartbeat, the failover cluster can restart or fail over the role to a different cluster node.

In Windows Server versions prior to Windows Server 2012, the failover cluster was not able to monitor applications that were running inside a virtual machine. For example, if you used a virtual machine as a print server, the failover cluster was not able to detect if the Print Spooler service in a virtual machine stopped. It would not take any action even though the print server did not work because the virtual machine was still responding to a heartbeat.

In Windows Server 2016 failover clustering, you can implement the following technologies for virtual machine maintenance and monitoring:

- Service and virtual machine health monitoring
- Network health detection (Windows Server 2012 R2 and later only)
- Virtual machine drain on shutdown (Windows Server 2012 R2 and later only)

Failover clustering in Windows Server 2016 has the ability to monitor and detect application health for applications and services that run inside a virtual machine. If a service in a virtual machine stops responding or if an event is added to the System, Application, or Security logs, the failover cluster can take actions such as restarting the virtual machine or failing it over to a different node to restore the service. The only requirement is that both the failover cluster node and the virtual machine must be running Windows Server 2012 or a later operating system, and have integration services installed.

You can configure virtual machine monitoring by using the Failover Cluster Manager or Windows PowerShell. By default, you configure a failover cluster to monitor virtual machine health. To enable heartbeat monitoring, you must install integration services on the virtual machine. You can verify the monitoring configuration on the **Settings** tab of the virtual machine resource properties. To enable monitoring of any specific services that are running on the virtual machine, you must right-click the virtual machine cluster role, click **More actions**, and then click **Configure Monitoring**. In the **Select Services** window, you can select services to monitor inside the virtual machine. The failover cluster will take action only if a service stops responding and if you have configured the service with the **Take No Actions** recovery setting in the Services Control Manager.

Windows Server 2016 also can monitor failure of virtual machine storage and loss of network connectivity with a technology called *network health detection*. Storage failure detection can detect the failure of a virtual machine boot disk or any other virtual hard disk that the virtual machine is using. If failure happens, the failover cluster moves and restarts the virtual machine on a different node.

You also can configure a virtual network adapter to connect to a protected network. If network connectivity to such a network is lost because of reasons such as physical switch failure or disconnected network cable, the failover cluster will move the virtual machine to a different node to restore network connectivity.

Windows Server 2012 R2 also enhances virtual machine availability in scenarios when one Hyper-V node shuts down before you place it in the maintenance mode and before draining any clustered roles from it. In Windows Server 2012, shutting down the cluster node before draining it results in virtual machines that are put into a saved state and then moved to other nodes and resumed. This causes an interruption to the availability of the virtual machines. In Windows Server 2016, if such a scenario occurs, the cluster automatically live-migrates all running virtual machines before the Hyper-V node shuts down.



Note: We still recommend that you drain clustered roles (and place the node in maintenance mode) before performing a shutdown operation.

Configuration of this functionality, called *virtual machine drain on shutdown*, is not accessible through Failover Cluster Manager. To configure it, you must use Windows PowerShell, and configure the **DrainOnShutdown** cluster property. It is enabled by default, and the value of this property is set to **1**. If you want to check the value, run Windows PowerShell as Administrator, and then run the following command:

```
(Get-Cluster).DrainOnShutdown
```

Question: What are some alternative Microsoft technologies that you can use for virtual machine monitoring and network monitoring?

Lesson 3

Implementing Windows Server 2016 Hyper-V virtual machine migration

Moving virtual machines from one location to another is a common procedure in the administration of Windows Server 2016 Hyper-V environments. Most of the techniques for moving virtual machines in previous versions of Windows Server required downtime. Windows Server 2016 enables virtual machine movement with no downtime. In this lesson, you will learn about virtual machine movement and migration options.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the migration options for virtual machines.
- Explain how Storage Migration works.
- Explain how Live Migration works.

Virtual machine migration options

In several scenarios you might want to migrate a virtual machine from one location to another. For example, you might want to move a virtual machine's virtual hard disk from one physical drive to another on the same host. Or you might want to move a virtual machine from one node in a cluster to another or just move a computer from one host server to another host server without the hosts being members of a cluster.

In Windows Server 2016, you can perform virtual machine migration by using the following methods:

Available options for moving virtual machines are:

- Virtual machine and storage migration
- Quick Migration
- Live Migration
- Hyper-V Replica
- Export or import of a virtual machine

- Shared nothing migration. With this method, you move a powered-on virtual machine from one location to another (or from one host to another) by using the migration wizard in the Hyper-V Manager. The virtual machine and storage do not require failover clustering or any other high availability technology with a shared nothing migration.
- Quick Migration. This feature has been available since Windows Server 2008. It requires that you have failover clustering installed and configured.
- Live Migration. This feature is an improvement over Quick Migration and has been available since Windows Server 2008 R2. You use it to migrate a virtual machine from one host to another without experiencing downtime.
- Hyper-V Replica. Use this feature both to replicate a virtual machine to another host instead of moving the virtual machine and to synchronize all virtual machine changes from the primary host to the host that holds the replica.

- Exporting and importing virtual machines. This is an established method of moving virtual machines without using a cluster. You export a virtual machine on one host, and then physically move exported files to another host by performing an import operation. This very time-consuming operation requires that a virtual machine be turned off during export and import. In Windows Server 2016, this migration method is improved. You can import a virtual machine to a Hyper-V host without exporting it before import. Windows Server 2016 Hyper-V is capable of configuring all of the necessary settings during the import operation.

Question: When will you export and import a virtual machine instead of migrating it?

How storage migration works

An administrator might want to move virtual machine files to another location in response to many different situations. For example, if the disk where a virtual machine hard disk resides runs out of space, you must move the virtual machine to another drive or volume. Moving a virtual machine to another host is a very common procedure.

In earlier versions of Windows Server, such as Windows Server 2008 R2 or Windows Server 2008, moving virtual machine storage resulted in downtime because the machine had to be turned off. If you moved a virtual machine's storage between two locations, then you also had to perform export and import operations for that specific machine. Export operations can be time-consuming, depending on the size of the virtual machine hard disks.

- Storage migration technology allows you to move a virtual machine and its storage to another location without downtime
- During migration, the virtual machine hard disk is copied from one location to another
- Changes are written to both the source and destination drives
- You can move virtual machine storage to the same host, another host, or a SMB share
- Storage and virtual machine configuration can be in different locations

In Windows Server 2012 and later, virtual machine and storage migration makes it possible for you to move a virtual machine to another location on the same host or to another host computer, without turning off the virtual machine.

To copy a virtual hard disk, an administrator starts live storage migration by using the Hyper-V console or Windows PowerShell and then either completes the **Live Migration Wizard** or specifies parameters in Windows PowerShell. Doing this creates a new virtual hard disk on the destination location, and the copy process starts. During the copy process, the virtual machine is fully functional. However, all changes that occur during copying are written to both the source and destination locations. Read operations are performed only from the source location.

As soon as the disk copy process is complete, Hyper-V switches virtual machines to run on the destination virtual hard disk. In addition, if the virtual machine is moving to another host, the computer configuration is copied, and the virtual machine is associated with another host. If a failure were to occur on the destination side, a fail-back option is always available to run on the source directory. After the virtual machine is successfully migrated and associated to a new location, the process deletes the source virtual hard disks.

The time that is required to move a virtual machine depends on the source and destination location, the speed of hard drives or storage, and the size of the virtual hard disks. The moving process accelerates if source and destination locations are on storage that supports Windows Offloaded Data Transfers (ODX).


When you move a virtual machine's virtual hard disks to another location, the **Move Wizard** presents three available options in Hyper-V Manager:

- **Move all the virtual machine's data to a single location.** You specify one single destination location, such as disk file, configuration, checkpoint, or smart paging.
- **Move the virtual machine's data to a different location.** You specify individual locations for each virtual machine item.
- **Move only the virtual machine's virtual hard disk.** You move only the virtual hard disk file.

The **Move Wizard** and these options are only available if the Hyper-V virtual machine is not part of failover cluster. All three of the options are achievable in Failover Cluster Manager using the Move Virtual Machine Storage options.

How Live Migration works

By using the Live Migration feature in Hyper-V failover clustering, you can move running virtual machines from one failover cluster node to another node in the same cluster. With Live Migration, users connected to the virtual machine should experience almost no server outages.


 **Note:** Shared-nothing live migrations are able to live-migrate without any shared storage or cluster. However, this relies on your network to carry reads and writes to both locations simultaneously and is why a clustered solution is better for production environments.

You can control Hyper-V Live Migration through **Hyper-V Settings** in Hyper-V Manager. On the **Live Migrations** tab, under **Advanced Features**, you can select the authentication protocol. The default selection is **Credential Security Support Provider (CredSSP)**. However, you also have the option of using Kerberos authentication.

CredSSP is the default configuration, and is easy to configure. However, it is less secure than Kerberos authentication. To live-migrate virtual machines CredSSP requires signing into the source server, remote desktop, or remote Windows PowerShell. Kerberos is the more secure of the two options, requires manual selection, and requires constrained delegation configuration for that host. However, this does not require sign-in to the Hyper-V host server for live migration.

You can start the live migration process by using one of the following methods:

- The Failover Cluster Management console.
- The Virtual Machine Manager Administrator console, if you use VMM to manage your physical hosts.
- A Windows Management Instrumentation (WMI) or Windows PowerShell script.

 **Note:** Use Live Migration to reduce the perceived outage of a virtual machine significantly during a planned failover. During a planned failover, you start the failover manually. Live Migration does not apply during an unplanned failover, such as when the node hosting the virtual machine fails.

The live migration process

The live migration process consists of the following steps:

1. **Migration setup.** When the administrator starts the virtual machine failover, the source node creates a TCP connection with the target physical host. This connection is used to transfer the virtual machine configuration data to the target physical host. Live Migration creates a temporary virtual machine on the target physical host and allocates memory to the destination virtual machine. The migration preparation also checks to determine whether a virtual machine can be migrated.
2. **Guest-memory transfer.** The guest memory transfers iteratively to the target host while the virtual machine is still running on the source host. Hyper-V on the source physical host monitors the pages in the working set. As the system modifies memory pages, it tracks and marks them as being modified. During this phase of the migration, the migrating virtual machine continues to run. Hyper-V iterates the memory copy process several times, and a smaller number of modified pages are copied to the destination physical computer every time. A final memory-copy process copies the remaining modified memory pages to the destination physical host. Copying stops as soon as the number of dirty pages drops below a threshold, or after 10 iterations are complete.
3. **Storage handle.** All storage attached to the virtual machine moves to the target host. This includes iSCSI and Virtual Fibre Channel attached storage and all virtual hard disks.
4. **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 begins running the virtual machine on the target host.
5. **Cleanup.** The cleanup stage finishes the migration by tearing down the virtual machine on the source host, terminating the worker threads, and signaling the completion of the migration.



Note: In Windows Server 2012 R2 and later, you can perform live migration of virtual machines by using SMB 3.0 as a transport. This means that you can utilize key SMB features, such as SMB Direct and SMB Multichannel, which provide high-speed migration with low CPU utilization.

Lesson 4

Implementing Hyper-V Replica

Hyper-V Replica is a disaster recovery feature that is built into Hyper-V. You can use it to replicate a running virtual machine to a secondary location. In Windows Server 2016, you can extend the replication to a third location. While the primary virtual machine is running, the replica virtual machine remains offline. Hyper-V Replica continues to update any changes, and, when needed, you can perform failover from a primary virtual machine to a replica virtual machine. You perform failovers manually, and they can be planned or unplanned. Planned failovers are without data loss, whereas unplanned failovers can cause loss of last changes, up to five minutes by default. In this lesson, you will learn how to implement and manage Hyper-V Replica. You also will learn how to perform both a test failover and a planned site failover.

Lesson Objectives

After completing this lesson, you will be able to:

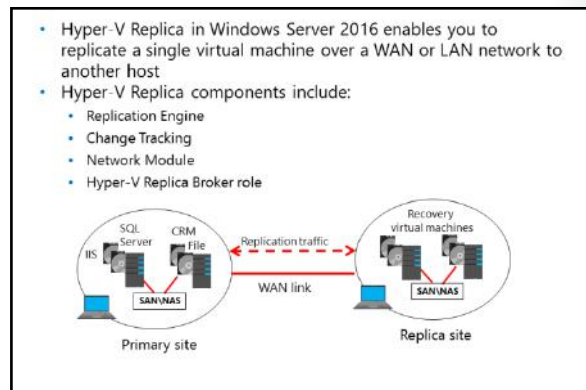
- Describe Hyper-V Replica.
- Describe the features in Hyper-V Replica in Windows Server 2016.
- Describe how to configure Hyper-V Replica.
- Explain how to perform failover with Hyper-V Replica.

What is Hyper-V Replica?

In some cases, you might want to have a spare copy of one virtual machine that you can run if the original virtual machine fails. By implementing high availability, you have one instance of a virtual machine. However, high availability does not prevent corruption of software that is running inside the virtual machine. One way to address the issue of corruption is to manually copy the virtual machine periodically. You also can back up the virtual machine and its storage. Although backing up achieves the desired result, it is resource-intensive and time-consuming. In addition, because backups are performed periodically, you never have the exact same copy as the running virtual machine.

To resolve this problem and to give administrators an up-to-date copy of a single virtual machine, Microsoft has included the *Hyper-V Replica* technology in Windows Server 2016. This technology enables virtual machines running at a primary site (or a location or host) to be replicated efficiently to a secondary site (location or host) across a WAN or LAN link. Hyper-V Replica gives you two instances of a single virtual machine residing on different hosts, one as the primary (live) copy and the other as a replica (offline) copy. These copies synchronize in real time, and a failover can be initiated whenever needed.

In the event of a failure at a primary site (caused by natural disaster, a power outage, or a server failure), an administrator can use Hyper-V Manager to execute a failover of production workloads to replica servers at a secondary location within minutes, thus incurring minimal downtime. With Hyper-V Replica, an administrator can restore virtualized workloads to a specific point in time depending on the Recovery History selections for the virtual machine.



The Hyper-V Replica technology consists of several components:

- **Replication Engine.** This component is the core of Hyper-V Replica. It manages the replication configuration details and initial replication, delta replication, failover, and test-failover operations. It also tracks virtual machine and storage mobility events and takes appropriate actions as required. For example, the replication engine pauses replication events until migration events complete and then resumes where the replication events left off.
- **Change Tracking.** This component tracks changes that are happening on the primary copy of the virtual machine. It is designed to make the scenario work regardless of where the virtual machine .vhd or .vhdx file(s) resides.
- **Network Module.** This module provides a secure and efficient way to transfer virtual machine replicas between the primary host and the replica host. Data compression is enabled by default. This communication is also secure because it relies on HTTPS and certification-based authentication.
- **Hyper-V Replica Broker role.** You configure this role in failover clustering, and it gives you Hyper-V Replica functionality even when the virtual machine that is being replicated is highly available and can move from one cluster node to another. The Hyper-V Replica Broker redirects all virtual machine-specific events to the appropriate node in the Replica cluster. The Broker queries the cluster database to determine which node should manage which events. This ensures that in the event that a quick migration, live migration, or storage migration process executes, all events redirect to the correct node in the cluster.

The site configurations do not have to use the same server or storage hardware. It is important, however, to ensure that sufficient hardware resources are available to run the replica virtual machine.

Hyper-V Replica in Windows Server 2016

In Windows Server 2012 R2 and Windows Server 2016, the Hyper-V Replica feature is improved with the following enhancements:

- **Ability to change the replication frequency.** In previous versions of Windows Server, Hyper-V Replica was set to a 5-minute replication interval, and you were not able to change this value. In Windows Server 2012 R2 and Windows Server 2016, you can set the replication interval to 30 seconds, 5 minutes, or 15 minutes. This means that you can configure your replication traffic based on your real environment. Keep in mind, however, that replicas with a higher latency (such as 15 minutes) generate more traffic when they replicate.
- **Extended replication.** In Windows Server 2012, it is possible to have only one replica of an existing virtual machine. Windows Server 2016 and Windows Server 2012 R2 provide you with the ability to replicate a single virtual machine to a third server. This means that you can replicate a running virtual machine to two independent servers. However, the replication does not happen from one server to two other servers. The server that is running an active copy of the virtual machine replicates to the replica server, and the replica server then replicates to the extended replica server. You create a second replica by running the **Extend Replication Wizard** on a passive copy. In this wizard, you can configure the same options that you configured when configuring the first replica.

Hyper-V Replica in Windows Server 2012 R2 and Windows Server 2016 is enhanced with the following features:

- **Ability to change the replication frequency:**
 - Available intervals are 30 seconds, 5 minutes, and 15 minutes
- **Extended replication:**
 - You can extend Hyper-V Replica to include a third host

Administrators can now benefit from these features because these features help to optimize the use of Hyper-V Replica and increase the availability of critical virtual machines.

Question: Do you see extended replication as a benefit for your environment?

Configuring Hyper-V Replica

Before you implement the Hyper-V Replica technology, ensure that the following prerequisites are met:

- The server hardware supports the Hyper-V role on Windows Server 2016.
- Sufficient storage exists on both the primary and replica servers to host the files that the replicated virtual machines use.
- Network connectivity exists between the locations hosting the primary and replica servers. This can be a WAN or a LAN link.
- Firewall rules are configured correctly to enable replication between the primary and replica sites (default traffic is over TCP port 80 or 443).
- An X.509v3 certificate exists to support mutual authentication with certificates, if desired.

To configure Hyper-V Replica, you should:

1. Configure authentication options
2. Configure ports
3. Select replica servers
4. Select a location for replica files
5. Enable replication on a virtual machine

You do not have to install Hyper-V Replica separately because it is not a Windows Server role or feature. Hyper-V Replica is implemented as part of the Hyper-V role. You can use it on Hyper-V servers that are standalone or on servers that are part of a failover cluster (in which case, you should configure Hyper-V Replica Broker). Unlike failover clustering, a Hyper-V role is not dependent on AD DS. You can use it with Hyper-V servers that are standalone or that are members of different Active Directory domains. (The exception to this is for servers that are part of a failover cluster.)

To enable the Hyper-V Replica technology, you first configure Hyper-V server settings. In the **Replication Configuration** group of options, you enable Hyper-V server as a replica server, and then select authentication and port options. You also should configure authorization options. You can choose to enable replication from any server that successfully authenticates (which is convenient in scenarios where all servers are part of same domain), or you can type fully qualified domain names (FQDNs) of servers that you accept as replica servers. In addition, you must configure the location for replica files. You configure these settings on each server that serves as a replica server.

After you configure options on the server level, you enable replication on a virtual machine. During this configuration process, you must specify both the replica server name and the options for connection. You can select which virtual hard disks to replicate (in cases where a virtual machine has more than one virtual hard disk), and you can also configure Recovery History and an initial replication method. In Windows Server 2016, you also can configure a replication interval.

After you have configured these options, you can start replication. After you make the initial replica, in Windows Server 2016 you also can make an extended replica to a third physical server running Hyper-V.

Failover with Hyper-V Replica

You can perform three types of failovers with Hyper-V Replica: test failover, planned failover, and failover. These three options offer different benefits, and are useful in different scenarios. This topic explains the difference between the three types of failovers.

Test failover

After you configure a Hyper-V Replica and after the virtual machines start replicating, you can perform a test failover. A test failover is a nondisruptive task that you can undertake to test a virtual machine on the replica server while the primary virtual machine is running and without interrupting the replication. You can initiate a test failover on the replicated virtual machine, which creates a new checkpoint. You can use this checkpoint to select a recovery point from which the new test virtual machine will be created. The test virtual machine has the same name as the replica but with “– Test” appended to the end of the name by default. This prevents potential conflicts with running the primary virtual machine, if the test virtual machine does not start and is disconnected.

After you finish testing, you can stop a test failover but only if test failover is running. 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 must 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. You might do this, for example, before site maintenance or before an expected disaster. Because this is a planned event, there is no data loss, but the virtual machine is 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 yet replicated to the replica server. The planned failover process then 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 is running on the replica server, and its changes do not replicate. If you want to establish replication again, you should reverse the replication. You must configure settings similar to when you turned replication on, and the existing virtual machine is 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 either unavailable or turned off. A *failover* is an unplanned event that can result in data loss because changes at the primary virtual machine might not have replicated before the disaster happened. (The replication frequency setting controls how often changes replicate.) Similar to a planned failover, during a failover, the virtual machine is running on a replica server. 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. When you reverse the replication direction, it removes the option to cancel failover.

- **Test failover:**
 - Is nondisruptive testing with zero downtime
 - Creates a new virtual machine in the recovery site from the replica checkpoint and 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
 - 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

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 replication for the virtual machine is paused.
- **View replication health.** This action provides data about the replication events for a virtual machine.
- **Extend replication.** This action is available on replicated virtual machines. It is available in Windows Server 2016, 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, it deletes all recovery points (checkpoints) for a replicated virtual machine and merges their differencing virtual hard disks.
- **Remove replication.** This action stops replication for the virtual machine.



Note: If you have implemented System Center 2012 R2 and are interested in using Hyper-V Replica for disaster recovery, you should consider using the Windows Azure Hyper-V Recovery Manager. The Hyper-V Recovery Manager helps to orchestrate the recovery of private cloud services across multiple locations in the event of an outage at the primary site.

Lab: Implementing failover clustering with Windows Server 2016 Hyper-V

Scenario

The initial deployment of virtual machines on Hyper-V has been successful for A. Datum Corporation. As a next step in virtual machine deployment, A. Datum is considering ways to ensure that the services and applications deployed on the virtual machines are highly available. As part of the implementation of high availability for most network services and applications, A. Datum is also considering options for making the virtual machines that run on Hyper-V highly available.

As one of the senior network administrators at A. Datum, you are responsible for integrating Hyper-V with failover clustering to ensure that the virtual machines deployed on Hyper-V are highly available. You are responsible for planning the virtual machine and storage configuration and for implementing the virtual machines as highly available services on the failover cluster. You also are considering other techniques, such as Hyper-V Replica, for ensuring high availability for virtual machines. You have limited hardware; so to facilitate testing before implementation in your production environment, you will enable nested virtualization to test clustering two Hyper-V Hosts and Hyper-V Replica without using physical hardware.

Objectives

After completing this lab, you will be able to:

- Configure Hyper-V Replicas.
- Configure a failover cluster for Hyper-V.
- Configure a highly available virtual machine.

Lab Setup

Estimated Time: **75 minutes**

Virtual machines: **20743A-LON-HOST2, 20743A-LON-DC1-C**

User name: **Adatum\Administrator**

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

To perform this lab, you must restart the host computer and then at the boot selection screen, select **20743A-LON-HOST2**.

Sign in to **LON-HOST2** as **Adatum\Administrator** with the password of **Pa\$\$w0rd**.

Exercise 1: The Hyper-V Failover clustering testing environment

Scenario

In order to test failover clustering for Hyper-V clusters you have decided to install the Hyper-V role on one server, implement nested virtualization on two virtual machines, create a shared VHD for both servers, and install Hyper-V on the two nested virtual machines.

The main tasks for this exercise are as follows:

1. Enabling Nested Virtualization.
2. Upgrade the configuration version and install Hyper-V on nested virtualization hosts.

► Task 1: Enabling Nested Virtualization

1. Install the **Hyper-V** role by running the following command in Windows PowerShell on **LON-HOST2**:

```
Install-WindowsFeature -Name Hyper-V,Hyper-V-Tools,Hyper-V-PowerShell -Restart
```

Your computer will restart a couple of times. Be sure to select **20743A-LON-HOST2** at the boot menu.

2. Sign into **LON-HOST2** as **Adatum\Administrator** with **Pa\$\$w0rd**.
3. Open **File Explorer** and then browse to **E:\Program Files\Microsoft Learning\20743\Drives**. (Note that the drive letter may vary based upon your host machine).
4. Run the following scripts in order: **CreateVirtualSwitches.ps1** **LON-HOST2_VM-Pre-Import-20743A.ps1**.



Note: These scripts will create the needed switches and import the virtual machines that need to be imported for this lab.

5. Start **Windows PowerShell** and then run the following commands to enable nested virtualization, replacing <virtual machine> with **20743A-LON-NVHOST3** and then **20743A-LON-NVHOST4**.

```
Set-VMProcessor -VMName <virtual machine> -ExposeVirtualizationExtensions $true -Count 2
```

```
Set-VMemory <virtual machine> -DynamicMemoryEnabled $false
```

```
Get-VMNetworkAdapter -VMName <virtual machine> | Set-VMNetworkAdapter -MacAddressSpoofing On
```

► Task 2: Upgrade the configuration version and install Hyper-V on nested virtualization hosts

1. Start **Hyper-V Manager**.
2. Upgrade the configuration version for **20743A-LON-DC1-C**, **20743A-LON-NVHOST3** and **20743A-LON-NVHOST4**.
3. Start **20743A-LON-DC1-C**, **20743A-LON-NVHOST3** and **20743A-LON-NVHOST4**.
4. By using Windows PowerShell Direct commands, install Hyper-V on both **20743A-LON-NVHOST3** and **20743A-LON-NVHOST4**. When prompted, use **Adatum\Administrator** with the password of **Pa\$\$w0rd**.

```
Enter-PSSession -VMName <Virtual Machine Name>
Install-WindowsFeature -Name Hyper-V,Hyper-V-Tools,Hyper-V-Powershell -Restart
Exit
```



Note: If this generates the following error "Command 'Exit' was not run as the session in which it was intended to run was either closed or broken" Ignore this error.

- By using the following Windows PowerShell command, configure networking on both **20743A-LON-NVHOST3** and **20743A-LON-NVHOST4**. When prompted, use **Adatum\Administrator** with the password of **Pa\$\$w0rd**. (Note that the path may differ on your host machine).

```
Invoke-Command -VMName <Virtual Machine Name> -FilePath "D:\Program Files\Microsoft Learning\20743\Drives\CreateVirtualSwitches.ps1"
```

- Connect to and then sign in to **20743A-LON-NVHOST3** as **Adatum\Administrator** with the password of **Pa\$\$w0rd**.
- Connect to and then sign in to **20743A-LON-NVHOST4** as **Adatum\Administrator** with the password of **Pa\$\$w0rd**.

Results: After completing this exercise, you should have successfully enabled nested virtualization on **20743A-LON-NVHOST3** and **20743A-LON-NVHOST4**.

Exercise 2: Configuring Hyper-V Replica

Scenario

Before you start with cluster deployment, you decide to evaluate the new Hyper-V Replica technology in Hyper-V for replicating virtual machines between hosts. You want to be able to mount a copy of a virtual machine manually onto another host if active copy (or host) fails.

The main tasks for this exercise are as follows:

- Import the **LON-CORE** virtual machine on **LON-NVHOST3**.
- Configure a replica on both host machines.
- Configure replication for the **LON-CORE** virtual machine.
- Validate a planned failover to the replica site.

► Task 1: Import the LON-CORE virtual machine on LON-NVHOST3

- On **LON-NVHOST3**, open **Server Manager**, open **Hyper-V Manager**, and then import the **20743A-LON-CORE** virtual machine.
 - Use path **C:\Program Files\Microsoft Learning\20743\Drives\20743A-LON-CORE**
 - Accept default values.



Note: The drive letter might be different based upon the number of drives on the physical host machine.

► **Task 2: Configure a replica on both host machines**

1. On **LON-NVHOST3** and **LON-NVHOST4**, open **Hyper-V Manager**, and then configure each server to be a Hyper-V Replica server.
 - Use Kerberos (HTTP) authentication.
 - Enable replication from any authenticated server.
 - Create and use folder **C:\VMReplica** as a default location to store replica files. (The drive letter might be different based upon the number of drives on the physical host machine)
2. Enable the firewall rule named **Hyper-V Replica HTTP Listener (TCP-In)** on both hosts.

► **Task 3: Configure replication for the LON-CORE virtual machine**

1. On **LON-NVHOST3**, enable replication for the **20743A-LON-CORE** virtual machine:
 - Use Kerberos authentication (HTTP).
 - Select **20743A-LON-CORE.vhd**.
 - Configure replication frequency for **30 seconds**.
 - Select to have only the latest recovery point available.
 - Start replication immediately.
2. Wait for the initial replication to finish, which could take up to ten minutes.
3. Once replication completes, ensure that the **20743A-LON-CORE** virtual machine now displays in the **Hyper-V Manager** console on **LON-NVHOST4**.

► **Task 4: Validate a planned failover to the replica site**

1. On **LON-NVHOST4**, view replication health for **20743A-LON-CORE**.
2. On **LON-NVHOST3**, perform planned failover on **LON-VNHOST4**.
3. Verify that **20743A-LON-CORE** is running on **LON-NVHOST4**.
4. On **LON-NVHOST3**, remove replication for **20743A-LON-CORE**.
5. On **LON-NVHOST4**, shut down **20743A-LON-CORE**.

Results: After completing this exercise, you should have successfully configured Hyper-V Replica.

Exercise 3: Configuring a failover cluster for Hyper-V

Scenario

A. Datum has several virtual machines that are hosting important services and, as a result, must be highly available. Because these services are not cluster-aware, A. Datum has decided to implement failover clustering on the Hyper-V host level. You plan to use iSCSI drives as storage for these virtual machines.

The main tasks for this exercise are as follows:

1. Connect to the iSCSI target from both host machines.
2. Configure failover clustering on both host machines.
3. Configure disks for a failover cluster.

► **Task 1: Connect to the iSCSI target from both host machines**

1. On **LON-NVHOST3**, in **Server Manager**, start the **iSCSI initiator**.
2. Use the **172.16.0.10** address to discover and connect to iSCSI target.
3. Switch to **LON-NVHOST4**, and use **Server Manager** to start the **iSCSI initiator**.
4. Use the **172.16.0.10** address to discover and connect to the iSCSI target.
5. In **Server Manager**, open **Disk Management**, and then initialize and bring all iSCSI drives online:
 - a. Format the first drive, and then name it **ClusterDisk**.
 - b. Format the second drive, and then name it **ClusterVMs**.
 - c. Format the third drive, and then name it **Quorum**.
6. Switch back to **LON-NVHOST3**, open **Disk Management**, and then bring all three iSCSI drives online.



Note: Disk numbers might vary based on the number of physical disks in the host computer. Choose the disks that are 20 GB in size.

► **Task 2: Configure failover clustering on both host machines**

1. On **LON-NVHOST3** and **LON-NVHOST4**, install the failover clustering feature.
2. On **LON-NVHOST3**, create a failover cluster using the following settings:
 - Add **LON-NVHOST3** and **LON-NVHOST4**.
 - Name the cluster **VMCluster**.
 - Assign the address **172.16.0.126**.
 - Clear the **Add all eligible storage to the cluster** option.

► **Task 3: Configure disks for a failover cluster**

1. In **Failover Cluster Manager**, on **LON-NVHOST3**, add all three iSCSI disks to the cluster.
2. Verify that all three iSCSI disks display as available for cluster storage.
3. Add the Cluster Disk 1 to **Cluster Shared Volumes**.
4. From the **VMCluster.adatum.com** node, click **More Actions**, and then configure the **Cluster Quorum Settings** to use the default quorum configuration.

Results: After completing this exercise, you should have successfully configured the failover clustering infrastructure for Hyper-V.

Exercise 4: Configuring a highly available virtual machine

Scenario

After you have configured the Hyper-V failover cluster, you want to add virtual machines as highly available resources. In addition, you want to evaluate live migration and test storage migration.

The main tasks for this exercise are as follows:

1. Move virtual machine storage to the iSCSI target.
2. Configure the virtual machine as highly available.
3. Perform a Live Migration for the virtual machine.
4. Perform a Storage Migration for the virtual machine.
5. Prepare for the end of the course.

► Task 1: Move virtual machine storage to the iSCSI target

1. Ensure that **LON-NVHOST3** is the owner of the disk that is assigned to Cluster Shared Volume. If it is not, move the disk to **LON-NVHOST3**.
2. On **LON-NVHOST3**, open **File Explorer**, and then browse to **C:\Program Files\Microsoft Learning\20743\Drives\20743A-LON-CORE\Virtual Hard Disks**.



Note: The drive letter might be different depending on the physical machine.

3. Move the **20743A-LON-CORE.vhd** virtual hard disk file to the **C:\ClusterStorage\Volume1** location.

► Task 2: Configure the virtual machine as highly available

1. In **Failover Cluster Manager**, click the **Roles** node, and then start the **New Virtual Machine Wizard**. If an error displays informing you that Microsoft Management has stopped working, restart this step.
2. In the **New Virtual Machine Wizard**, use the following settings:
 - Select **LON-NVHOST3** as the cluster node.
 - Name the computer as **TestClusterVM**.
 - Store the file in **C:\ClusterStorage\Volume1**.
 - Select **Generation 1**.
 - Assign **512** megabytes (MB) of **RAM** to **TestClusterVM**.
 - Leave the Network as **Not Connected**.
 - Connect the machine to the existing virtual hard disk **20743A-LON-CORE.vhd**, located at **C:\ClusterStorage\Volume1**.
 - Open **Settings** for **TestClusterVM**.
3. Enable the option for migration to computers with a different processor version.
4. On the **Roles** node, start the virtual machine.

► **Task 3: Perform a Live Migration for the virtual machine**

1. On **LON-NVHOST4**, in **Failover Cluster Manager**, start **Live Migration**.
2. Move **TestClusterVM** from **LON-NVHOST3** to **LON-NVHOST4**.
3. Connect to **TestClusterVM**, and then ensure that you can operate it.

► **Task 4: Perform a Storage Migration for the virtual machine**

1. On **LON-NVHOST4**, open **Hyper-V Manager**, and then start **20743A-LON-CORE**.
2. Perform a move operation on **20743A-LON-CORE** by moving the virtual machine from its current location to **C:\LON-CORE**.
3. Verify whether the virtual machine is operational during the move process.
4. When the move completes, shut down all running virtual machines.

Results: After completing this exercise, you should have successfully configured the virtual machine as highly available.

► **Task 5: Prepare for the end of the course**

- Shut down the host computer.

Question: How can you extend Hyper-V Replica in Windows Server 2016?

Question: What is the difference between live migration and storage migration?

Module Review and Takeaways

Review Question

Question: Do you have to implement CSV to provide high availability for virtual machines in VMM in Windows Server 2016?

Tools

Tools for implementing failover clustering with Hyper-V include:

- Failover Cluster Manager
- Hyper-V Manager
- VMM console

Best Practices

- Develop standard configurations before you implement highly available virtual machines. You should configure the host computers to be as close to identical as possible. To ensure that you have a consistent Hyper-V platform, you should configure standard network names, and use consistent naming standards for CSV volumes.
- Use new features in Hyper-V Replica to extend your replication to more than one server.
- Consider using Scale-Out File Server clusters as storage for highly available virtual machines.
- Implement Virtual Machine Manager. Virtual Machine Manager provides a management layer on top of Hyper-V and Failover Cluster Manager that can stop you from making mistakes when you manage highly available virtual machines. For example, it stops 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 implementing CSV and migrating the shared storage to CSV.	
A virtual machine fails over to another node in the host cluster, but loses all network connectivity.	
Four hours after restarting a Hyper-V host that is a member of a host cluster, there are still no virtual machines running on the host.	

Course Evaluation

Your evaluation of this course will help Microsoft understand the quality of your learning experience.

Please work with your training provider to access the course evaluation form.

Microsoft will keep your answers to this survey private and confidential and will use your responses to improve your future learning experience. Your open and honest feedback is valuable and appreciated.

MCT USE ONLY. STUDENT USE PROHIBITED

Module 1: Installing and configuring Windows Server 2016

Lab: Installing and configuring Nano Server

Exercise 1: Installing Nano Server

► Task 1: Copy the required Windows PowerShell scripts

1. Switch to **LON-DC1**.
2. Right-click **Start**, and then click **Windows PowerShell (Admin)**.
3. In the **Windows PowerShell** window, type **cd**, and then press Enter.
4. In the **Windows PowerShell** window, type **md Nano**, and then press Enter.
5. In the **Windows PowerShell** window, type the following command, and then press Enter:

```
copy d:\NanoServer\NanoServerImageGenerator\*.ps* c:\nano
```

► Task 2: Import Windows PowerShell modules

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

```
Import-Module c:\nano\NanoServerImageGenerator.psm1
```

► Task 3: Create a virtual hard disk

1. In the **Windows PowerShell** window, type the following command, and then press Enter:

```
new-NanoServerImage -Edition Standard -mediapath D:\ -Basepath c:\nano -targetpath  
c:\nano\nano-svr1.vhdx -DeploymentType Guest -computername NANO-SVR1 -storage -  
packages Microsoft-NanoServer-IIS-Package
```

2. At the **AdministratorPassword** prompt, type **Pa\$\$w0rd**, and then press Enter.
3. When the process is completed, on the taskbar, click **File Explorer**, navigate to **C:\Nano**, and then examine the files listed. Verify that **nano-svr1.vhdx** exists.



Note: Normally, you would now create a virtual machine to use the nano-svr1.vhdx file. However, to expedite the process, you will start a virtual machine that has already been created.

► Task 4: Sign in to the NANO-SVR1 virtual machine

1. On **NANO-SVR1**, in the **User name** box, type **Administrator**, and then press the Tab key.
2. In the **Password** box, type **Pa\$\$w0rd**, and then press Enter.

Results: After completing this exercise, you will have successfully created the required virtual hard disk for Nano Server.

Exercise 2: Completing post-installation tasks on Nano Server

► Task 1: Use the Nano Server Recovery Console to view basic settings

1. On **NANO-SVR1**, in the **Nano Server Recovery Console**, observe that the computer name is **Nano-Svr1** and that the computer is in a workgroup. Press the Tab key until **Networking** is selected, and then press Enter.
2. Press Enter on the **Ethernet** adapter. In **Network Adapter Settings**, notice that DHCP is obtaining the IP configuration.
3. Make a note of the IP address.
4. Press Esc twice.

► Task 2: Add Nano Server to the domain

1. Switch to **LON-DC1**.
2. Switch to the **Administrator: Windows PowerShell** window.
3. At the command prompt, type the following cmdlet, and then press Enter:

```
djoin.exe /provision /domain adatum /machine nano-svr1 /savefile .\odjblob
```



Note: Replace the IP address 172.16.0.X in the following commands with the IP address that you recorded earlier from your Nano Server installation.

4. At the command prompt, type the following cmdlet, and then press Enter. Your IP address might be different.

```
Set-Item WSMAN:\localhost\Client\TrustedHosts "172.16.0.X"
```

5. Type **Y**, and when prompted, press Enter.
6. At the command prompt, type the following cmdlet, and then press Enter. Your IP address might be different.

```
$ip = "172.16.0.X"
```

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

```
Enter-PSSession -ComputerName $ip -Credential $ip\Administrator
```

8. In the **Windows PowerShell credential request** dialog box, in the **Password** box, type **Pa\$\$w0rd**, and then click **OK**.
9. At the command prompt, type the following cmdlet, and then press Enter:

```
netsh advfirewall firewall set rule group="File and Printer Sharing" new enable=yes
```

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

```
Exit-PSSession
```

11. At the command prompt, type the following command, and then press Enter. Your IP address might be different.

```
net use z: \\172.16.0.X\c$
```

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

```
copy c:odjblob
```

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

```
Enter-PSSession -ComputerName $ip -Credential $ip\Administrator
```

15. In the **Windows PowerShell credential request** dialog box, in the **Password** box, type **Pa\$\$w0rd**, and then click **OK**.
16. At the command prompt, type **cd**, and then press Enter.
17. At the command prompt, type the following cmdlet, and then press Enter:

```
djoin /requestodj /loadfile c:\odjblob /windowspath c:\windows /localos
```

18. At the command prompt, type the following cmdlet, and then press Enter. Nano Server restarts.

```
shutdown /r /t 5
```

19. Switch to **NANO-SVR1**.
20. In the **User name** box, type **Administrator**, and then press the Tab key.
21. In the **Password** box, type **Pa\$\$w0rd** and then press Tab.
22. In the **Domain** box, type **Adatum**, and then press Enter.
23. In the **Nano Server Recovery Console**, observe that the computer is in the adatum.com domain.

► Task 3: Use Windows PowerShell to configure the Nano Server settings

1. Switch to **LON-DC1**, and then close Windows PowerShell.
2. Right-click **Start**, and then click **Windows PowerShell (Admin)**.
3. At the command prompt, type the following cmdlet, and then press Enter:

```
get-windowsfeature -comp Nano-svr1
```

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

```
install-windowsfeature Fs-fileserver -comp Nano-svr1
```



Note: If you see error "Warning: Failed to start automatic updating for installed components. Error: 0x80040154" you can ignore this.

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

```
get-windowsfeature -comp Nano-svr1
```

6. At the command prompt, type the following cmdlet, and then press Enter. Substitute the X for the last octet of the IP address on the Nano server.

```
$ip = "172.16.0.X"
```

- At the command prompt, type the following cmdlet, and then press Enter:

```
Enter-PSSession -ComputerName $ip -Credential $ip\Administrator
```

- In the **Windows PowerShell credential request** dialog box, in the **Password** box, type **Pa\$\$w0rd**, and then click **OK**.

- At the command prompt, type the following cmdlet, and then press Enter:

```
get-netipaddress
```

- At the command prompt, type the following cmdlet, and then press Enter:

```
bcdedit /enum
```

- At the command prompt, type the following cmdlet, and then press Enter:

```
net share
```

- At the command prompt, type the following cmdlet, and then press Enter:

```
Exit-PSSession
```

► **Task 4: Enable remote management with Server Manager**

- On **LON-DC1**, in **Server Manager**, in the navigation pane, right-click **All Servers**, and then click **Add Servers**.
- In the **Add Servers** dialog box, in the **Name (CN):** box, type **Nano-SVR1**, and then click **Find Now**.
- In the **Name** list, click **Nano-svr1**, and then to add the computer to the **Computer** list, click the right arrow key.
- Click **OK**.
- In **Server Manager**, expand **File and Storage Services**.
- Click **Shares**, and then in the **TASKS** list, click **New Share**.
- In the **New Share Wizard**, click **SMB Share – Quick**, and then click **Next**.
- On the **Select the server and path for this share** page, in the **Server** list, click **nano-svr1**, and then click **Next**.
- On the **Specify share name** page, in the **Share name** box, type **Data**, and then click **Next**.
- To complete this installation, click **Next** twice, and then click **Create**.
- Click **Close**.

► **Task 5: Test the file server and web server on Nano Server**

- On **LON-DC1**, switch to the **Administrator: Windows PowerShell** window.
- At the command prompt, type the following command, and then press Enter:

```
net use z: /d
```

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

```
net use z: \\Nano-svr1\c$
```

- Click **Start**, type **Notepad**, and then press Enter.

5. In Notepad, type **<H1> Nano Server Website </H1>**.
6. Click **File** and then click **Save As**.
7. In the **Save As** dialog box, in the **File name** box, type **z:\inetpub\wwwroot**, and then press Enter.
8. In the **Save as type** list, click **All Files**.
9. In the **File name** box, type **Default.htm**, and then click **Save**.
10. Close Notepad.
11. Click **Start**, click **All apps**, click **Windows Accessories**, and then click **Internet Explorer**.
12. Navigate to **http://nano-svr1**. Does your webpage display?
13. Close Windows Internet Explorer.
14. On **LON-DC1**, at the command prompt, type the following command, and then press Enter:

```
net use y: \\nano-svr1\data
```

15. Type **cmd**, and then press Enter.
16. Type **write**, and then press Enter.
17. In WordPad, type **This is my document**, click **File**, and then click **Save**.
18. In the **Save As** dialog box, in the **File name** box, type **Y:**, and then press Enter.
19. In the **File name** box, type **My document**, and then click **Save**.
20. In File Explorer, navigate to **data (\\nano-svr1) (Y:)**. Is your file listed?

► Task 6: Prepare for the next module

When you have finished the lab, revert the virtual machines to their initial state.

1. On the host computer, start **Hyper-V Manager**.
2. In the **Virtual Machines** list, right-click **20743A-NANO-SVR1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20743A-LON-DC1**.

Results: After completing this exercise, you will have successfully configured the domain and network settings of Nano Server and installed an additional role.

MCT USE ONLY. STUDENT USE PROHIBITED

Module 2: Overview of storage in Windows Server 2016

Lab A: Implementing and managing storage

Exercise 1: Implementing FSRM

► Task 1: Create a quota template

1. On **LON-DC1**, in **Server Manager**, click **Add roles and features**.
2. In the **Add Roles and Features Wizard**, click **Next**.
3. Confirm that **Role-based or feature-based installation** is selected, and then click **Next**.
4. Confirm that **LON-DC1.Adatum.com** is selected, and then click **Next**.
5. On the **Select server roles** page, expand **File and Storage Services (2 of 12 Installed)**, expand **File and iSCSI Services (1 of 11 Installed)**, and then select the **File Server Resource Manager** check box.
6. In the **Add role and features Wizard** dialog box, click **Add Features**.
7. Click **Next** twice to confirm the role service and feature selection.
8. On the **Confirm installation selections** page, click **Install**.
9. When the installation completes, click **Close**.
10. In **Server Manager**, click **Tools**, and then click **File Server Resource Manager**.
11. In the **File Server Resource Manager** console, expand **Quota Management**, and then click **Quota Templates**.
12. Right-click **Quota Templates**, and then click **Create Quota Template**.
13. In the **Create Quota Template** dialog box, in the **Template name** text box, type **500 MB Limit - Mail to user and log to Event Viewer**.
14. In the **Space Limit** section, type **500**, and ensure that the unit is **MB** and the **Hard quota. Do not allow users to exceed limit** option is selected.
15. In the **Notification thresholds** section, click **Add**.
16. In the **Add Threshold** dialog box, in the **Generate notifications when usage reaches (%)** text box, type **75**.
17. On the **E-mail Message** tab, select the **Send e-mail to the user who exceeded the threshold** check box.
18. Click the **Event Log** tab. In the **File Server Resource Manager** dialog box, click **Yes**.
19. On the **Event Log** tab, select the **Send warning to event log** check box, and then click **OK**.
20. In the **File Server Resource Manager** dialog box, click **Yes**.
21. In the **Create Quota Template** window, click **OK**.

► Task 2: Configure a quota based on the quota template

1. In the **File Server Resource Manager** console, click **Quotas**.
2. Right-click **Quotas**, and then click **Create Quota**.
3. In the **Create Quota** window, in the **Quota path** box, type **E:\Labfiles\Mod02\Home**.
4. Select the **Auto apply template and create quotas on existing and new subfolders** option.

5. In the **Derive properties from this quota template (Recommended)** drop-down list, click **500 MB Limit - Mail to user and log to Event Viewer**.

6. Click **Create**.

► **Task 3: Test that the quota is functional**

1. Switch to **LON-CL1**, and then sign in as **Adatum\Abbi** with the password **Pa\$\$w0rd**.
2. On **LON-CL1**, on the taskbar, click the **File Explorer** icon.
3. In the **File Explorer** window, in the tree pane, click **This PC**.
4. Notice that the drive P only has 500 MB available space.

► **Task 4: Create a file screen**

1. Switch to **LON-DC1**.
2. In the **File Server Resource Manager** console tree, expand **File Screening Management**, and then click **File Screens**.
3. Right-click **File Screens**, and then click **Create File Screen**.
4. In the **Create File Screen** window, in the **File screen path** text box, type **E:\Labfiles\Mod02\Home**.
5. In the **Create File Screen** window, click the **Derive properties from this file screen template (recommended)** drop-down list, and then click **Block Audio and Video Files**.
6. Click **Create**.

► **Task 5: Create a file group**

1. On **LON-DC1**, right-click **File Server Resource Manager (Local)**, and then click **Configure Options**.
2. In the **File Server Resource Manager Options** dialog box, click the **File Screen Audit** tab.
3. On the **File Screen Audit** tab, select the **Record file screening activity in auditing database** check box, and then click **OK**.
4. In the **File Server Resource Manager** console tree, click **File Groups**.
5. Right-click **File Groups**, and then click **Create File Group**.
6. In the **Create File Group Properties** window, in the **File group name** box, type **Adatum Media Files Group**.
7. In the **Files to include** box, type ***.mp***, and then click **Add**.
8. In the **Files to exclude** box, type ***.mpp**, click **Add**, and then click **OK**.



Note: For the purposes of this course, only a sample of digital-media extensions is being added here. To limit all digital media files, you need to add more extensions to the file group.

9. In the **File Server Resource Manager** console tree, click **File Screen Templates**.
10. In the content pane, right-click the **Block Audio and Video Files** template, and then click **Edit Template Properties**.
11. On the **Settings** tab, under **File groups**, click to clear the **Audio and Video Files** check box.
12. Select the **Adatum Media Files Group** check box.
13. Click **OK**. At the message prompt, click **Yes**, and then click **OK**.

► Task 6: Test the file screen

1. Switch to **LON-CL1**. In the **File Explorer** window, in the right pane, double-click **Abbi (\\LON-DC1\Home) (P:)**.
2. In the right pane, right-click and point to **New**, and then click **Text Document**.
3. On the ribbon, click the **View** tab and then select the **File name extensions** check box.
4. Right-click **New Text Document.txt**, and then click **Rename**.
5. Select the whole file name, type **musicfile.mp3**, and then press Enter.
6. Click **Yes** to change the file-name extension.
7. Notice the **File Access Denied** dialog box that opens.
8. Click **Cancel**.

► Task 7: Generate an on-demand storage report

1. Switch to **LON-DC1**. In the **File Server Resource Manager** console, click **Storage Reports Management**.
2. Right-click **Storage Reports Management**, and then click **Generate Reports Now**.
3. In the **Report data** section, under **Select reports to generate**, select the **File Screening Audit** check box.
4. Click the **Scope** tab, and then click **Add**.
5. In the **Browse for Folder** dialog box, browse to **E:\Labfiles\Mod02\Home**, and then click **OK**.
6. Click **OK** to close the **Storage Reports Task Properties** dialog box.
7. In the **Generate Storage Reports** dialog box, verify that **Wait for reports to be generated and then display them** is selected, and then click **OK**.
8. Double-click the HTML document that starts with **FileScreenAudit**, and then review the generated html reports.
9. Close all open windows on **LON-DC1**.

Results: After completing this exercise, you should have successfully configured FSRM quotas and file screening, and generated a storage report.

Exercise 2: Implementing Data Deduplication

► Task 1: Install the Data Deduplication role service

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, click **Next**.
5. On the **Select server roles** page, in the **Roles** list, expand **File and Storage Services (3 of 12 installed)**.
6. Expand **File and iSCSI Services (2 of 11 installed)**.

7. Select the **Data Deduplication** check box, and then click **Next**.
8. On the **Select features** page, click **Next**.
9. On the **Confirm installation selections** page, click **Install**.
10. When installation is complete, on the **Installation progress** page, click **Close**.

► Task 2: Enable and configure Data Deduplication

1. On **LON-DC1**, right-click **Start** and then click **Run**.
2. In the **Run** dialog box, in the **Open** text box, type **E:\Labfiles\Mod02\CreateLabFiles.cmd**, and then click **OK**.
3. On the taskbar, click the **File Explorer** icon.
4. If necessary, click **This PC**.
5. Notice that **Allfiles (E:)** has less than **50%** free space.
6. Switch to the **Server Manager** window.
7. In **Server Manager**, in the navigation pane, click **File and Storage Services**, and then click **Disks**.
8. In the **Disks** pane, click **1**.
9. Beneath **VOLUMES**, click **E**.
10. Right-click **E**, and then click **Configure Data Deduplication**.
11. In the **Allfiles (E:) Deduplication Settings** dialog box, in the **Data deduplication** list, click **General purpose file server**.
12. In the **Deduplicate files older than (in days)** text box, type **0**.
13. Click **Set Deduplication Schedule**.
14. In the **LON-DC1 Deduplication Schedule** dialog box, select the **Enable throughput optimization** check box, and then click **OK**.
15. In the **Allfiles (E:) Deduplication Settings** dialog box, click **OK**.

► Task 3: Test Data Deduplication

1. On **LON-DC1**, click **Start** and then click **Windows PowerShell**.
2. In the Windows PowerShell window, type the following command, and then press Enter:

```
Start-DedupJob E: -Type Optimization -Memory 50
```

3. In the Windows PowerShell window, type the following command, and then press Enter:

```
Get-DedupJob -Volume E:
```

4. Switch to the **File Explorer** window.
5. In **File Explorer**, expand **Allfiles (E:)**, expand **Labfiles**, expand **Mod02**, click **Data**, right-click **report.docx**, and then select **Properties**.
6. In the **Properties** window, observe the values of **Size** and **Size on disk** and note any differences.
7. Wait for five minutes to allow the deduplication job to run.
8. Switch to the Windows PowerShell prompt.

9. In the Windows PowerShell command prompt window, type the following command, and then press Enter:

```
Get-DedupStatus -Volume E: | fl
```

10. In the Windows PowerShell command prompt window, type the following command, and then press Enter:

```
Get-DedupVolume -Volume E: | fl
```

11. In the Windows PowerShell command prompt window, type the following command, and then press Enter:

```
Get-DedupMetaData -Volume E: | fl
```

12. Switch to **Server Manager**.
13. In **Server Manager**, in the navigation pane, click **File and Storage Services**, and then click **Disks**.
14. In the **DISKS** pane, click **1**.
15. Beneath **VOLUMES**, click **E**.
16. Observe the values for **Deduplication Rate** and **Deduplication Savings**.
17. Close all open windows except **Server Manager**.

Results: After completing this exercise, you should have successfully installed and configured the Data Deduplication role service for the appropriate data volume on LON-DC1.

Exercise 3: Configuring iSCSI storage

► Task 1: Install the iSCSI target role service

1. Switch to **LON-SVR2**. On the taskbar, click **Start**, and then click **Server Manager**.
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, click **Next**.
6. On the **Select server roles** page, expand **File And Storage Services (2 of 12 installed)**, expand **File and iSCSI Services (1 of 11 installed)**, select the **iSCSI Target Server** check box, and then click **Next**.
7. On the **Select features** page, click **Next**.
8. On the **Confirm installation selections** page, click **Install**.
9. When installation is complete, click **Close**.

► Task 2: Configure the iSCSI targets

1. On **LON-SVR2**, in **Server Manager**, in the navigation pane, click **File and Storage Services**.
2. In the **File and Storage Services** pane, click **iSCSI**.
3. In the **iSCSI VIRTUAL DISKS** pane, click **TASKS**, and then in the **TASKS** drop-down list, click **New iSCSI Virtual Disk**.
4. In the **New iSCSI Virtual Disk Wizard**, on the **Select iSCSI virtual disk location** page, under **Storage location**, click **C:**, and then click **Next**.
5. On the **Specify iSCSI virtual disk name** page, in the **Name** box, type **iSCSIDisk1**, and then click **Next**.
6. On the **Specify iSCSI virtual disk size** page, in the **Size** box, type **5**, ensure **GB** is selected in the drop-down list, and then click **Next**.
7. On the **Assign iSCSI target** page, click **New iSCSI target**, and then click **Next**.
8. On the **Specify target name** page, in the **Name** box, type **lon-dc1**, and then click **Next**.
9. On the **Specify access servers** page, click **Add**.
10. In the **Select a method to identify the initiator** dialog box, click **Enter a value for the selected type**, and in the **Type** drop-down list, click **IP Address**, in the **Value** box, type **172.16.0.10**, and then click **OK**.
11. On the **Specify access servers** page, click **Add**.
12. In the **Select a method to identify the initiator** dialog box, click **Enter a value for the selected type**, and in the **Type** drop-down list, select **IP Address**, and then in the **Value** box, type **10.10.0.10**, and then click **OK**.
13. On the **Specify access servers** page, click **Next**.
14. On the **Enable Authentication** page, click **Next**.
15. On the **Confirm selections** page, click **Create**.
16. On the **View results** page, wait until the creation is completed, and then click **Close**.
17. In the **iSCSI VIRTUAL DISKS** pane, click **TASKS**, and then in the **TASKS** drop-down list, select **New iSCSI Virtual Disk**.
18. In the **New iSCSI Virtual Disk Wizard**, on the **Select iSCSI virtual disk location** page, under **Storage location**, click **C:**, and then click **Next**.
19. On the **Specify iSCSI virtual disk name** page, in the **Name** box, type **iSCSIDisk2**, and then click **Next**.
20. On the **Specify iSCSI virtual disk size** page, in the **Size** box, type **5**, ensure **GB** is selected in the drop-down list, and then click **Next**.
21. On the **Assign iSCSI target** page, click **lon-dc1**, and then click **Next**.
22. On the **Confirm selections** page, click **Create**.
23. On the **View results** page, wait until the creation is completed, and then click **Close**.
24. In the **iSCSI VIRTUAL DISKS** pane, click **TASKS**, and then in the **TASKS** drop-down list, select **New iSCSI Virtual Disk**.
25. In the **New iSCSI Virtual Disk Wizard**, on the **Select iSCSI virtual disk location** page, under **Storage location**, click **C:**, and then click **Next**.

26. On the **Specify iSCSI virtual disk name** page, in the **Name** box, type **iSCSIDisk3**, and then click **Next**.
27. On the **Specify iSCSI virtual disk size** page, in the **Size** box, type **5**, ensure **GB** is selected in the drop-down list, and then click **Next**.
28. On the **Assign iSCSI target** page, click **lon-dc1**, and then click **Next**.
29. On the **Confirm selections** page, click **Create**.
30. On the **View results** page, wait until the creation is completed, and then click **Close**.
31. In the **iSCSI VIRTUAL DISKS** pane, click **TASKS**, and then in the **TASKS** drop-down list, select **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** box, type **iSCSIDisk4**, and then click **Next**.
34. On the **Specify iSCSI virtual disk size** page, in the **Size** box, type **5**, ensure **GB** is selected in the drop-down list, and then click **Next**.
35. On the **Assign iSCSI target** page, click **lon-dc1**, and then click **Next**.
36. On the **Confirm selections** page, click **Create**.
37. On the **View results** page, wait until the creation is completed, and then click **Close**.
38. In the **iSCSI VIRTUAL DISKS** pane, click **TASKS**, and then in the **TASKS** drop-down list, click **New iSCSI Virtual Disk**.
39. In the **New iSCSI Virtual Disk Wizard**, on the **Select iSCSI virtual disk location** page, under **Storage location**, click **C:**, and then click **Next**.
40. On the **Specify iSCSI virtual disk name** page, in the **Name** box, type **iSCSIDisk5**, and then click **Next**.
41. On the **Specify iSCSI virtual disk size** page, in the **Size** box, type **5**, ensure **GB** is selected in the drop-down list, and then click **Next**.
42. On the **Assign iSCSI target** page, click **lon-dc1**, and then click **Next**.
43. On the **Confirm selections** page, click **Create**.
44. On the **View results** page, wait until the creation is completed, and then click **Close**.

► **Task 3: Configure Multipath I/O (MPIO)**

1. Switch to **LON-DC1**.
2. Open **Network Connections** and enable **Ethernet 2**.
3. Switch to **LON-SVR2** and open **Network Connections**. Enable **Ethernet 2**.
4. Switch to **LON-DC1**.
5. On the taskbar, click **Start**, and then click **Server Manager**.
6. In **Server Manager**, click **Add roles and features**.
7. In the **Add Roles and Features Wizard**, on the **Before you begin** page, click **Next**.
8. On the **Select installation type** page, click **Next**.

9. On the **Select destination server** page, ensure that **Select server from the server pool** is selected, and then click **Next**.
10. On the **Select server roles** page, click **Next**.
11. On the **Select features** page, click **Multipath I/O**, and then click **Next**.
12. On the **Confirm installation selections** page, click **Install**.
13. When installation is complete, click **Close**.
14. Click **Start**, click **Power**, and then click **Restart**.
15. Click **Continue**.
16. After the computer restarts, sign in as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
17. In **Server Manager**, on the menu bar, click **Tools**, and then click **iSCSI Initiator**.
18. In the **Microsoft iSCSI** dialog box, click **Yes**.
19. In the **iSCSI Initiator Properties** dialog box, on the **Targets** tab, in the **Target** box, type **LON-SVR2**, click **Quick Connect**, and then in the **Quick Connect** box, click **Done**.
20. To close the **iSCSI Initiator Properties** dialog box, click **OK**.
21. In **Server Manager**, on the menu bar, click **Tools**, and then click **MPIO**.
22. In **MPIO Properties** dialog box, click the **Discover Multi-Paths** tab.
23. Select the **Add support for iSCSI devices** check box, and then click **Add**. When you are prompted to restart the computer, click **Yes**.
24. After the computer restarts, sign in as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
25. In **Server Manager**, on the menu bar, click **Tools**, and then click **MPIO**.
26. In the **MPIO Properties** dialog box, on the **MPIO Devices** tab, notice that **Device Hardware ID** **MSFT2005iSCSIBusType_0x9** appears on the list.
27. To close the **MPIO Properties** dialog box, click **OK**.

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

1. On **LON-DC1**, in **Server Manager**, click **Tools**, and then click **iSCSI Initiator**.
2. In the **iSCSI Initiator Properties** dialog box, on the **Targets** tab, click **Disconnect**.
3. In the **Disconnect From All Sessions** dialog box, click **Yes**.
4. In the **iSCSI Initiator Properties** dialog box, on the **Targets** tab, click **Connect**.
5. In the **Connect to Target** window, select the **Enable multi-path** check box, verify that the **Add this connection to the list of Favorite Targets** check box is selected, and then click **Advanced**.
6. In the **Advanced Settings** dialog box, on the **General** tab, change the **Local Adapter** from **Default** to **Microsoft iSCSI Initiator**. In the **Initiator IP** drop-down list, click **172.16.0.10**, and then in the **Target Portal IP** drop-down list, click **172.16.0.12 / 3260**.
7. In the **Advanced Settings** dialog box, click **OK**.
8. In the **Connect to Target** window, click **OK**.
9. In the **iSCSI Initiator Properties** dialog box, on the **Targets** tab, click **Connect**.
10. In the **Connect to Target** window, select the **Enable multi-path** check box, verify that the **Add this connection to the list of Favorite Targets** check box is selected, and then click **Advanced**.

11. In the **Advanced Settings** dialog box, on the **General** tab, change the **Local Adapter** from **Default** to **Microsoft iSCSI Initiator**. In the **Initiator IP** drop-down list, select **10.10.0.10**, and then in the **Target Portal IP** drop-down list, select **10.10.0.12 / 3260**.
 12. In the **Advanced Settings** dialog box, click **OK**.
 13. In the **Connect to Target** window, click **OK**.
 14. In the **iSCSI Initiator Properties** dialog box, click the **Volumes and Devices** tab.
 15. In the **iSCSI Initiator Properties** dialog box, on the **Volumes and Devices** tab, click **Auto Configure**.
 16. In the **iSCSI Initiator Properties** dialog box, click the **Targets** tab.
 17. In the **Targets** list, select **iqn.1991-05.com.microsoft:lon-svr2-lon-dc1-target**, and then click **Devices**.
 18. In the **Devices** dialog box, click **MPIO**.
 19. Verify that in the **Load balance policy**, **Round Robin** is selected. Under **This device has the following paths**, notice that two paths are listed. Select the first path, and then click **Details**.
 20. Note the IP address of the Source and Target portals, and then click **OK**.
 21. Select the second path, and then click **Details**.
 22. Verify that the Source IP address is the address for the second network adapter, and then click **OK**.
 23. To close the **Device Details** dialog box, click **OK**.
 24. To close the **Devices** dialog box, click **OK**.
 25. Close the **iSCSI Initiator Properties** dialog box.
- **Task 5: Verify the presence of iSCSI disks**
1. In **Server Manager**, click **File and Storage Services**.
 2. Click **Disks**.
 3. Notice that the five disks of iSCSI bus type are present in the list.

Results: After completing this exercise, you should have successfully installed iSCSI Target Server, configured MPIO, and connected to the iSCSI target by using iSCSI initiators.

► **Task 6: Prepare for the next lab**

- Leave the virtual machines running for the next lab in this module.

Lab B: Implementing and managing advanced storage solutions

Exercise 1: Configuring redundant storage spaces

► Task 1: Create a storage pool by using the iSCSI disks attached to the server

1. On **LON-DC1**, in **Server Manager**, in the navigation pane, click **Storage Pools**.
2. In the **STORAGE POOLS** pane, click **TASKS**, and then in the **TASKS** drop-down list, click **New Storage Pool**.
3. In the **New Storage Pool Wizard** window, on the **Before you begin** page, click **Next**.
4. On the **Specify a storage pool name and subsystem** page, in the **Name** box, type **StoragePool1**, ensure that in the available disks group **LON-DC1 – Primordial** is selected, and then click **Next**.
5. On the **Select physical disks for the storage pool** page, select all five physical disks, and then click **Next**.
6. On the **Confirm selections** page, click **Create**.
7. On the **View results** page, wait until the creation is completed, and then click **Close**.

► Task 2: Create a three-way mirrored disk

1. In **Server Manager**, in the **STORAGE POOLS** pane, click **StoragePool1**.
2. In the **VIRTUAL DISKS** pane, click **TASKS**, and then in the **TASKS** drop-down list, click **New Virtual Disk**.
3. In the **New Virtual Disk Wizard** window, on the **Before you begin** page, click **Next**.
4. On the **Select the storage pool** page, click **StoragePool1**, and then click **Next**.
5. On the **Specify the virtual disk name** page, in the **Name** box, type **Mirrored vDisk**, and then click **Next**.
6. On the **Specify enclosure resiliency** page click **Next**.
7. On the **Select the storage layout** page, in the **Layout** list, click **Mirror**, and then click **Next**.
8. On the **Configure the resiliency settings** page, select the **Three-way mirror** option, and then click **Next**.
9. On the **Specify the provisioning type** page, select the **Thin** option, and then click **Next**.
10. On the **Specify the size of the virtual disk** page, in the **Specify size** box, type **10**, and then click **Next**.
11. On the **Confirm selections** page, click **Create**.
12. On the **View results** page, wait until the creation is completed, ensure **Create a volume when this wizard closes** is selected, and then click **Close**.
13. In the **New Volume Wizard** window, on the **Before you begin** page, click **Next**.
14. On the **Select the server and disk** page, in the **Disk** pane, click the virtual disk named **Mirrored vDisk**, and then click **Next**.
15. On the **Specify the size of the volume** page, to confirm the default selection, click **Next**.
16. On the **Assign to a drive letter or folder** page, ensure **F** is selected in the **Drive letter** drop-down list, and then click **Next**.

17. On the **Select file system settings** page, in the **File system** drop-down list, select **ReFS**, and in the **Volume label** box, type **Mirrored Volume**, and then click **Next**.
18. On the **Confirm selections** page, click **Create**.
19. On the **Completion** page, wait until the creation is completed, and then click **Close**.

► **Task 3: Copy a file to the volume, and verify visibility in Windows Explorer**

1. On the taskbar, click **Start** and then click **Windows PowerShell**.
2. At the Windows PowerShell prompt, type the following command, and then press Enter:

```
Copy C:\windows\system32\write.exe F:\
```

3. Close the Windows PowerShell prompt.
4. On the taskbar, click the **File Explorer** icon, expand **This PC**, and then click **Mirrored Volume (F:)**. You should now see Write.exe in the file list.
5. Close **File Explorer**.

► **Task 4: Disconnect an iSCSI disk**

1. Switch to **LON-SVR2**.
2. Open **Server Manager**, and in the **iSCSI VIRTUAL DISKS** pane, in the **LON-SVR2** list, right-click **iSCSIDisk1.vhdx**, and then click **Disable iSCSI Virtual Disk**.
3. In the **Disable iSCSI Virtual Disk warning** message box, click **Yes**.

► **Task 5: Verify that the file still is accessible, and check the virtual disk's health**

1. Switch to **LON-DC1**.
2. On the taskbar, click the **File Explorer** icon, and then click **Mirrored Volume (F:)**.
3. In the file list pane, double-click **write.exe** to ensure access to the volume still is available.
4. Close the **Document - WordPad** window.
5. Close **File Explorer**.
6. In **Server Manager**, on the menu bar, click the **Refresh** button. Wait until all panes are refreshed. Notice the warning that appears next to **Mirrored vDisk**. The result may vary slightly. A warning may also appear in the **physical disks** section. If the status for **StoragePool1** does not change:
 - a. Restart **LON-DC1**.
 - b. Sign in as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
 - c. In **Server Manager**, click **File and Storage Services**, and then click **Storage Pools**.
7. In the **VIRTUAL DISKS** pane, right-click **Mirrored vDisk**, and then in the drop-down list, select **Properties**.
8. In the **Mirrored vDisk Properties** window, in the navigation pane, click **Health**.
9. Click **OK**.

► **Task 6: Add a new iSCSI virtual disk**

1. Switch to **LON-SVR2**.
2. In **Server Manager**, in the navigation pane, click **File and Storage Services**.
3. In the **File and Storage Services** pane, click **iSCSI**.

4. In the **iSCSI VIRTUAL DISKS** pane, click **TASKS**, and then in the **TASKS** drop-down list, select **New iSCSI Virtual Disk**.
5. In the **New iSCSI Virtual Disk Wizard**, on the **Select iSCSI virtual disk location** page, under **Storage location**, click **C:**, and then click **Next**.
6. On the **Specify iSCSI virtual disk name** page, in the **Name** box, type **iSCSIDisk6**, and then click **Next**.
7. On the **Specify iSCSI virtual disk size** page, in the **Size** box, type **5**, ensure **GB** is selected in the drop-down list, and then click **Next**.
8. On the **Assign iSCSI target** page, click **lon-dc1**, and then click **Next**.
9. On the **Confirm selections** page, click **Create**.
10. On the **View results** page, wait until the creation is completed, and then click **Close**.

► **Task 7: Add the new disk to the storage pool, and extend the virtual disk**

1. Switch to **LON-DC1**.
2. In **Server Manager**, on the menu bar, click the **Refresh** button.
3. Wait for all of the panes to refresh.
4. In the **STORAGE POOLS** pane, right-click **StoragePool1**, and then in the drop-down list, click **Add Physical Disk**.
5. In the **Add Physical Disk** window, select the disk that you see in the list, and then click **OK**.
6. In the **VIRTUAL DISKS** pane, right-click **Mirrored vDisk**, and then in the drop-down list, select **Extend Virtual Disk**.
7. In the **Extend Virtual Disk** window, in the **Specify size** box, type **15**, and then click **OK**.

Results: After completing this exercise, you should have successfully created a storage pool, added a new disk to the storage pool, and extended the disk.

► **Task 8: Prepare for the next exercise**

When you finish Exercise 1, revert the virtual machines to their initial state. To do this, perform the following steps:

1. On the host computer, start **Hyper-V Manager**.
2. In the **Virtual Machines** list, right-click **20743A-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20743A-LON-SVR2** and **20743A-LON-CL1**.

Exercise 2: Implementing the Storage Spaces Direct feature

► **Task 1: Install the Windows Server roles and features**

1. On the host computer, start **Hyper-V Manager**.
2. In **Hyper-V Manager**, click **20743A-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 and 3 for **20743A-LON-SVR1**, **20743A-LON-SVR2**, and **20743A-LON-SVR3**.
6. Repeat step 4 for **LON-SVR1**.
7. Switch to **LON-DC1**.
8. In **Server Manager**, click **All Servers**.
9. Right-click **All Servers**, and then click **Add Servers**.
10. In the **Add Servers** window, click **Find Now**.
11. Select **LON-SVR1**, **LON-SVR2**, and **LON-SVR3** and click the **arrow** button to move the three servers to the **Selected** list. Click **OK**.
12. Verify that all three servers have a **Manageability** of **Online – Performance counters not started** before continuing.
13. Click **Start**, and then click **Windows PowerShell ISE**.
14. In the **Administrator: Windows PowerShell ISE** window, click **File**, and then click **Open**.
15. In the **File name** box, type **E:\Labfiles\Mod02**, and then press Enter.
16. Double-click the **Implement-StorageSpacesDirect** file.
17. Select the first line in step 0, starting with **Install-Windowsfeature**, and then press F8. Wait until the installation finishes.
18. If the column **Restart Needed** has a value of **Yes**, select the line starting with **Restart-Computer**, and then press F8. After the server restarts, sign in as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
19. Repeat steps 13 through to 16. If a dialog box opens, click **OK**.
20. Select the line in step 1, starting with **Invoke-Command**, and then press F8. Wait until the installation finishes.
21. Verify that the output of the command includes four lines with **Success** as **True**.
22. Switch to **LON-SVR1**. Click **Start**, and then click **Server Manager**.
23. In **Server Manager**, click **Tools**, and then click **Failover Cluster Manager**.
24. Notice that the tool opens. This verifies that the command was successful.

► Task 2: Validate cluster configuration

1. Switch to **LON-DC1**.
2. In the **Administrator: Windows PowerShell ISE** window, select the line in step 2, starting with **Test-Cluster**, and then press F8. Wait until the installation finishes.
3. Verify that the output of the command only includes warnings and that the last line is a validation report in html format. This validates that the command was successful.

► Task 3: Create a cluster

1. In the **Administrator: Windows PowerShell ISE** window, select the line in step 3, starting with **New-Cluster**, and then press F8. Wait until the installation finishes.
2. Verify that the output of the command only includes warnings, and that the last line has a **Name** column with the value **S2DCluster**. This validates that the command was successful.
3. Switch to **LON-SVR1**, and then in the **Failover Cluster Manager** window, in the **Actions** pane, click **Connect to Cluster**.
4. In the **Select Cluster** dialog box, type **S2DCluster**, and then click **OK**.
5. Notice that the cluster opens. This verifies that the command was successful. If you receive an error message stating **Cluster S2DCluster not found**, perform the following steps:
 - a. Restart **LON-SVR1**.
 - b. Sign in as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
 - c. Click **Start**, and then click **Server Manager**.
 - d. In **Server Manager**, click **Tools**, and then click **Failover Cluster Manager**.
 - e. **S2DCluster.Adatum.com** should now appear in the navigation pane. If necessary, perform steps 3 and 4.

► Task 4: Enable the Storage Spaces Direct feature

1. Switch to **LON-DC1**.
2. In the **Administrator: Windows PowerShell ISE** window, select the line in step 4, starting with **Invoke-Command**, and then press F8. Wait until the installation finishes.
3. In the **Confirm** dialog box, click **Yes**.
4. Verify that there is no output of the command. This validates that the command was successful.

► Task 5: Create a storage pool

1. In the **Administrator: Windows PowerShell ISE** window, select the line in step 5, starting with **Invoke-Command**, and then press F8. Wait until the installation finishes.
2. In the output of the command, verify that the **FriendlyName** attribute has a value of **S2DStoragePool**. This validates that the command was successful.
3. Switch to **LON-SVR1**, and then in the **Failover Cluster Manager** window, expand **S2DCluster.adatum.com**.
4. Expand **Storage**, and then click **Pools**.
5. Verify the existence of **Cluster Pool 1**. This verifies that the command was successful.

► Task 6: Create a virtual disk

1. Switch to **LON-DC1**. In the **Administrator: Windows PowerShell ISE** window, select the line in step 6, starting with **Invoke-Command**, and then press F8. Wait until the installation finishes.
2. Verify that in the output of the command is an attribute **FileSystemLabel**, with a value of **CSV**. This validates that the command was successful.
3. Switch to **LON-SVR1**, and then in the **Failover Cluster Manager** window, click **Disks**.
4. Verify the existence of **Cluster Virtual Disk (CSV)**. This verifies that the command was successful.

► Task 7: Create a file server and file share

1. Switch to **LON-DC1**. In the **Administrator: Windows PowerShell ISE** window, select the line in step 7, starting with **Invoke-Command**, and then press F8. Wait until the installation finishes.
2. Verify that in the output of the command is an attribute **FriendlyName**, with a value of **S2D-SOFS**. This validates that the command was successful.
3. Switch to **LON-SVR1**, and then in the **Failover Cluster Manager** window, click **Roles**.
4. Verify the existence of **S2D-SOFS**. This verifies that the command was successful.
5. Switch to **LON-DC1**. In the **Administrator: Windows PowerShell ISE** window, select the three lines in step 8, starting with **Invoke-Command**, and then press F8. Wait until the installation finishes.
6. Verify that there in the output of the command is an attribute **Path** with a value of **C:\ClusterStorage\Volume1\VM01**. This validates that the command was successful.
7. Switch to **LON-SVR1**, and then in the **Failover Cluster Manager** window, click **S2D-SOFS**, and then click the **Shares** tab.
8. Verify the existence of **VM01**. This verifies that the command was successful.

► Task 8: Test high availability for the storage

1. Switch to **LON-DC1**.
2. On the taskbar, click **File Explorer**.
3. In the **address** text box, type **\\s2d-sofs\VM01** and then press Enter.
4. On the **Home** tab, click **New Folder**, and then type **VMFolder**.
5. Double-click **VMFolder**.
6. Switch to the **Administrator: Windows PowerShell ISE** window and at the Windows PowerShell prompt, type the following cmdlet, and then press Enter:

```
Stop-Computer -computername LON-SVR2
```

7. Switch to the **Server Manager** window and click **All Servers**.
8. In the **Servers** list, click **LON-SVR2**. Verify that **Manageability** changes to **Target computer not accessible**. This verifies that the computer is no longer turned on.
9. Switch to the **File Explorer** window.
10. On the **Home** tab, click **New item**, click **Text Document** and then press Enter. This verifies that the storage is still available with one server turned off.
11. Switch to **LON-SVR1**. In **Failover Cluster Manager**, click **Disks** and then click **Cluster Virtual Disk (CSV)**. Verify that for the **Cluster Virtual Disk (CSV)**, the **Health Status** is **Warning** and **Operational Status** is **Degraded**. **Operational Status** may also display as **Incomplete**.

Results: After completing this exercise, you should have implemented Storage Spaces Direct successfully.

► Task 9: Prepare for the next module

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

1. On the host computer, start **Hyper-V Manager**.
2. In the **Virtual Machines** list, right-click **20743A-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20743A-LON-SVR1**, **20743A-LON-SVR2**, and **20743A-LON-SVR3**.

Module 3: Implementing Directory Services

Lab: Implementing and managing AD DS

Exercise 1: Cloning a domain controller

► Task 1: Prepare a source domain controller to be cloned

1. On **LON-DC1**, in **Server Manager**, click **Tools**, and then click **Active Directory Administrative Center**.
2. In **Active Directory Administrative Center**, double-click **Adatum (local)**, and then in the list, double-click the **Domain Controllers** organizational unit (OU).
3. In the list, select **LON-DC1** if it is not already selected, and then in the **Tasks** pane, in the **LON-DC1** section, click **Add to group**.
4. In the **Select Groups** dialog box, in the **Enter the object names to select** text box, type **Cloneable**, and then click **Check Names**.
5. Ensure that the group name is expanded to **Cloneable Domain Controllers**, and then click **OK**.
6. Click **Start**, and then click **Windows PowerShell**.
7. At the command prompt in the Windows PowerShell command-line interface, type the following command, and then press Enter:

```
Get-ADDCCloningExcludedApplicationList
```

8. Verify the list of critical apps. In production, you need to verify each app or use a domain controller that has fewer apps installed by default. Type the following command, and then press Enter:

```
Get-ADDCCloningExcludedApplicationList -GenerateXML
```

9. Run the following command to create the **DCCloneConfig.xml** file:

```
New-ADDCCloneConfigFile -CloneComputerName "LON-DC3"
```

10. Type the following command to shut down **LON-DC1**, and then press Enter:

```
Stop-Computer
```

11. Wait for the machine to shut down. You might be asked to confirm the shutdown.

► Task 2: Export the source virtual machine

1. On the host computer, in **Hyper-V Manager**, in the **details** pane, select the **20743A-LON-DC1** virtual machine.
2. In the **Actions** pane, in the **20743A-LON-DC1** section, click **Export**.
3. In the **Export Virtual Machine** dialog box, navigate to **D:\Program Files\Microsoft Learning\20743**, and then click **Export**. Wait until the export finishes. This can take from 10 to 15 minutes.
4. Start and connect to **20743A-LON-DC1** and sign in as **Adatum\Administrator** with the password of **Pa\$\$w0rd**.

► Task 3: Create and start the cloned domain controller

1. On the host computer, in **Hyper-V Manager**, in the **Actions** pane, in the section that is named for the host computer, click **Import Virtual Machine**.
2. In the **Import Virtual Machine Wizard**, on the **Before You Begin** page, click **Next**.
3. On the **Locate Folder** page, click **Browse**, go to **D:\Program Files\Microsoft Learning\20743\20743A-LON-DC1**, click **Select Folder**, and then click **Next**.
4. On the **Select Virtual Machine** page, select **20743A-LON-DC1** if it is not already selected, and then click **Next**.
5. On the **Choose Import Type** page, select **Copy the virtual machine (create a new unique ID)**, and then click **Next**.
6. On the **Choose Folders for Virtual Machine Files** page, select the **Store the virtual machine in a different location** check box. For each folder location, specify **D:\Program Files\Microsoft Learning\20743** as the path, and then click **Next**.
7. On the **Choose Folders to Store Virtual Hard Disks** page, specify the **D:\Program Files\Microsoft Learning\20743** path, and then click **Next**.
8. On the **Completing Import Wizard** page, click **Finish**. The machine imports. This can take from 10 to 15 minutes or longer.



Note: You can continue with the next exercise while the import proceeds.

9. After the import completes, in the **management** list, identify and select the newly imported **20743A-LON-DC1** virtual machine, which has the **State** value as **Off**. In the lower section of the **Actions** pane, click **Rename**.
10. Type **20743A-LON-DC3** as the name, and then press Enter.
11. In the **Actions** pane, in the **20743A-LON-DC3** section, click **Start**, and then click **Connect** to see the virtual machine starting.
12. While the server is starting, a "Domain Controller cloning is at x% completion" message displays.

Results: After completing this exercise, you should have successfully cloned a domain controller.

Exercise 2: Implementing service accounts

► Task 1: Create and associate a managed service account

1. On **LON-DC1**, click **Start**, and then click **Windows PowerShell**.
2. In the Windows PowerShell window, at the command prompt, type the following command, and then press Enter:

```
Add-KdsRootKey -EffectiveTime ((get-date).addhours(-10))
```

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

```
New-ADServiceAccount -Name Webservice -DNSHostName LON-DC1 -  
PrincipalsAllowedToRetrieveManagedPassword LON-DC1$
```

4. Type the following command, and then press Enter:

```
Add-ADComputerServiceAccount -identity LON-DC1 -ServiceAccount Webservice
```

5. Type the following command, and then press Enter:

```
Get-ADServiceAccount -Filter *
```

6. Note the output of the command.
7. Type the following command, and then press Enter:

```
Install-ADServiceAccount -Identity Webservice
```

8. Minimize the Windows PowerShell window.

► **Task 2: Configure the Web server application pool to use the group managed service account**

1. On **LON-DC1**, in **Server Manager**, click the **Tools** menu, and then click **Internet Information Services (IIS) Manager**.
2. In the **Internet Information Services (IIS) Manager** console, expand **LON-DC1 (Adatum\Administrator)**. Click **Application Pools**.
3. In the details pane, right-click **DefaultAppPool**, and then click **Advanced Settings**.
4. In the **Advanced Settings** dialog box, click **Identity**, and then click the ellipses (...).
5. In the **Application Pool Identity** dialog box, click **Custom Account**, and then click **Set**.
6. In the **Set Credentials** dialog box, in the **User name** text box, type **Adatum\Webservice\$**, and then click **OK** three times.
7. In the **Actions** pane, click **Stop** to stop the application pool.
8. To start the application pool, click **Start**.
9. Verify that the identity of the **DefaultAppPool** is set to **adatum\webservice\$**.
10. Close the **Internet Information Services (IIS) Manager**.



Note: If you did not complete Exercise 1, "Cloning a domain controller," do so before reverting the virtual machines.

Results: After completing this exercise, you should have successfully implemented service accounts.

► **Task 3: Prepare for the next module**

When you have finished the lab, revert the virtual machines to their initial state:

1. On the host computer, start **Hyper-V Manager**.
2. In the **Virtual Machines** list, right-click **20743A-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Shut down **20743A-LON-DC3**.

MCT USE ONLY. STUDENT USE PROHIBITED

Module 4: Implementing AD FS

Lab A: Implementing AD FS

Exercise 1: Installing and configuring AD FS

► Task 1: Create a DNS record for AD FS

1. On **LON-DC1**, in **Server Manager**, click **Tools**, and then click **DNS**.
2. In **DNS Manager**, expand **LON-DC1**, expand **Forward Lookup Zones**, and then click **Adatum.com**.
3. Right-click **Adatum.com**, and then click **New Host (A or AAAA)**.
4. In the **New Host** window, in the **Name** box, type **adfs**.
5. In the **IP address** box, type **172.16.0.12**, and then click **Add Host**.
6. In the **DNS** window, click **OK**.
7. Click **Done**, and then close **DNS Manager**.

► Task 2: Install AD FS

1. On **LON-SVR2**, click **Start**, and then click **Server Manager**, click **Manage**, and then click **Add Roles and Features**.
2. In the **Add Roles and Features Wizard**, on the **Before you begin** page, click **Next**.
3. On the **Select installation type** page, click **Role-based or feature-based installation**, and then click **Next**.
4. On the **Select destination server** page, click **Select a server from the server pool**, click **LON-SVR2.Adatum.com**, and then click **Next**.
5. On the **Select server roles** page, select the **Active Directory Federation Services** check box, and then click **Next**.
6. On the **Select features** page, click **Next**.
7. On the **Active Directory Federation Services (AD FS)** page, click **Next**.
8. On the **Confirm installation selections** page, click **Install**.
9. When the installation is complete, click **Close**.

► Task 3: Configure AD FS

1. On **LON-SVR2**, 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**.



Note: The adfs.adatum.com certificate was preconfigured for this task. In your own environment, you must obtain this certificate.

► Task 4: Verify AD FS functionality

1. On **LON-CL1**, sign in as **Adatum\Beth** with the password **Pa\$\$w0rd**.
2. Click **Start**, click **All apps**, click **Windows Accessories**, and then click **Internet Explorer**.
3. In **Internet Explorer**, in the **Address** bar, type **https://adfs.adatum.com/federationmetadata/2007-06/federationmetadata.xml**, and then press Enter.
4. Verify that the file loads, and then close **Internet Explorer**.

Results: After completing this exercise, you will have installed and configured AD FS. You also will have verified that it is functioning by viewing the FederationMetaData.xml file contents.

Exercise 2: Configuring an internal application for AD FS

► Task 1: Configure the Active Directory claims provider trust

1. On **LON-SVR2**, in **Server Manager**, click **Tools**, and then click **AD FS Management**.
2. In the **AD FS Management** console, click **Claims Provider Trusts**.
3. In the middle pane, right-click **Active Directory**, and then click **Edit Claim Rules**.
4. In the **Edit Claims Rules for Active Directory** window, on the **Acceptance Transform Rules** tab, click **Add Rule**.
5. In the **Add Transform Claim Rule Wizard**, on the **Select Rule Template** page, in the **Claim rule template** list, select **Send LDAP Attributes as Claims**, and then click **Next**.
6. On the **Configure Rule** page, in the **Claim rule name** box, type **Outbound LDAP Attributes Rule**.
7. In the **Attribute store** drop-down list, select **Active Directory**.
8. In the **Mapping of LDAP attributes to outgoing claim types** section, select the following values for the LDAP Attribute and the Outgoing Claim Type, and then click **Finish**:
 - E-Mail-Addresses: **E-Mail Address**
 - User-Principal-Name: **UPN**
 - Display-Name: **Name**
9. In the **Edit Claim Rules for Active Directory** window, click **OK**.

► **Task 2: Configure the application to trust incoming claims**

1. On **LON-SVR1**, open **Server Manager**, click **Tools**, and then click **Windows Identity Foundation Federation Utility**.
2. On the **Welcome to the Federation Utility Wizard** page, in the **Application configuration location** box, type **C:\inetpub\wwwroot\AdatumTestApp\web.config** for the location of the sample web.config file.
3. In the **Application URI** box, type **https://lon-svr1.adatum.com/AdatumTestApp/** to indicate the path to the sample application that will trust the incoming claims from the federation server, and then click **Next** to continue.
4. On the **Security Token Service** page, click **Use an existing STS**, in the **STS WS-Federation metadata document location** box, type **https://adfs.adatum.com/federationmetadata/2007-06/federationmetadata.xml**, and then click **Next** to continue.
5. On the **STS signing certificate chain validation error** page, click **Disable certificate chain validation**, and then click **Next**.
6. On the **Security token encryption** page, click **No encryption**, and then click **Next**.
7. On the **Offered claims** page, review the claims that will be offered by the federation server, and then click **Next**.
8. On the **Summary** page, review the changes that will be made to the sample application by the **Federation Utility Wizard**, scroll through the items to understand what each item is doing, and then click **Finish**.
9. In the **Success** window, click **OK**.

► **Task 3: Configure a relying party trust for the claims-aware application**

1. On **LON-SVR2**, in the **AD FS** console, click **Relying Party Trusts**.
2. In the **Actions** 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, click **Import data about the relying party published online or on a local network**.
5. In the **Federation Metadata address (host name or URL)** box, type **https://lon-svr1.adatum.com/adatumtestapp/**, and then click **Next**. This downloads the metadata configured in the previous task.
6. On the **Specify Display Name** page, in the **Display name** box, type **A. Datum Test App**, and then click **Next**.
7. On the **Choose Access Control Policy** page, click **Permit everyone**, and then click **Next**.
8. On the **Ready to Add Trust** page, review the relying party trust settings, and then click **Next**.
9. On the **Finish** page, click **Close**.
10. Leave the **Edit Claims Issuance Policy for A. Datum Test App** window open for the next task. (This might be hidden behind Server Manager.)

► **Task 4: Configure claim rules for the relying party trust**

1. On **LON-SVR2**, in the **AD FS Management** console, in the **Edit Claim Issuance Policy for A. Datum Test App** window, on the **Issuance Transform Rules** tab, click **Add Rule**.
2. In the **Claim rule template** list, select **Pass Through or Filter an Incoming Claim**, and then click **Next**.
3. In the **Claim rule name** box, type **Pass through Windows account name**.
4. In the **Incoming claim type** drop-down list, click **Windows account name**, and then click **Finish**.
5. On the **Issuance Transform Rules** tab, click **Add Rule**.
6. In the **Claim rule template** box, select **Pass Through or Filter an Incoming Claim**, and then click **Next**.
7. In the **Claim rule name** box, type **Pass through E-Mail Address**.
8. In the **Incoming claim type** drop-down list, click **E-Mail Address**, and then click **Finish**.
9. On the **Issuance Transform Rules** tab, click **Add Rule**.
10. In the **Claim rule template** box, select **Pass Through or Filter an Incoming Claim**, and then click **Next**.
11. In the **Claim rule name** box, type **Pass through UPN**.
12. In the **Incoming claim type** drop-down list, click **UPN**, and then click **Finish**.
13. On the **Issuance Transform Rules** tab, click **Add Rule**.
14. In the **Claim rule template** box, select **Pass Through or Filter an Incoming Claim**, and then click **Next**.
15. In the **Claim rule name** box, type **Pass through Name**.
16. In the **Incoming claim type** drop-down list, click **Name**, and then click **Finish**.
17. On the **Issuance Transform Rules** tab, click **OK**.

► **Task 5: Test access to the claims-aware application**

1. On **LON-CL1**, click **Start**, click **All apps**, click **Windows Accessories**, and then click **Internet Explorer**.
2. In **Internet Explorer**, in the **Address** bar, type **https://lon-svr1.adatum.com/AdatumTestApp/**, and then press Enter.



Note: It is critical to use the trailing slash in the URL for Step 2.

3. In the **Windows Security** window, sign in as **Adatum\Beth** with the password **Pa\$\$w0rd**.
4. Review the claim information that is displayed by the application.
5. Close **Internet Explorer**.

► **Task 6: Configure Internet Explorer to pass local credentials to the application automatically**

1. On **LON-CL1**, click **Start**, type **Internet Options**, and then click **Internet Options**.
2. In the **Internet Properties** window, on the **Security** tab, click **Local intranet**, and then click **Sites**.
3. In the **Local intranet** window, click **Advanced**.
4. In the **Local intranet** window, in the **Add this website to the zone** box, type **https://adfs.adatum.com**, and then click **Add**.
5. In the **Add this website to the zone** box, type **https://lon-svr1.adatum.com**, click **Add**, and then click **Close**.
6. In the **Local intranet** window, click **OK**.
7. In the **Internet Properties** window, click **OK**.
8. Open **Internet Explorer**.
9. In **Internet Explorer**, in the Address bar, type **https://lon-svr1.adatum.com/AdatumTestApp/**, and then press Enter.



Note: It is critical to use the trailing slash in the URL for Step 9.

10. Notice that you did not receive a prompt for credentials.
11. Review the claim information that is displayed by the application.
12. Close **Internet Explorer**.

► **Task 7: Prepare for the next lab**

- Leave all the virtual machines running at the end of this lab.

Results: After completing this exercise, you will have configured AD FS to support authentication for an application.

Lab B: Implementing Web Application Proxy

Exercise 1: Implementing Web Application Proxy

► Task 1: Install Web Application Proxy

1. On **LON-SVR3**, open **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-SVR3.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 2: Add the adfs.adatum.com certificate to LON-SVR3

1. On **LON-SVR2**, click **Start**, 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, to close the success message, click **OK**.
16. Close the **Microsoft Management** console, and do not save the changes.
17. On **LON-SVR3**, click **Start**, 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-SVR2\c\$\adfs.pfx**, and then click **Next**.
27. On the **Private key protection** page, in the **Password** box, type **Pa\$\$w0rd**.
28. Select the **Mark this key as exportable** check box, and then click **Next**.
29. On the **Certificate Store** page, click **Place all certificates in the following store**.
30. In the **Certificate store** box, select **Personal**, and then click **Next**.
31. On the **Completing the Certificate Import Wizard** page, click **Finish**, and then, to clear the success message, click **OK**.

► **Task 3: Add the LON-SVR1.adatum.com certificate to LON-SVR3**

1. On **LON-SVR1**, click **Start**, type **mmc**, and then press Enter.
2. In the **Microsoft Management** console, click **File**, and then click **Add/Remove Snap-in**.
3. In the **Add or Remove Snap-ins** window, in the **Available snap-ins** column, double-click **Certificates**.
4. In the **Certificates snap-in** window, click **Computer account**, and then click **Next**.
5. In the **Select Computer** window, click **Local Computer (the computer this console is running on)**, and then click **Finish**.
6. In the **Add or Remove Snap-ins** window, click **OK**.
7. In the **Microsoft Management** console, expand **Certificates (Local Computer)**, expand **Personal**, and then click **Certificates**.
8. Right-click **LON-SVR1.adatum.com**, point to **All Tasks**, and then click **Export**.
9. In the **Certificate Export Wizard**, click **Next**.
10. On the **Export Private Key** page, click **Yes, export the private key**, and then click **Next**.
11. On the **Export File Format** page, click **Next**.
12. On the **Security** page, select the **Password** check box.

13. In the **Password** and **Confirm password** boxes, type **Pa\$\$w0rd**, and then click **Next**.
14. On the **File to Export** page, in the **File name** box, type **C:\lon-svr1.pfx**, and then click **Next**.
15. On the **Completing the Certificate Export Wizard** page, click **Finish**, and then, to close the success message, click **OK**.
16. Close the **Microsoft Management** console, and do not save the changes.
17. On **LON-SVR3**, switch to **Console1**.
18. In the **Microsoft Management** console, expand **Certificates (Local Computer)**, and then click **Personal**.
19. Right-click **Personal**, point to **All Tasks**, and then click **Import**.
20. In the **Certificate Import Wizard**, click **Next**.
21. On the **File to Import** page, in the **File name** box, type **\\LON-SVR1\c\$\lon-svr1.pfx**, and then click **Next**.
22. On the **Private key protection** page, in the **Password** box, type **Pa\$\$w0rd**.
23. Select the **Mark this key as exportable** check box, and then click **Next**.
24. On the **Certificate Store** page, click **Place all certificates in the following store**.
25. In the **Certificate store** box, select **Personal**, and then click **Next**.
26. On the **Completing the Certificate Import Wizard** page, click **Finish**, and then, to clear the success message, click **OK**.
27. Close the **Microsoft Management** console, and do not save the changes.

► Task 4: Configure Web Application Proxy

1. On **LON-SVR3**, in **Server Manager**, click the **Notifications** icon, and then click **Open the Web Application Proxy Wizard**.
2. In the **Web Application Proxy Configuration Wizard**, on the **Welcome** page, click **Next**.
3. On the **Federation Server** page, enter the following, and then click **Next**:
 - Federation service name: **adfs.adatum.com**
 - User name: **Adatum\Administrator**
 - Password: **Pa\$\$w0rd**
4. On the **AD FS Proxy Certificate** page, in the **Select a certificate to be used by the AD FS proxy** box, select **adfs.adatum.com**, and then click **Next**.
5. On the **Confirmation** page, click **Configure**.
6. On the **Results** page, click **Close**.
7. The **Remote Access Management** console opens automatically. Leave it open for the next task.


► Task 5: Configure the AD FS Proxy for the test application

1. On **LON-SVR3**, 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 **Active Directory Federation Services (AD FS)**, and then click **Next**.
4. On the **Supported Clients** page, click **Next**.


5. On the **Relying Party** page, click **A. Datum Test App**, and then click **Next**.
6. On the **Publishing Settings** page, in the **Name** box, type **A. Datum Test App Rule**.
7. In the **External URL** box, type **https://lon-svr1.adatum.com/adatumtestapp/**.
8. In the **External certificate** box, select **lon-svr1.adatum.com**.
9. In the **Backend server URL** box, type **https://lon-svr1.adatum.com/adatumtestapp/**, and then click **Next**.
10. On the **Confirmation** page, click **Publish**.
11. On the **Results** page, click **Close**.

► Task 6: Test Web Application Proxy

1. Switch to **LON-CL3**, and sign in as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. Click **Start**, and 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 two lines, click **File**, and then click **Save**:
 - **172.16.0.13 adfs.adatum.com**
 - **172.16.0.13 lon-svr1.adatum.com**

 **Note:** You edit the hosts to force **LON-CL3** to access the application through Web Application Proxy. In a production environment, you do this by using split DNS. The IPv4 address is for **LON-SVR3**, the Web Application Proxy.

7. Close Notepad.
8. Open **Internet Explorer**.
9. In **Internet Explorer**, in the Address bar, type **https://lon-svr1.adatum.com/adatumtestapp/**, and then press Enter.

 **Note:** If you receive two certificate errors, ignore them.

10. In the **A. Datum Corporation** dialog box, sign in as **Adatum\Beth** with the password **Pa\$\$w0rd**.
11. It is not necessary to store the password for this site. Review the claims, and then close **Internet Explorer**.

► Task 7: Prepare for the next module

When you have finished the lab, revert the virtual machines to their initial state.

1. On the host computer, start **Hyper-V Manager**.
2. In the **Virtual Machines list**, right-click **20743A-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20743A-LON-SVR1**, **20743A-LON-SVR2**, **20743A-LON-SVR3**, **20743A-LON-CL3**, and **20743A-LON-CL1**.

Results: After completing this exercise, you will have configured Web Application Proxy to secure access to AdatumTestApp from the Internet.

Module 5: Implementing network services

Lab: Implementing network services

Exercise 1: Configuring DNS policies

► Task 1: Check DNS name resolution before configuring DNS policies

1. Switch to **LON-CL1**.
2. On **LON-CL1**, right-click **Start**, and then click **Command Prompt**.
3. At the command prompt, type the following two commands, pressing Enter after each command:

```
ipconfig /flushdns  
nslookup www.adatum.com
```

4. Verify that the last command returns an IP address of **172.16.0.10**.
5. Switch to **TREY-DC1**.
6. On **TREY-DC1**, click **Start**, and then click **Windows PowerShell**.
7. In the **Windows PowerShell** window, type the following command, and then press Enter:

```
Add-DnsServerConditionalForwarderZone -Name "adatum.com" -MasterServers 172.16.0.10
```

8. In the **Windows PowerShell** window, type the following three commands, pressing Enter after each command:

```
Clear-DNSServerCache  
Clear-DNSClientCache  
Resolve-DNSName adatum.com
```

9. When you receive a prompt, type **Y**, and then press Enter. Verify that the last command returns an IP address of **172.16.0.10**. If the last command returns an error, retry the last command.

► Task 2: Configure DNS policies

1. Switch to **LON-DC1**.
2. On **LON-DC1**, click **Start**, and then click **Windows PowerShell**.
3. In the **Windows PowerShell** window, type the following two commands, pressing Enter after each command:

```
Add-DnsServerZoneScope -ZoneName "adatum.com" -Name "partners"  
Add-DNSServerClientSubnet -Name "TreyPartner" -IPv4Subnet 172.16.10.0/24
```

4. In the **Windows PowerShell** window, type the following command, and then press Enter:

```
Add-DnsServerResourceRecord -ZoneName "adatum.com" -A -Name "www" -IPv4Address  
"131.107.0.200" -ZoneScope "partners"
```

5. In the **Windows PowerShell** window, type the following command, and then press Enter:

```
Add-DnsServerQueryResolutionPolicy -Name "SplitBrainZonePolicy" -Action ALLOW -  
ClientSubnet "eq,TreyPartner" -ZoneScope "partners,1" -ZoneName adatum.com
```

► Task 3: Check DNS name resolution after configuring DNS policies

1. Switch to **LON-CL1**.
2. On **LON-CL1**, at the command prompt, type the following two commands, pressing Enter after each command:

```
ipconfig /flushdns  
nslookup www.adatum.com
```

3. Verify that the last command still returns an IP address of **172.16.0.10**.
4. Switch to **TREY-DC1**.
5. On **TREY-DC1**, in the **Windows PowerShell** window, type the following three commands, pressing Enter after each command:

```
Clear-DNSServerCache  
Clear-DNSClientCache  
Resolve-DNSName www.adatum.com
```

6. When you receive a prompt, press **Y**, and then press Enter. Verify that the last command returns an IP address of **131.107.0.200**. If the last command returns an error, retry the last command.

Results: After completing this exercise, you should have successfully configured DNS policies for split-brain DNS and then tested that the policies work.

Exercise 2: Configuring DHCP failover**► Task 1: Install DHCP on an additional server**

1. Switch to **LON-SVR1**.
2. On **LON-SVR1**, click **Start**, and then click **Server Manager**.
3. In the **Server Manager** console, click **Add roles and features**.
4. In **Add Roles and Features Wizard**, click **Next**.
5. On the **Select installation type** page, click **Next**.
6. On the **Select destination server** page, verify that in the **Server Pool** list, **LON-SVR1.adatum.com** is selected, and then click **Next**.
7. On the **Select server roles** page, select the **DHCP Server** check box.
8. In the **Add Roles and Features Wizard** dialog box, click **Add features**.
9. On the **Select server roles** page, click **Next**.
10. On the **Select features** page, click **Next**.
11. On the **DHCP Server** page, click **Next**.
12. On the **Confirm installation selections** page, click **Install**.
13. When the installation finishes, click **Close**.
14. In the menu bar, click the **Notifications** flag, and then click **Complete DHCP Configuration**.
15. In the **DHCP Post-Install configuration wizard**, click **Next**.

16. On the **Authorization** page, click **Commit**.
17. On the **Summary** page, click **Close**.
18. In the **Server Manager** console, click **Tools**, and then click **DHCP**.
19. In the **DHCP** console, expand **lon-svr1.adatum.com**, and then click **IPv4**.
20. Verify that no scopes exist.

► **Task 2: Configure DHCP scopes**

1. Switch to **LON-DC1**.
2. On **LON-DC1**, switch to the **Server Manager** console, click **Tools**, and then click **Active Directory Users and Computers**.
3. In the **Active Directory Users and Computers** console, in the navigation pane, expand **Adatum.com**, and then click **Users**.
4. In the content pane, right-click **DnsUpdateProxy**, and then click **Properties**.
5. In the **DnsUpdateProxy Properties** window, click the **Members** tab, and then click **Add**.
6. In the **Select Users, Contacts, Computers, Service Accounts or Groups** window, click **Object Types**.
7. In the **Object Types** dialog box, select the **Computers** check box, and then click **OK**.
8. In the **Enter the object names to select** dialog box, type **lon-dc1; lon-svr1**, click **Check Names**, and then click **OK**.
9. In the **DnsUpdateProxy Properties** dialog box, click **OK**.
10. Close the **Active Directory Users and Computers** console.
11. In the **Server Manager** console, click **Tools**, and then click **DHCP**.
12. In the **DHCP** console, expand **lon-dc1.adatum.com**, and then click **IPv4**.
13. Right-click **IPv4**, and then click **New Scope**.
14. In **New Scope Wizard**, click **Next**.
15. On the **Scope Name** page, in the **Name** text box, type **Building 2 Scope**, and in the **Description** text box, type **For new building at Adatum HQ**, and then click **Next**.
16. In the **IP Address Range** page, in the **Start IP Address** text box, type **192.168.2.11**, and in the **End IP Address** text box, type **192.168.2.100**. Verify that the **Subnet mask** text box contains **255.255.255.0**, and then click **Next**.
17. On the **Add Exclusions and Delay** page, click **Next**.
18. On the **Lease Duration** page, click **Next**.
19. On the **Configure DHCP options** page, verify that the **Yes, I want to configure these options now** option is selected, and then click **Next**.
20. On the **Router (Default Gateway)** page, click **Next**.
21. On the **Domain Name and DNS Servers** page, in the **IP Address** list, click **172.16.0.10**, and then click **Remove**.
22. In the **IP Address** list, change the IP address to **172.16.0.9**, and then click **Add**.
23. In the **DHCP** dialog box, click **Yes**, and then click **Next**.
24. On the **WINS Servers** page, click **Next**.

25. On the **Activate Scope** page, verify that the **Yes, I want to activate the scope now** option is selected, and then click **Next**.

26. On the **Completing the New Scope Wizard** page, click **Finish**.

► **Task 3: Configure DHCP failover**

1. Switch to **LON-SVR1**.

2. Click **Start**, and then click **Server Manager**.

3. In the **Server Manager** console, click **Tools**, and then click **DHCP**.

4. In the **DHCP** console, right-click **DHCP**, and then click **Add Server**.

5. In the **Add Server** dialog box, select the **This authorized DHCP server** option, click **lon-dc1.adatum.com**, and then click **OK**.

6. In the **DHCP** console, expand **lon-dc1**, click and then right-click **IPv4**, and then click **Configure Failover**.

7. In the **Configure Failover Wizard**, click **Next**.

8. On the **Specify the partner server to use for failover** page, in the **Partner Server** text box, type **172.16.0.11**, and then click **Next**.

9. On the **Create a new failover relationship** page, in the **Relationship Name** text box, enter **Adatum**.

10. In the **Maximum Client Lead Time** field, in the **hours** text box, type **0**, and then in the **minutes** text box, type **15**.

11. Ensure that the **Mode** field is set to **Load balance**.

12. Ensure that the **Load Balance Percentage** is set to **50%**.

13. Select the **State Switchover Interval** check box. Do not change the default value of 60 minutes.

14. In the **Enable Message Authentication Shared Secret** field, type **Pa\$\$w0rd**, and then click **Next**.

15. Click **Finish**, and then click **Close**.

16. In the **lon-svr1.adatum.com** node, click **IPv4**, and then press F5. Verify that both scopes are listed.

17. Expand the **IPv4** node, and then expand **Scope [172.16.0.0] Adatum**.

18. Click **Address Pool**, and notice that the address pool is configured.

19. Click **Scope Options**, and notice that the scope options are configured.

► **Task 4: Test DHCP failover**

1. Switch to **LON-CL1**.

2. On **LON-CL1**, right-click **Start**, and then click **Network Connections**.

3. In the **Network Connections** window, right-click **Ethernet**, and then click **Properties**.

4. In the **Ethernet Properties** dialog box, click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.

5. In the **Internet Protocol Version 4 (TCP/IPv4) Properties** dialog box, select the **Obtain an IP address automatically** option, select the **Obtain DNS server address automatically** option, and then click **OK**.

6. In the **Ethernet Properties** dialog box, click **Close**.

7. Switch to the command prompt.
8. At the command prompt, type **ipconfig /all**, and then press Enter. Record your IP address and the IP address of your DHCP server.
9. Switch to **LON-SVR1**.
10. If the DHCP server address recorded in step 8 is **172.16.0.10**, perform this step. On **LON-SVR1**, in the **DHCP** console, right-click **lon-dc1**, point to **All tasks**, and then click **Stop**.
11. If the DHCP server address recorded in step 8 is **172.16.0.11**, perform this step. On **LON-SVR1**, in the **DHCP** console, right-click **lon-svr1.adatum.com**, point to **All tasks**, and then click **Stop**.
12. Switch to **LON-CL1**.
13. On **LON-CL1**, at the command prompt, type the following three commands, pressing Enter after each command:

```
ipconfig /release
ipconfig /renew
ipconfig /all
```
14. Record your IP address and the IP address of your DHCP server. The DHCP server IP address should have changed, but the IP address of **LON-CL1** should be the same.
15. Switch to **LON-SVR1**.
16. If the DHCP Server address recorded in step 8 is **172.16.0.10**, perform this step. On **LON-SVR1**, right-click **lon-dc1**, point to **All tasks**, and then click **Start**.
17. If the DHCP server address recorded in step 8 is **172.16.0.11**, perform this step. On **LON-SVR1**, right-click **lon-svr1.adatum.com**, point to **All tasks**, and then click **Start**.

Results: After completing this exercise, you should have:

- Installed an additional DHCP server.
- Created a new DHCP scope.
- Configured DHCP failover.
- Verified DHCP failover.

Exercise 3: Configuring IPAM

► Task 1: Install the IPAM feature

1. Switch to **LON-SVR2**.
2. On **LON-SVR2**, click **Start**, and then click **Server Manager**.
3. In the **Server Manager** console, click **Add roles and features**.
4. In the **Add Roles and Features Wizard**, 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, click **Next**.
8. On the **Select features** page, select the **IP Address Management (IPAM) Server** check box.

9. In the **Add Roles and Features Wizard** dialog box, click **Add Features**, and then click **Next**.
10. On the **Confirm installation selections** page, click **Install**.
11. When the installation finishes, click **Close**.

► **Task 2: Configure IPAM provisioning**

1. On **LON-SVR2**, in the **Server Manager** navigation pane, click **IPAM**.
2. In the **IPAM Overview** pane, click **Connect to IPAM server**, click **LON-SVR2.Adatum.com**, and then click **OK**.
3. In the **Server Manager** content pane, click **Provision the IPAM server**.
4. In the **Provision IPAM Wizard**, on the **Before You Begin** page, click **Next**.
5. On the **Configure database** page, click **Next**.
6. On the **Select provisioning method** page, ensure that the **Group Policy Based** method is selected. In the **GPO name prefix** text box, type **IPAM**, and then click **Next**.
7. On the **Summary** page, click **Apply**. Provisioning will take a few minutes to complete.
8. When provisioning completes, click **Close**.

► **Task 3: Configure IPAM server discovery**

1. On **LON-SVR2**, on the **IPAM Overview** pane, click **Configure server discovery**.
2. In the **Configure Server Discovery** dialog box, click **Get forests**, in the **Configure Server Discovery** dialog box, click **OK**, and then click **Cancel**. The yellow status bar indicates when discovery is complete.
3. In the **IPAM Overview** content pane, click **Configure server discovery**.
4. Click **Add** to add the Adatum.com domain, and then click **OK**.
5. In the **IPAM Overview** content pane, click **Start server discovery**. Discovery might take 5 to 10 minutes to run. The yellow status bar indicates when discovery is complete.

► **Task 4: Configure managed servers**

1. On **LON-SVR2**, in the **IPAM Overview** pane, click **Select or add servers to manage and verify IPAM access**. Notice that the **IPAM Access Status** is blocked. If no servers are listed, click the **Server Manager** console **refresh** button.
2. Click **Start**, and then click **Windows PowerShell**.
3. At the Windows PowerShell prompt, type the following command, and then press Enter:

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

4. When you receive a prompt to confirm the action, type **Y**, and then press Enter. The command will take a few minutes to complete.
5. Close Windows PowerShell.
6. In **Server Manager**, in the content pane, right-click **LON-DC1**, and then click **Edit Server**.
7. In the **Add or Edit Server** dialog box, set the **Manageability** status to **Managed**, and then click **OK**.
8. Switch to **LON-DC1**.
9. On **LON-DC1**, click **Start**, and then click **Windows PowerShell**.

10. In the **Windows PowerShell** window, type the following command, and then press Enter:

```
Gpupdate /force
```

11. Wait for the **gpupdate** command to finish.
12. Close Windows PowerShell.
13. Switch to **LON-SVR2**.
14. Repeat steps 6 through 13 for **LON-SVR1**.
15. In the **Server Manager** console, in the content pane, right-click **LON-DC1**, and then click **Refresh Server Access Status**. The yellow status bar indicates when refresh is complete.
16. In the **Server Manager** console, in the content pane, right-click **LON-SVR1**, and then click **Refresh Server Access Status**. The yellow status bar indicates when refresh is complete.
17. After the refresh completes, click the **Server Manager** console **refresh** button. It might take up to 10 minutes for the status to change. If necessary, repeat both refresh tasks, as needed, until a green check mark displays next to **LON-DC1** and **LON-SVR1**, and the **IPAM Access Status** shows **Unblocked** for both servers.
18. In the **Server Inventory** page, right-click **LON-DC1**, and then click **Retrieve ALL Server Data**. This action will take a few minutes to complete. The yellow status bar indicates when data retrieval is complete.
19. In the **Server Inventory** page, right-click **LON-SVR1**, and then click **Retrieve ALL Server Data**. This action will take a few minutes to complete. The yellow status bar indicates when data retrieval is complete.

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

1. On **LON-SVR2**, in the **IPAM** navigation pane, under **MONITOR AND MANAGE**, click **DNS and DHCP Servers**. Verify that both servers have a **Server Availability** status of **Running** before continuing. If one or both servers have a **Server Availability** of something other than **Running**, click the **Server Manager** console **refresh** button.
2. In the content pane, right-click **LON-SVR1.Adatum.com**, and then click **Create DHCP Scope**.
3. In the **Create DHCP Scope** dialog box, in the **Scope Name** text box, type **Building 3 Scope**.
4. In the **Start IP address** text box, type **10.0.0.50**.
5. In the **End IP address** text box, type **10.0.0.100**.
6. In the **Subnet mask** text box, type **255.255.255.0**.
7. In the **Create scope** pane, click **Options**.
8. On the **DHCP Scope Options** page, click **New**.
9. In the **New Configuration** pane, click the **Option** drop-down arrow, and then select **003 Router**.
10. Under **Values**, in the **IP Address** box, type **10.0.0.1**, and then click **Add Configuration**.
11. On the **DHCP Scope Options** page, click **New**.
12. In the **New Configuration** section, click the **Option** drop-down arrow, and then select **006 DNS Servers**.
13. Under **Values**, in the **IP Address** box, type **172.16.0.9**, click **Add Configuration**, and then click **OK**.
14. In the navigation pane, click **DHCP Scopes**.

15. Right-click **Building 3 Scope**, and then click **Configure DHCP Failover**.
16. In the **Configure DHCP Failover Relationship** dialog box, in the **Configuration option** drop-down list, select **Use an existing relationship**.
17. In the **Relationship name** drop-down list, select **Adatum**.
18. Click **OK**.
19. Switch to **LON-DC1**.
20. On **LON-DC1**, switch to the **DHCP** console, if necessary.
21. In the **DHCP** console, expand **lon-dc1.adatum.com**, click **IPv4**, and then verify that **Building 3 Scope** exists. If **Building 3 Scope** is not listed, press F5.

► **Task 6: Import IP addresses into IPAM**

1. Switch to **LON-SVR2**.
2. In the **IPAM** navigation pane, under **IP ADDRESS SPACE**, click **IP Address Inventory**.
3. In the content pane, click **TASKS**, and then click **Import IP Addresses**.
4. In the **Open** dialog box, in the **File name** text box, type **\\LON-SVR1\Labfiles\Mod05**, and then press Enter.
5. Double-click the **IP-addresses.csv** file.
6. In the **Import IP Addresses** dialog box, verify that 30 records were successfully imported, and then click **OK**.
7. In the content pane, verify that the list contains 30 records.

► **Task 7: Manage DHCP scopes in IPAM**

1. On **LON-SVR2**, in the **IPAM** navigation pane, under **MONITOR AND MANAGE**, click **DHCP Scopes**.
2. In the content pane, below **TASKS**, click the arrow to expand the search area.
3. Click **Add criteria**, select the **[006] DNS Servers** check box, and then click **Add**.
4. In the **contains** text box, type **172.16.0.9**.
5. In the list, select the lines **Building 2 Scope** and **Building 3 Scope** that in the **Server Name** column have the value **lon-svr1.adatum.com**.
6. Right-click one of the selected scopes, and then click **Edit DHCP Scope**.
7. In the **Edit DHCP Scope** dialog box, in the navigation pane, click **Options**, and then in the content pane, in the **Configuration action** drop-down list, select **Find and replace**.
8. Click **New**, and then in the **New Configuration** section, in the **Option** drop-down list, select **006 DNS Servers**.
9. In the **Value to find** text box, type **172.16.0.9**, and in the **Replace with** text box, type **172.16.0.10**. Click **Add Configuration**, and then click **OK**.

► **Task 8: Manage DNS zones in IPAM**

1. On **LON-SVR2**, in the **IPAM** navigation pane, under **MONITOR AND MANAGE**, click **DNS Zones**.
2. In the content pane, click **TASKS**, and then click **Retrieve Server Data**. After the refresh completes, click the **Server Manager** console **refresh** button. Verify that the **Zone Status** for the **Adatum.com** zone now is green.
3. In the navigation pane, click the arrow next to **Forward Lookup**, and then click **Adatum.com**. Click an empty space in the content pane.
4. In the content pane, right-click **LON-DC1.Adatum.com**, and then click **Edit DNS zone**.
5. In the **Edit DNS zone** dialog box, in the navigation pane, click **Name Servers**, and then in the **Fully qualified domain name (FQDN)** text box, type **LON-DC2.Adatum.com**, click **Add Record**, and then click **OK**.
6. In the content pane, click the **Current view** drop-down list box, and then select **Resource Records**.
7. Right-click the **www** record, and then click **Delete DNS resource record**.
8. In the **Delete DNS resource record** dialog box, click **OK**.
9. In the **IPAM** navigation pane, click **DNS Zones**.
10. Right-click **Adatum.com**, and then click **Add DNS resource record**.
11. In the **Add DNS resource record** dialog box, click **New**.
12. In the **New resource record** section, in the **Resource record type** drop-down list box, select **A**.
13. In the **Name** text box, type **lon-dc2**.
14. In the **IP address** text box, type **172.16.0.20**, click **Add resource record**, and then click **OK**.
15. Switch to **LON-DC1**.
16. On **LON-DC1**, switch to the **Server Manager** console, click **Tools**, and then click **DNS**.
17. In the **DNS Manager** console, expand **Forward Lookup Zones**, and then click **Adatum.com**.
18. In the content pane, verify that there is both an NS record for **lon-dc2.adatum.com**, an A record for **lon-dc2** and that the A record for **www** is missing.

► **Task 9: Implement IP address tracking**

1. Switch to **LON-SVR2**.
2. On **LON-SVR2**, in the **Server Manager** console, in the **IPAM** navigation pane, click **EVENT CATALOG**.
3. In the lower half of the **IPAM** navigation pane, under **IP Address Tracking**, click **By Host Name**.
4. In the content pane, click **TASKS**, and then click **Retrieve Event Catalog Data** to get the latest events from the DHCP servers. The yellow status bar indicates when discovery is complete.
5. In the **By Host Name** text box, type **LON-CL1**.
6. In the **first date** text box, type yesterday's date in the format m/d/yyyy. For example, *12/31/2016*.
7. In the **second date** text box, type tomorrow's date in the format m/d/yyyy (such as *12/31/2016*), and then click **Search**.
8. Verify that at least one event is shown below the search area.

► **Task 10: Prepare for the next module**

1. On the host computer, start **Hyper-V Manager**.
2. In the **Virtual Machines** list, right-click **20743A-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20743A-LON-SVR1**, **20743A-LON-SVR2**, **20743A-LON-CL1**, and **20743A-TREY-DC1**.

Results: After completing this exercise, you should have:

- Installed and configured the IPAM feature.
- Configured IPAM provisioning.
- Configured IPAM server discovery.
- Configured managed servers.
- Configured and verified a new DHCP scope with IPAM.
- Imported IP addresses into IPAM.
- Managed DHCP scopes in IPAM.
- Managed DNS scopes in IPAM.
- Implemented IP address tracking.

Module 6: Implementing Hyper-V

Lab: Implementing server virtualization with Hyper-V

Exercise 1: Installing the Hyper-V server role

► Task 1: Install the Hyper-V server role

1. On **LON-HOST1**, click **Start** and then click **Server Manager**.
2. On the **Manage** menu, click **Add Roles and Features**.
3. On the **Before you begin** page of the **Add Roles and Features Wizard**, click **Next**.
4. On the **Select installation type** page, select **Role-based or feature-based installation**, and then click **Next**.
5. On the **Select destination server** page, ensure that **LON-HOST1.Adatum.com** is selected, and then click **Next**.
6. On the **Server Roles** page, select **Hyper-V**.
7. In the **Add Roles and Features Wizard**, click **Add Features**.
8. On the **Select Server Roles** page of the **Add Roles and Features Wizard**, click **Next**.
9. On the **Select features** page, click **Next**.
10. On the **Hyper-V** page, click **Next**.
11. On the **Create Virtual Switches** page, verify that no selections have been made, and then click **Next**.
12. On the **Virtual Machine Migration** page, click **Next**.
13. On the **Default Stores** page, review the location of **Default Stores**, and then click **Next**.
14. On the **Confirm Installation Selections** page, select **Restart the destination server automatically if required**.
15. In the **Add Roles and Features Wizard**, review the message about automatic restarts, and then click **Yes**.
16. On the **Confirm Installation Selections** page, click **Install**.
17. After a few minutes, the server will automatically restart. Ensure that you restart the machine by using the **Boot** menu and then selecting **20743A-LON-HOST1**. The computer will restart several times.

► Task 2: Complete Hyper-V role installation and verify settings

1. Sign in to **LON-HOST1** as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. Click **Start** and then click **Server Manager**.
3. When the installation of Hyper-V is complete, click **Close** to close the **Add Roles and Features Wizard**.
4. Click the **Tools** menu, and then click **Hyper-V Manager**.
5. In Hyper-V Manager, click the Hyper-V host server name (**LON-HOST1**).
6. In the **Actions** pane, click **Hyper-V Settings**.

7. In the **Hyper-V Settings** dialog box, click the **Keyboard** item. Verify that the **Keyboard** is set to use the **Use on the virtual machine** option.
8. In the **Hyper-V Settings** dialog box, click **Virtual Hard Disks**. Verify that the location of the default folder is configured to use the Virtual Hard Disks folder, and then click **OK**.

Results: After completing this exercise, you should have deployed the Hyper-V role to a physical server.

Exercise 2: Configuring virtual networking

► Task 1: Configure the external network

1. On **LON-HOST1** in Hyper-V Manager, in the **Actions** pane, click **Virtual Switch Manager**.
2. In the **Virtual Switch Manager** dialog box, select **New virtual network switch**. Ensure that **External** is selected, and then click **Create Virtual Switch**.
3. In the **Virtual Switch Properties** area of the **Virtual Switch Manager** dialog box, specify the following information, and then click **OK**:
 - o Name: **Corporate Network**
 - o Connection Type: **External Network**: Mapped to the host computer's physical network adapter. This will vary depending on the host computer.
4. In the **Apply Networking Changes** dialog box, review the warning, and then click **Yes**.

► Task 2: Create a private network

1. On **LON-HOST1** in Hyper-V Manager, in the **Actions** pane, click **Virtual Switch Manager**.
2. Under **Virtual Switches**, select **New virtual network switch**.
3. Under **Create virtual switch**, select **Private**, and then click **Create Virtual Switch**.
4. In the **Virtual Switch Properties** section, configure the following settings, and then click **OK**:
 - o Name: **Private Network**
 - o Connection type: **Private network**

► Task 3: Create an internal network

1. On **LON-HOST1** in Hyper-V Manager, in the **Actions** pane, click **Virtual Switch Manager**.
2. Under **Virtual Switches**, select **New virtual network switch**.
3. Under **Create virtual switch**, select **Internal**, and then click **Create Virtual Switch**.
4. In the **Virtual Switch Properties** section, configure the following settings, and then click **OK**:
 - o Name: **Internal Network**
 - o Connection type: **Internal network**

► Task 4: Configure network settings on LON-HOST1

1. Open **Server Manager**, and then click **Local Server**.
2. In the **Properties** pane, next to **vEthernet (Internal Network)** click the **IPv4 address assigned by DHCP** link.

3. In the **Network Connections** dialog box, right-click the **vEthernet (Internal Network)** object, and then click **Properties**.
4. In the **vEthernet (Internal Network) Properties** dialog box, click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
5. On the **General** tab, click **Use the following IP address**, and then configure the following:
 - o IP Address: **172.16.0.31**
 - o Subnet mask: **255.255.0.0**
 - o Default gateway: **172.16.0.1**
6. On the **General** tab, click **Use the following DNS server addresses**, and then configure the following:
 - o Preferred DNS server: **172.16.0.10**
7. To close the **Properties** dialog box, click **OK**.
8. If prompted in the **Microsoft TCP/IP** dialog box, click **Yes**.
9. Click **Close**.
10. In the **Network Connections** box, right-click **Ethernet 2** and then click **Disable**.
11. Close the **Network Connections** dialog box.

Results: After completing this exercise, you should have configured virtual switch options on a physically deployed Windows Server 2016 server that is running the Hyper-V role.

Exercise 3: Creating and configuring a virtual machine

► Task 1: Import virtual machines



Note: Perform the following steps only on LON-HOST1.

1. On **LON-HOST1**, click the **Start** button, right-click **Windows PowerShell ISE**, point to **More**, and then click **Run as administrator**.
2. On the **File** menu, click **Open**, go to **E:\Program Files\Microsoft Learning\20743\Drives**, and then open **LON-HOST1_VM-Pre-Import-20743A.ps1**.



Note: The drive letter might depend on the number of drives on the physical host machine.

3. Select the **LON-HOST1_VM-Pre-Import-20743A.ps1** tab, click **File**, and then click **Run**.
4. At the **On which disk drive are the base images extracted** prompt, type the letter that represents the host computer's **Base** folder, and then press Enter.
5. At the **On which disk drive are the course virtual machines extracted** prompt, type the letter that represents the host computer's **20743** folder, and then press Enter.
6. Press Enter to complete the script.

7. Close the Windows PowerShell ISE.



Note: The drive letter might depend on the number of drives on the physical host machine.

► Task 2: Configure virtual machine storage

1. On **LON-HOST1** on the taskbar, click **File Explorer**, click **This PC**, and then go to **E:\Program Files\Microsoft Learning\20743\Drives**.



Note: The drive letter might depend on the number of drives on the physical host machine.

2. Verify that the **20743A-BASE.vhd** hard disk image file is present.
3. Click the **Home** tab, and then click the **New Folder** icon twice to create two new folders. Right-click each folder and rename them:
 - **LON-GUEST1**
 - **LON-GUEST2**
4. Close File Explorer.
5. Switch to **Hyper-V Manager**.
6. In the **Actions** pane, click **New**, and then click **Hard Disk**.
7. On the **Before You Begin** page of the **New Virtual Hard Disk Wizard**, click **Next**.
8. On the **Choose Disk Format** page, select **VHD**, and then click **Next**.
9. On the **Choose Disk Type** page, select **Differencing**, and then click **Next**.
10. On the **Specify Name and Location** page, specify the following details, and then click **Next**:
 - Name: **LON-GUEST1.vhd**
 - Location: **E:\Program Files\Microsoft Learning\20743\Drives\LON-GUEST1**
11. On the **Configure Disk** page, type the location: **E:\Program Files\Microsoft Learning\20743\Drives\20743A-BASE.vhd**, and then click **Finish**.
12. Click the **Start** button, and then click the **Windows PowerShell** icon.
13. At the Windows PowerShell command prompt, type the following command to create a new differencing disk to use with LON-GUEST2, and then press Enter:


```
New-VHD "E:\Program Files\Microsoft Learning\20743\Drives\LON-GUEST2\LON-GUEST2.vhd"
-ParentPath "E:\Program Files\Microsoft Learning\20743\Drives\20743A-BASE.vhd"
```
14. Close the Windows PowerShell window.
15. In the **Actions** pane of Hyper-V Manager, click **Inspect Disk**.
16. In the **Open** dialog box, go to **E:\Program Files\Microsoft Learning\20743\Drives\LON-GUEST2**, click **LON-GUEST2.vhd**, and then click **Open**.
17. In the **Virtual Hard Disk Properties** dialog box, verify that **LON-GUEST2.vhd** is configured as a differencing virtual hard disk with **E:\Program Files\Microsoft Learning\20743\Drives\20743A-BASE.vhd** as a parent, and then click **Close**.

► Task 3: Create virtual machines

1. On **LON-HOST1** in Hyper-V Manager, in the **Actions** pane, click **New**, and then click **Virtual Machine**.
2. On the **Before You Begin** page of the **New Virtual Machine Wizard**, click **Next**.
3. On the **Specify Name and Location** page of the **New Virtual Machine Wizard**, select **Store the virtual machine in a different location**, enter the following values, and then click **Next**:
 - o Name: **LON-GUEST1**
 - o Location: **E:\Program Files\Microsoft Learning\20743\Drives\LON-GUEST1**
4. On the **Specify Generation** page, click **Next**.
5. On the **Assign Memory** page of the **New Virtual Machine Wizard**, enter a value of **1024 MB**, select the **Use Dynamic Memory for this virtual machine** option, and then click **Next**.
6. On the **Configure Networking** page of the **New Virtual Machine Wizard**, select **Private Network**, and then click **Next**.
7. On the **Connect Virtual Hard Disk** page, choose **Use an existing virtual hard disk**.
8. Click **Browse**, and then go to **E:\Program Files\Microsoft Learning\20743\Drives\LON-GUEST1\lon-guest1.vhd**.
9. Click **Open**, and then click **Finish**.
10. Click the **Start** button, and then click the **Windows PowerShell** icon.
11. At the Windows PowerShell command prompt, enter the following command to create a new virtual machine named **LON-GUEST2**:

```
New-VM -Name LON-GUEST2 -MemoryStartupBytes 1024MB -VHDPATH "E:\Program Files\Microsoft Learning\20743\Drives\LON-GUEST2\LON-GUEST2.vhd" -SwitchName "Private Network"
```

12. Close the **Windows PowerShell** window.
13. In Hyper-V Manager, click **LON-GUEST2**.
14. In the **Actions** pane, under **LON-GUEST2**, click **Settings**.
15. In the **Settings for the LON-GUEST2** dialog box, click **Automatic Start Action**, and then set the **Automatic Start Action** setting to **Nothing**.
16. In the **Settings for the LON-GUEST2** dialog box, click **Automatic Stop Action**, and then set the **Automatic Stop Action** setting to **Shut down the guest operating system**.
17. To close the **Settings for the LON-GUEST2** dialog box, click **OK**.

► Task 4: Configure virtual LANs (VLANs) and network bandwidth settings



Note: Perform the following steps only on LON-HOST1.

1. On **LON-HOST1**, in Hyper-V Manager, start **20743A-LON-DC1-B**.
2. In Hyper-V Manager, in the **Actions** pane, click **Virtual Switch Manager**.
3. Click **Internal Network**, and then select the **Enable virtual LAN identification for management operating system** check box.
4. In the **VLAN ID** text box, type **4**, and then click **OK**.

5. Click **LON-GUEST2**, and then click **Settings**.
6. In settings for **LON-GUEST2**, click **Network Adapter**, change the **Virtual switch** to **Internal Network**, and then click **Enable virtual LAN identification**.
7. In the **VLAN identifier** text box, type **4**, and then click **OK**.
8. In Hyper-V Manager, right-click **LON-GUEST2**, and then click **Start**.



Note: You should see the machine start up and enter a **running** state.

9. Right-click **LON-GUEST2** and click **Connect**.
10. Click **Next**, on the next page, click **Do this later**, and then click **Accept**.
11. Set the password as **Pa\$\$w0rd**, and click **Finish**.
12. Sign in to **LON-GUEST2** as **Administrator** with the password **Pa\$\$w0rd**.
13. Click the **Start** button, type **cmd**, and then press Enter to open the command prompt.
14. At the command prompt, ping **172.16.0.10** by running the following command:

```
Ping 172.16.0.10
```



Note: You should not be able to reach the machine. You will receive a **General failure** message.

15. Repeat steps 5-7 on **20743A-LON-DC1-B**.
16. Switch to **LON-GUEST2**.
17. At the command prompt, run the following command:

```
Ipconfig /renew
```

18. At the command prompt, ping **172.16.0.10** again from **LON-GUEST2** by running the following command:

```
Ping 172.16.0.10
```



Note: You should now see the message, **Reply from 172.16.0.10**.

► Task 5: Configure virtual machine static memory

1. On **LON-GUEST2**, open **Task Manager**, click **More Details**, and then from the **Performance** tab, confirm that there is currently only 1 GB of memory.
2. In Hyper-V Manager, right-click **LON-GUEST2**, and then click **Settings**.
3. Click **Memory**, change the **RAM** to **2048**, and then click **OK**.
4. Return to the still-connected **LON-GUEST2**, note that the Task Manager has increased the memory, and then close **Task Manager**.

► **Task 6: Configure and test virtual machine checkpoints**

1. On **LON-GUEST2**, right-click the desktop, click **New**, and then click **Folder**. In the folder name box, type **Sydney**, and then press Enter.
2. Repeat step 1, and create a second folder named **Melbourne**.
3. Repeat step 1, and create a third folder named **Brisbane**.
4. On the **Action** menu of the **Virtual Machine Connection** window, click **Checkpoint**.
5. In the text box, type **Before Change**, click **Yes**, and then click **OK**.
6. Drag the **Sydney** and **Brisbane** folders to the **Recycle Bin**.
7. Right-click the **Recycle Bin**, and then click **Empty Recycle Bin**.
8. In the **Delete Multiple Items** dialog box, click **Yes**.
9. On the **Action** menu of the **Virtual Machine Connection** window, click **Revert**.
10. In the **Revert Virtual Machine** dialog box, click **Revert**.
11. On the **Action** menu of the **Virtual Machine Connection** window, click **Start**.
12. Right-click **LON-GUEST2**, and then click **Connect**.
13. Sign in with the user name as **Administrator** and the password as **Pa\$\$w0rd**.
14. Verify that the following folders are present on the desktop:
 - **Sydney**
 - **Melbourne**
 - **Brisbane**

Results: After completing this exercise, you should have deployed two separate virtual machines on a virtual hard disk file to act as a parent disk for two differencing disks. The System Preparation Tool (Sysprep) has generalized these virtual machines. You also should have imported a specially prepared virtual machine.

► **Task 7: Prepare for the next module**

- When you finish the lab, shut down **LON-GUEST2**. If you are continuing to the next module, keep **LON-DC1** running. Otherwise you can also shut down **LON-DC1**.

MCT USE ONLY. STUDENT USE PROHIBITED

Module 7: Configuring advanced networking features

Lab: Configuring advanced Hyper-V networking features

Exercise 1: Creating and using Hyper-V virtual switches

► Task 1: Verify the current Hyper-V network configuration

1. On **LON-HOST1**, click **Start**, **All Apps**, scroll down to **Windows Administrative Tools**, and then click **Hyper-V Manager**.
2. In Hyper-V Manager, in the **Actions** pane, click **Virtual Switch Manager**.
3. In the **Virtual Switch Manager for LON-HOST1** window, note the **Corporate Network**, **Private Network**, and **Internal Network** switches that have been created for **LON-HOST1**.
4. Select **Internal Network** and then remove the check mark next to **Enable virtual LAN identification for management operating system**. Click **OK**.
5. Right-click **20743A-LON-DC1-B** and then click **Settings**.
6. In the **Settings for 20743A-LON-DC1-B on LON-HOST1** dialog box, select **Network Adapter**.
7. In the details pane, under **VLAN ID**, remove the check mark next to **Enable virtual LAN identification**.
8. Click **OK**.
9. Repeat steps 5-8 for **LON-GUEST2**.

► Task 2: Create virtual network adapters

1. On **LON-HOST1**, click **Start**, and then click **Windows PowerShell**.
2. At the **Windows PowerShell** command prompt, type the following commands, and then press Enter after each line:

```
Add-VMNetworkAdapter -VMName LON-GUEST2 -Name "Network Adapter 2"
```

```
Connect-VMNetworkAdapter -VMName LON-GUEST2 -Name "Network Adapter 2" -SwitchName "Corporate Network"
```

```
Add-VMNetworkAdapter -VMName LON-GUEST2 -Name "Network Adapter 3"
```

```
Connect-VMNetworkAdapter -VMName LON-GUEST2 -Name "Network Adapter 3" -SwitchName "Corporate Network"
```

► Task 3: Use the Hyper-V virtual switches

1. On **LON-HOST1**, open **Hyper-V Manager**.
2. In Hyper-V Manager, in the **Virtual Machines** pane, find and then right-click **LON-GUEST2**, and then click **Settings**.
3. In the **Settings for LON-GUEST2 on LON-HOST1** window, in the console tree, select the **Network Adapter 2**.

4. Note that the virtual switch assigned is **Corporate Network**.
5. In the **Settings** window, click **Cancel**.
6. In the **Hyper-V Manager** console, start and then connect to **LON-GUEST2**.
7. Sign in as **Administrator** with the password **Pa\$\$w0rd**.
8. In the **Server Manager** console tree, select the **Local Server** node.
9. Notice the network adapters; **Ethernet**, **Ethernet 2**, and **Ethernet 3**, which are configured to use DHCP to obtain IP address information.

► Task 4: Add NIC Teaming

1. On **LON-GUEST2** in the **Server Manager** console tree, select the **Local Server** node.
2. In the **Properties details** pane, next to **NIC Teaming**, click the **Disabled** hyperlink.
3. In the **NIC Teaming** dialog box, in the **Adapters and Interfaces** pane, select **Ethernet 2**, click **Tasks** and then click **Add to new team**.
4. In the **NIC Teaming** dialog box, in the **Team name** box, type **LON-GUEST2 NIC Team**, select **Ethernet 2** and **Ethernet 3**, and then click **OK**.
5. In the **NIC Teaming** dialog box, in the **Teams** pane, note the following:
 - Team: **LON-GUEST2 NIC Team**
 - Status: **OK** (This may show **fault** depending upon if you have external connectivity)
 - Teaming Mode: **Switch Independent**
 - Load Balancing: **Address Hash**
 - Adapters: **2**

Results: After completing this exercise, you should have successfully configured the Hyper-V virtual switch.

Exercise 2: Configuring and using the advanced features of a virtual switch

► Task 1: Configure DHCP Guard

1. On **LON-HOST1**, open **Hyper-V Manager**.
2. In Hyper-V Manager, in the **Virtual Machines** pane, select and right-click **LON-GUEST1**, and then click **Settings**.
3. In the **Settings for LON-GUEST1 on LON-HOST1** window, in the console tree, select and then expand **Network Adapter**.
4. Under **Network Adapter**, click **Advanced Features**.
5. In the details pane, in the **DHCP Guard** area, click **Enable DHCP Guard**, and then click **OK**.
6. Repeat steps 2-5 for **LON-GUEST2**.

► **Task 2: Configure and use bandwidth management**

1. While still on **LON-HOST1**, in Hyper-V Manager, in the **Virtual Machines** pane, find and then right-click **LON-GUEST2**, and then click **Settings**.
2. In the **Settings for LON-GUEST2 on LON-HOST1** window, in the console tree, select **Network Adapter 2**.
3. In the details pane, in the **Bandwidth Management** area, select **Enable bandwidth management**.
4. In the **Maximum bandwidth** box, type **100**, and then click **OK**.

► **Task 3: Prepare for the next module**

1. Shut down the following virtual machines:
 - **LON-GUEST2**
 - **20743A-LON-DC1-B**
2. Restart the host computer and then at the boot menu select **Windows Server 2012**.

Results: After completing this exercise, you should have successfully configured the advanced features of the Hyper-V virtual switch.

MCT USE ONLY. STUDENT USE PROHIBITED

Module 8: Implementing Software Defined Networking

Lab: Deploying Network Controller

Exercise 1: Preparing to deploy Network Controller

► Task 1: Create the required AD DS security groups

1. Switch to **LON-DC1**.
2. In **Server Manager**, click **Tools**, and then click **Active Directory Users and Computers**.
3. In **Active Directory Users and Computers**, expand **Adatum.com**, and then click **IT**.
4. Right-click **IT**, click **New**, and then click **Group**.
5. In the **New Object – Group** dialog box, in the **Group name** box, type **Network Controller Admins**, and then click **OK**.
6. In the details pane, double-click **Network Controller Admins**, and in the **Network Controller Admins Properties** dialog box, on the **Members** tab, click **Add**.
7. In the **Select Users, Contacts, Computers, Service Accounts, or Groups** dialog box, in the **Enter the object names to select (examples)** box, type **administrator; Beth**, and then click **OK** twice.
8. Right-click **IT**, click **New**, and then click **Group**.
9. In the **New Object – Group** dialog box, in the **Group name** box, type **Network Controller Ops**, and then click **OK**.
10. In the details pane, double-click **Network Controller Ops**, and in the **Network Controller Ops Properties** dialog box, on the **Members** tab, click **Add**.
11. In the **Select Users, Contacts, Computers, Service Accounts, or Groups** dialog box, in the **Enter the object names to select (examples)** box, type **administrator; Beth**, and then click **OK** twice.
12. Close **Active Directory Users and Computers**.

► Task 2: Request a certificate for authenticating Network Controller

1. Switch to **LON-SVR2**.
2. Right-click **Start**, and then click **Run**.
3. In the **Run** dialog box, type **mmc.exe**, and then press Enter.
4. In the **Console1 – [Console Root]** window, click **File**, and then click **Add/Remove Snap-in**.
5. In the **Add or Remove Snap-ins** dialog box, in the **Snap-in** list, double-click **Certificates**.
6. Click **Computer account**, click **Next**, and then click **Finish**.
7. Click **OK**.
8. In the navigation pane, expand **Certificates**, and then expand **Personal**.
9. Right-click **Personal**, click **All Tasks**, and then click **Request New Certificate**.
10. In the **Certificate Enrollment** dialog box, on the **Before you Begin** page, click **Next**.
11. On the **Select Certificate Enrollment Policy** page, click **Next**.

12. Select the **Computer** check box, and then click **Enroll**.
13. Click **Finish**.
14. Close the management console and do not save changes.

Results: After completing this exercise, you should have successfully prepared your environment for Network Controller.

Exercise 2: Deploying Network Controller

► Task 1: Add the Network Controller role

1. On **LON-SVR2**, click **Start**, and then click **Server Manager**.
2. In **Server Manager**, in the details pane, click **Add roles and features**.
3. In the **Add Roles and Features Wizard**, on the **Before you begin** page, click **Next**.
4. On the **Select installation type** page, click **Next**.
5. On the **Select destination server** page, click **Next**.
6. On the **Select server roles** page, in the **Roles** list, select the **Network Controller** check box.
7. Click **Add Features**, and then click **Next**.
8. On the **Select features** page, click **Next**.
9. On the **Network Controller** page, click **Next**.
10. On the **Confirm installation selections** page, click **Install**.
11. When the role is installed, click **Close**.
12. Right-click **Start**, point to **Shut down or sign out**, and then click **Restart**.
13. In the **Choose a reason that best describes why you want to shut down this computer** dialog box, click **Continue**.
14. After **LON-SVR2** has restarted, sign in as **Adatum\administrator** with the password as **Pa\$\$w0rd**.

► Task 2: Configure the Network Controller cluster



Note: These steps are duplicated in the high level steps for this lab.

1. On **LON-SVR2**, right-click **Start**, and then click **Windows PowerShell (Admin)**.
2. At the Windows PowerShell (Admin) command prompt, type the following command, and then press Enter:

```
$node=New-NetworkControllerNodeObject -Name "Node1" -Server "LON-SVR2.Adatum.com" -
FaultDomain "fd:/rack1/host1" -RestInterface "Ethernet"
```

3. At the Windows PowerShell (Admin) command prompt, type the following command, and then press Enter:

```
$Certificate = Get-Item Cert:\LocalMachine\My | Get-ChildItem | where {$_.Subject -
imatch "LON-SVR2" }
```

- At the Windows PowerShell (Admin) command prompt, type the following command, and then press Enter:

```
Install-NetworkControllerCluster -Node $node -ClusterAuthentication Kerberos -
ManagementSecurityGroup "Adatum\Network Controller Admins" -
CredentialEncryptionCertificate $Certificate
```

► Task 3: Configure the Network Controller application



Note: This step is duplicated in the high level steps for this lab.

- On **LON-SVR2**, at the Windows PowerShell (Admin) command prompt, type the following command, and then press Enter:

```
Install-NetworkController -Node $node -ClientAuthentication Kerberos -
ClientSecurityGroup "Adatum\Network Controller Ops" -RestIpAddress "172.16.0.99/24" -
ServerCertificate $Certificate
```



Note: This command can take quite a while to complete.

► Task 4: Verify the deployment



Note: These steps are duplicated in the high level steps for this lab.

- On **LON-SVR2**, at the Windows PowerShell (Admin) command prompt, type the following command, and then press Enter:

```
$cred=New-Object Microsoft.Windows.Networkcontroller.credentialproperties
```

- At the Windows PowerShell (Admin) command prompt, type the following command, and then press Enter:

```
$cred.type="usernamepassword"
```

- At the Windows PowerShell (Admin) command prompt, type the following command, and then press Enter:

```
$cred.username="admin"
```

- At the Windows PowerShell (Admin) command prompt, type the following command, and then press Enter:

```
$cred.value="abcd"
```

- At the Windows PowerShell (Admin) command prompt, type the following command, and then press Enter:

```
New-NetworkControllerCredential -ConnectionUri https://LON-SVR2.Adatum.com -
Properties $cred -ResourceId cred1
```

- Press **Y**, and then press Enter when prompted.

7. At the Windows PowerShell (Admin) command prompt, type the following command, and then press Enter:

```
Get-NetworkControllerCredential -ConnectionUri https://LON-SVR2.Adatum.com -
ResourceId cred1
```

8. You should receive output that looks similar to the output below:

```
Tags                :
ResourceRef         : /credentials/cred1
CreatedTime         : 1/1/0001 12:00:00 AM
InstanceId          : e16ffe62-a701-4d31-915e-7234d4bc5a18
Etag                : W/"1ec59631-607f-4d3e-ac78-94b0822f3a9d"
ResourceMetadata    :
ResourceId          : cred1
Properties           : Microsoft.Windows.NetworkController.CredentialProperties
```

► Task 5: Prepare for the next module

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

1. On the host computer, start **Hyper-V Manager**.
2. In the **Virtual Machines** list, right-click **20743A-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20743A-LON-SVR2**.

Results: After completing this exercise, you should have successfully deployed Network Controller.

Module 9: Implementing remote access

Lab: Implementing DirectAccess

Exercise 1: Configuring DirectAccess using the Getting Started Wizard

► Task 1: Run the Getting Started Wizard on LON-RTR

1. Switch to **LON-RTR**.
2. Open **Server Manager**, click **Tools**, and then select **Remote Access Management**.
3. In the Remote Access Management console, under **Configuration**, click **DirectAccess and VPN**, and then click **Run the Getting Started Wizard**.
4. On the **Configure Remote Access** page, click **Deploy DirectAccess only**.
5. Verify that **Edge** is selected, and in the **Type the public name or IPv4 address used by clients to connect to the Remote Access server** text box, type **131.107.0.200**, and then click **Next**.
6. On the **Configure Remote Access** page, click the **here** link.
7. On the **Remote Access Review** page, verify that two Group Policy Objects (GPOs) have been created:
 - DirectAccess Server Settings
 - DirectAccess Client Settings
8. Next to **Remote Clients**, click **Change**.
9. In the **Remote Access Setup** window, click **Domain Computers (ADATUM\Domain Computers)**, click **Remove**, and then click **Add**.
10. In the **Select Groups** window, type **DA_Clients**, and then click **OK**.
11. Clear the **Enable DirectAccess for mobile computers only** check box, and click **Next**.
12. On the **DirectAccess Client Setup** page, in the **DirectAccess connection name** text box, add **Windows 10** to the existing phrase so that it says **Windows 10 Workplace Connection**, and then click **Finish**.
13. On the **Remote Access Review** page, click **OK**.
14. On the **Configure Remote Access** page, to finish the **DirectAccess** wizard, click **Finish**.
15. In the **Applying Getting Started Wizard Settings** dialog box, click **Close**.
16. Right-click **Start**, and then click **Windows PowerShell (Admin)**.
17. At the Windows PowerShell command prompt, type the following cmdlet, and then press Enter:

```
Restart-Computer
```

18. Sign in to **LON-RTR** as **Adatum\Administrator** with the password of **Pa\$\$w0rd**

► Task 2: Verify that the client is configured

1. Switch to **LON-CL1**.
2. When you configured the DirectAccess server, the wizard created two Group Policies and linked them to the domain. To apply them, restart **LON-CL1**, and then sign in as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
3. Right-click the **Start** button and then click **Command Prompt**.

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

```
gpupdate /force
```

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

```
gpresult /R
```

- Look through the resulting output that the command generated. Under the **Computer Settings** section, verify that the DirectAccess Client Settings GPO is applied.



Note: If the DirectAccess Client Settings GPO is not applied, restart **LON-CL1**, and then repeat steps 2 through 5 on **LON-CL1**.

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

```
netsh name show effectivepolicy
```

- Verify that the "DNS Effective Name Resolution Policy Table Settings" message appears.
- On **LON-CL1**, open **Internet Explorer**.
- In the Address bar, type **http://lon-svr1.adatum.com**, and then press Enter. Verify that the default IIS web page appears.
- Leave the **Internet Explorer** window open.
- Open **File Explorer**, type **\\LON-SVR1\DA**, and then press Enter. Note that you are able to access the folder.

► Task 3: Move LON-CL1 to the external network

- On **LON-CL1**, right-click the **Start** button, and click **Network Connections**.
- Right-click **Ethernet**, and click **Disable**.
- Right-click **Ethernet 2**, and click **Enable**.
- Right-click **Ethernet 2**, and click **Properties**.
- In the **Ethernet 2 Properties** dialog box, double-click **Internet Protocol Version 4 (TCP/IPv4)**.
- In the **Internet Protocol Version 4 (TCP/IPv4) Properties** dialog box, configure the following, and then click **OK**.
 - IP address: **131.107.0.2**
 - Subnet mask: **255.255.0.0**
 - Preferred DNS server: **131.107.0.100**
- In the **Ethernet 2 Properties** dialog box, click **OK**.

Results: After completing this exercise, you will have successfully deployed DirectAccess.

Exercise 2: Testing DirectAccess

► Task 1: Verify connectivity to the internal network resources

1. On **LON-CL1**, type the following command at the command prompt, and then press Enter:

```
netsh name show effectivepolicy
```

2. Verify that DNS Effective Name Resolution Policy Table Settings present two entries, one for .Adatum.com and one for Directaccess-NLS.Adatum.com.
3. On **LON-CL1**, switch to Internet Explorer.
4. In the Internet Explorer address bar, type **http://lon-SVR1.adatum.com**, and then press Enter. Verify that the default IIS web page appears.
5. Leave the **Internet Explorer** window open.
6. On the Start screen, type **\\LON-SVR1\DA**, and then press Enter. Note that you are able to access the folder content.



Note: When you verify the DirectAccess deployment, you might not be able to connect to the internal file share or the internal Web site since DirectAccess is not fully functional in Windows Server 2016 TP5. This issue will be resolved in Windows Server 2016 RTM.

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

```
ipconfig
```



Note: Notice the IP address for Tunnel adapter is **IPHTTPSInterface** starting with **2002**. This is an IP-HTTPS address.

► Task 2: Verify connectivity to the DirectAccess server

1. On **LON-CL1**, at the command prompt, type the following command, and then press Enter:

```
Powershell
```

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

```
Get-DAClientExperienceConfiguration
```



Note: Review the DirectAccess client settings.

► Task 3: Verify client connectivity to the DirectAccess server

1. Switch to **LON-RTR** and open **Server Manager**.
2. Click **Tools** and then click **Remote Access Management**.
3. In the console pane, click **Remote Client Status**.



Note: Notice that a client is connected via **IPHttps**. In the Connection Details pane, in the lower right of the screen, note the use of Kerberos authentication for the machine and the user.

4. Close all open windows.

► Task 4: Prepare for the next module

When you have finished the lab, revert the virtual machines to their initial state.

1. On the host computer, start **Hyper-V Manager**.
2. In the **Virtual Machines** list, right-click **20743A-LON-CL1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20743A-LON-DC1**, **20743A-LON-RTR**, **20743A-INET1**, and **20743A-LON-SVR1**.

Results: After completing this exercise, you will have successfully verified your DirectAccess deployment.

Module 11: Implementing failover clustering

Lab: Implementing failover clustering

Exercise 1: Configuring iSCSI storage

► Task 1: Configure the iSCSI targets

1. On **LON-SVR1** click on the taskbar, click the **Windows** button, and then click **Server Manager**.
2. In **Server Manager**, in the navigation pane, click **File and Storage Services**.
3. In the **File and Storage Services** pane, click **iSCSI**.
4. In the **iSCSI VIRTUAL DISKS** pane, click **TASKS**, and then in the **TASKS** drop-down list box, click **New iSCSI Virtual Disk**.
5. In the **New iSCSI Virtual Disk Wizard**, on the **Select iSCSI virtual disk location** page, under **Storage location**, click **C:**, and then click **Next**.
6. On the **Specify iSCSI virtual disk name** page, in the **Name** box, type **iSCSIDisk1**, and then click **Next**.
7. On the **Specify iSCSI virtual disk size** page, in the **Size** box, type **5**, ensure that **GB** is selected in the drop-down list box, and then click **Next**.
8. On the **Assign iSCSI target** page, click **New iSCSI target**, and then click **Next**.
9. On the **Specify target name** page, in the **Name** box, type **lon-svr1**, and then click **Next**.
10. On the **Specify access servers** page, click **Add**.
11. In the **Select a method to identify the initiator** dialog box, click **Enter a value for the selected type**, and then in the **Type** drop-down list box, click **IP Address**. In the **Value** box, type **172.16.0.12**, and then click **OK**.
12. On the **Specify access servers** page, click **Add**.
13. In the **Select a method to identify the initiator** dialog box, click **Enter a value for the selected type**, and then in the **Type** drop-down list box, select **IP Address**. In the **Value** box, type **172.16.0.13**, and then click **OK**.
14. On the **Specify access servers** page, click **Next**.
15. On the **Enable Authentication** page, click **Next**.
16. On the **Confirm selections** page, click **Create**.
17. On the **View results** page, wait until the creation completes, and then click **Close**.
18. In the **iSCSI VIRTUAL DISKS** pane, click **TASKS**, and then in the **TASKS** drop-down list box, select **New iSCSI Virtual Disk**.
19. In the **New iSCSI Virtual Disk Wizard**, on the **Select iSCSI virtual disk location** page, under **Storage location**, click **C:**, and then click **Next**.
20. On the **Specify iSCSI virtual disk name** page, in the **Name** box, type **iSCSIDisk2**, and then click **Next**.
21. On the **Specify iSCSI virtual disk size** page, in the **Size** box, type **5**, ensure that **GB** is selected in the drop-down list box, and then click **Next**.
22. On the **Assign iSCSI target** page, click **lon-svr1**, and then click **Next**.

23. On the **Confirm selections** page, click **Create**.
24. On the **View results** page, wait until the creation completes, and then click **Close**.
25. In the **iSCSI VIRTUAL DISKS** pane, click **TASKS**, and then in the **TASKS** drop-down list box, select **New iSCSI Virtual Disk**.
26. In the **New iSCSI Virtual Disk Wizard**, on the **Select iSCSI virtual disk location** page, under **Storage location**, click **C:**, and then click **Next**.
27. On the **Specify iSCSI virtual disk name** page, in the **Name** box, type **iSCSIDisk3**, and then click **Next**.
28. On the **Specify iSCSI virtual disk size** page, in the **Size** box, type **5**, ensure that **GB** is selected in the drop-down list box, and then click **Next**.
29. On the **Assign iSCSI target** page, click **lon-svr1**, and then click **Next**.
30. On the **Confirm selections** page, click **Create**.
31. On the **View results** page, wait until the creation completes, and then click **Close**.

Results: After completing this exercise, you should have successfully installed an iSCSI Target Server.

Exercise 2: Configuring a failover cluster

► Task 1: Connect clients to the iSCSI targets

1. On **LON-SVR2**, open **Server Manager**, click **Tools**, and then click **iSCSI Initiator**.
2. In the **Microsoft iSCSI** dialog box, click **Yes**.
3. In the **iSCSI Initiator** window, click the **Discovery** tab.
4. Click **Discover Portal**.
5. In the **IP address or DNS name** box, type **172.16.0.11**, and then click **OK**.
6. Click the **Targets** tab, and then click **Refresh**.
7. In the **Targets** list, click **iqn.1991-05.com.microsoft:lon-svr1-lon-svr1-target**, and then click **Connect**.
8. Click **Add this connection to the list of Favorite Targets**, and then click **OK** two times.
9. On **LON-SVR3**, open **Server Manager**, click **Tools**, and then click **iSCSI Initiator**.
10. In the **Microsoft iSCSI** dialog box, click **Yes**.
11. In the **iSCSI Initiator** window, click the **Discovery** tab.
12. Click **Discover Portal**.
13. In the **IP address or DNS name** box, type **172.16.0.11**, and then click **OK**.
14. Click the **Targets** tab, and then click **Refresh**.
15. In the **Targets** list, click **iqn.1991-05.com.microsoft:lon-svr1-lon-svr1-target**, and then click **Connect**.
16. Click **Add this connection to the list of Favorite Targets**, and then click **OK** two times.
17. On **LON-SVR2**, in **Server Manager**, click **Tools**, and then click **Computer Management**.

18. Expand **Storage**, and then click **Disk Management**.
19. Right-click **Disk 3**, and then click **Online**.
20. Right-click **Disk 3**, and then click **Initialize disk**.
21. In the **Initialize Disk** dialog box, click **OK**.
22. Right-click the unallocated space next to **Disk 3**, and then click **New Simple Volume**.
23. On the **Welcome** page, click **Next**.
24. On the **Specify Volume Size** page, click **Next**.
25. On the **Assign Drive Letter or Path** page, click **Next**.
26. On the **Format Partition** page, in the **Volume Label** box, type **Data**. Select the **Perform a quick format** check box, and then click **Next**.
27. Click **Finish**.



Note: If a window appears with a prompt to format the disk, click **Cancel**.

28. Repeat steps 19 through 27 for **Disk 4** and **Disk 5**, using **Data2** and **Data3**, respectively, for the volume labels.
29. Close the **Computer Management** window.
30. On **LON-SVR3**, in **Server Manager**, click **Tools**, and then click **Computer Management**.
31. Expand **Storage**, and then click **Disk Management**.
32. Select and right-click **Disk Management**, and then click **Refresh**.
33. Right-click **Disk 3**, and then click **Online**.
34. Right-click **Disk 4**, and then click **Online**.
35. Right-click **Disk 5**, and then click **Online**.
36. Close the **Computer Management** window.

► Task 2: Install the Failover Clustering feature

1. On **LON-SVR2**, if **Server Manager** is not already open, click **Start** and then click the **Server Manager** icon.
2. 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, in the **Features** list, click **Failover Clustering**. In the **Add features that are required for Failover Clustering** window, click **Add Features**, and then click **Next**.
8. On the **Confirm installation selections** page, click **Install**.

9. When the installation completes and you receive the **Installation succeeded on LON-SVR2.Adatum.com** message, click **Close**.
10. On **LON-SVR3**, repeat steps 1 through 8. When the installation completes and you receive the **Installation succeeded on LON-SVR3.Adatum.com** message, click **Close**.

► **Task 3: Validate and create a failover cluster**

1. On **LON-SVR2**, in **Server Manager**, click **Tools**, and then click **Failover Cluster Manager**.
2. In the **Failover Cluster Manager** console, in the **Actions** pane, click **Validate Configuration**.
3. In the **Validate a Configuration Wizard**, click **Next**.
4. In the **Enter Name** box, type **LON-SVR2**, and then click **Add**.
5. In the **Enter Name** box, type **LON-SVR3**.
6. Click **Add**, and then click **Next**.
7. Verify that **Run all tests (recommended)** is selected, and then click **Next**.
8. On the **Confirmation** page, click **Next**.
9. Wait for the validation tests to finish (it might take from 5 through 7 minutes).
10. Verify that all the tests completed without errors. Some warnings are expected.
11. Select the check box next to **Create the cluster now using the validated nodes**.
12. On the **Summary** page, click **Finish**.
13. On the **Before You Begin** page, click **Next**.
14. On the **Access Point for Administering the Cluster** page, in the **Cluster Name** box, type **Cluster1**.
15. Under **Address**, type **172.16.0.125**, and then click **Next**.
16. On the **Confirmation** page, click **Next**.
17. On the **Summary** page, click **Finish**.

Results: After this exercise, you should have installed and configured the Failover Clustering feature.

Exercise 3: Deploying and configuring a highly available file server

► **Task 1: Add the file server application to the failover cluster**

1. On **LON-SVR2**, in the **Failover Cluster Manager** console, expand **Cluster1.Adatum.com**, expand **Storage**, and then click **Disks**.
2. Ensure that three disks are present and online (with the names **Cluster Disk 1**, **Cluster Disk 2**, and **Cluster Disk 3**).
3. Right-click **Roles**, and then click **Configure Role**.
4. On the **Before You Begin** page, click **Next**.
5. On the **Select Role** page, click **File Server**, and then click **Next**.
6. On the **File Server Type** page, click **File Server for general use**, and then click **Next**.

7. On the **Client Access Point** page, in the **Name** box, type **AdatumFS**, in the **Address** box, type **172.16.0.130**, and then click **Next**.
8. On the **Select Storage** page, select the **Cluster Disk 2** check box, and then click **Next**.
9. On the **Confirmation** page, click **Next**.
10. On the **Summary** page, click **Finish**.

► **Task 2: Add a shared folder to a highly available file server**

1. On **LON-SVR3**, in the **Server Manager** console, click **Tools**, and then click **Failover Cluster Manager**.
2. Expand **Cluster1.Adatum.com**, click **Roles**, right-click **AdatumFS**, and then click **Add File Share**.
3. In the **New Share Wizard**, on the **Select the profile for this share** page, click **SMB Share - Quick**, and then click **Next**.
4. On the **Select the server and the path for this share** page, click **Next**.
5. On the **Specify share name** page, in the **Share name** box, type **Docs**, and then click **Next**.
6. On the **Configure share settings** page, review the available options but do not make any changes, and then click **Next**.
7. On the **Specify permissions to control access** page, click **Next**.
8. On the **Confirm selections** page, click **Create**.
9. On the **View results** page, click **Close**.

► **Task 3: Configure the failover and failback settings**

1. On **LON-SVR3**, in the **Failover Cluster Manager** console, click **Roles**, right-click **AdatumFS**, and then click **Properties**.
2. In the **AdatumFS Properties** dialog box, click the **Failover** tab, and then click **Allow failback**.
3. Click **Failback between**, and then set the values to **4** and **5** hours.
4. Click the **General** tab.
5. Select both **LON-SVR2** and **LON-SVR3** as the preferred owners.
6. Move **LON-SVR3** up so that it is first in the preferred owners list.
7. To close the **AdatumFS Properties** dialog box, click **OK**.

Results: After this exercise, you should have configured a highly available file server.

Exercise 4: Validating the deployment of the highly available file server

► **Task 1: Validate the highly available file server deployment**

1. On **LON-DC1**, open **File Explorer**, type **\\AdatumFS** in the address bar, and then press Enter.
2. Verify that you can access the location and that you can open the **Docs** folder.
3. Create a test text document inside this folder.
4. On **LON-SVR2**, switch to the **Failover Cluster Manager**.

5. In the **Failover Cluster Manager** console, expand **Cluster1.adatum.com**, and then click **Roles**.
6. In the **Owner Node** column, note the current owner of **AdatumFS**.



Note: The owner is either **LON-SVR2** or **LON-SVR3**.

7. Right-click **AdatumFS**, click **Move**, and then click **Select Node**.
8. In the **Move Clustered Role** dialog box, select the cluster node (it is either **LON-SVR2** or **LON-SVR3**), and then click **OK**.
9. Verify that **AdatumFS** has moved to a new owner.
10. Switch to **LON-DC1**.
11. To verify that you can still access the **\\AdatumFS** location, open **File Explorer**, type **\\AdatumFS** in the address bar, and then press Enter.

► **Task 2: Validate the failover and quorum configuration for the File Server role**

1. On **LON-SVR2**, in the **Failover Cluster Manager** console, click **Roles**.
2. In the **Owner Node** column, verify the current owner for the **AdatumFS** role.



Note: The owner is either **LON-SVR2** or **LON-SVR3**.

3. Click **Nodes**, and then select the node that is the current owner of the **AdatumFS** role.
4. Right-click the node, click **More Actions**, and then click **Stop Cluster Service**.
5. In the **Failover Cluster Manager** console, click **Roles**, and then verify that **AdatumFS** is running.



Note: This confirms that **AdatumFS** has moved to another node.

6. Switch to **LON-DC1**.
7. On **LON-DC1**, to verify that you can still access the **\\AdatumFS** location, open **File Explorer**, type **\\AdatumFS** in the address bar, and then press Enter.
8. Switch to **LON-SVR2**.
9. In the **Failover Cluster Manager** console, click **Nodes**, right-click the stopped node, click **More Actions**, and then click **Start Cluster Service**.
10. Expand **Storage**, and then click **Disks**.
11. In the **Disks** pane, find the disk that is assigned to **Disk Witness in Quorum**.



Note: You can view this in the **Assigned To** column.

12. Right-click the disk, click **Take Offline**, and then click **Yes**.
13. Switch to **LON-DC1**.
14. On **LON-DC1**, to verify that you can still access the **\\AdatumFS** location, open **File Explorer**, type **\\AdatumFS** in the address bar, and then press Enter.



Note: This verifies that the cluster is still running even if the witness disk is offline.

15. Switch to **LON-SVR2**.
16. In the **Failover Cluster Manager** console, expand **Storage**, click **Disks**, right-click the disk that is in **Offline** status, and then click **Bring Online**.
17. Right-click **Cluster1.Adatum.com**, click **More Actions**, and then click **Configure Cluster Quorum Settings**.
18. On the **Before You Begin** page, click **Next**.
19. On the **Select Quorum Configuration Option** page, click **Advanced quorum configuration**, and then click **Next**.
20. On the **Select Voting Configuration** page, review the available settings but do not make any changes.



Note: Notice that you can select a node or nodes that will or will not have a vote in the cluster.

21. Click **Next**.
22. On the **Select Quorum Witness** page, ensure that **Configure a disk witness** is selected, and then click **Next**.
23. On the **Configure Storage Witness** page, click **Cluster Disk 3**, and then click **Next**.
24. On the **Confirmation** page, click **Next**.
25. On the **Summary** page, click **Finish**.

Results: After this exercise, you should have tested the failover scenarios.

Exercise 5: Configuring CAU on the failover cluster

► 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, ensure that **Select a 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**.
7. In the **Add features that are required for Failover Clustering** dialog box, click **Add Features**, and then click **Next**.
8. On the **Confirm installation selections** page, click **Install**.
9. When the installation completes, click **Close**.

10. Switch to **LON-SVR2**.
11. Open **Server Manager**.
12. In **Server Manager**, click **Tools**, and then click **Windows Firewall with Advanced Security**.
13. In the **Windows Firewall with Advanced Security** window, click **Inbound Rules**.
14. In the **Rules** list, find the **Inbound Rule for Remote Shutdown (RPC-EP-In)** rule. Ensure that the rule is enabled. Right-click the rule, and then click **Enable Rule**.
15. In the **Rules** list, find the **Inbound Rule for Remote Shutdown (TCP-In)** rule. Right-click the rule, and then click **Enable Rule**.
16. Close the **Windows Firewall with Advanced Security** window.
17. Switch to **LON-SVR3**, and then repeat steps 11 through 16.
18. On **LON-DC1**, in **Server Manager**, click **Tools**, and then click **Cluster-Aware Updating**.
19. In the **Cluster-Aware Updating** window, in the **Connect to a failover cluster** drop-down list box, click **CLUSTER1**, and then click **Connect**.
20. In the **Cluster Actions** pane, click **Preview updates for this cluster**.
21. In the **Cluster1-Preview Updates** window, click **Generate Update Preview List**.
22. After a minute or two, when the updates display in the list, review the updates, and then click **Close**. (Note that there may be no updates currently. Updates will show as available.)

► **Task 2: Update the failover cluster and configure self-updating**

1. On **LON-DC1**, in the **Cluster-Aware Updating** console, click **Apply updates to this cluster**.
2. On the **Getting Started** page, click **Next**.
3. On the **Advanced options** page, review the options for updating, and then click **Next**.
4. On the **Additional Update Options** page, click **Next**.
5. On the **Confirmation** page, click **Update**, and then click **Close**.
6. In the **Cluster nodes** pane, review the progress of the updating.
7. Wait until the process finishes.



Note: This updating process might require a restart of both nodes. The process is finished when both nodes display **Succeeded** in the **Last Run status** column.

8. On **LON-SVR2**, in **Server Manager**, click **Tools**, and then click **Cluster-Aware Updating**.
9. In the **Cluster-Aware Updating** dialog box, in the **Connect to a failover cluster** drop-down list box, click **CLUSTER1**, and then click **Connect**.
10. In the **Cluster Actions** pane, click **Configure cluster self-updating options**.
11. On the **Getting Started** page, click **Next**.
12. 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**.
13. On the **Specify self-updating schedule** page, click **Weekly**. In the **Time of day** drop-down list box, click **4:00 AM**. In the **Day of the week** drop-down list box, click **Sunday**, and then click **Next**.
14. On the **Advanced Options** page, click **Next**.

15. On the **Additional Update Options** page, click **Next**.
16. On the **Confirmation** page, click **Apply**.
17. After the clustered role is successfully added, click **Close**.

► **Task 3: Prepare for the next module**

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

1. On the host computer, start **Hyper-V Manager**.
2. In the **Virtual Machines** list, right-click **20743A-LON-DC1**, and then click **Revert**.
3. In the **Revert VM** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20743A-LON-SVR1**, **20743A-LON-SVR2**, **20743A-LON-SVR3**, and **MSL-TMG1**.

Results: After this exercise, you should have configured CAU.

MCT USE ONLY. STUDENT USE PROHIBITED

Module 12: Implementing failover clustering with Windows Server 2016 Hyper-V

Lab: Implementing failover clustering with Windows Server 2016 Hyper-V

Exercise 1: The Hyper-V Failover clustering testing environment

► Task 1: Enabling Nested Virtualization

1. On **LON-HOST2**, click **Start**, and then click **Windows PowerShell**.
2. In the **Windows PowerShell** window, type the following command, and then press Enter:

```
Install-WindowsFeature -Name Hyper-V,Hyper-V-Tools,Hyper-V-PowerShell -Restart
```

Your computer will restart a couple of times. Be sure to select **20743A-LON-HOST2** at the boot menu.

3. Sign into **LON-HOST2** as **Adatum\Administrator** with **Pa\$\$w0rd**.
4. Open **File Explorer** and then browse to **E:\Program Files\Microsoft Learning\20743\Drives**. (Note that the drive letter may vary based upon your host machine).
5. Run the following scripts in order to create the needed switches and import the virtual machines that need to be imported for this lab: **CreateVirtualSwitches.ps1** **LON-HOST2_VM-Pre-Import-20743A.ps1**.
6. Start Windows PowerShell, and then run the following commands to enable nested virtualization on **20743A-LON-NVHOST3** and **20743A-LON-NVHOST4**.

```
Set-VMProcessor -VMName 20743A-LON-NVHOST3 -ExposeVirtualizationExtensions $true -
Count 2
Set-VMemory 20743A-LON-NVHOST3 -DynamicMemoryEnabled $false
Get-VMNetworkAdapter -VMName 20743A-LON-NVHOST3 | Set-VMNetworkAdapter -
MacAddressSpoofing On
Set-VMProcessor -VMName 20743A-LON-NVHOST4 -ExposeVirtualizationExtensions $true -
Count 2
Set-VMemory 20743A-LON-NVHOST4 -DynamicMemoryEnabled $false
Get-VMNetworkAdapter -VMName 20743A-LON-NVHOST4 | Set-VMNetworkAdapter -
MacAddressSpoofing On
```

► Task 2: Upgrade the configuration version and install Hyper-V on nested virtualization hosts

1. Open **Server Manager**, click **Tools**, and then click **Hyper-V Manager**.
2. In the left pane, select **LON-HOST2**.
3. In the **Virtual Machines** pane, select **20743A-LON-NVHOST3**.
4. In the **Actions** menu, click **Upgrade Configuration Version** and then click **Upgrade**.
5. In the **Actions** menu, click **Start**.
6. Repeat steps 3–5 for **20743-LON-NVHOST4**, and **20743A-LON-DC1-C**.
7. Click **Start**, right-click **Windows PowerShell**, and then click **Run as Administrator**.

8. To open a PSSession utilizing Windows PowerShell Direct, at the Windows PowerShell command prompt, type the following command, and then press Enter:

```
Enter-PSSession -VMName 20743A-LON-NVHOST3
```

9. In the **Credentials request** dialog box, in the **Username** text box, type **Adatum\Administrator**, and in the **Password** text box, type **Pa\$\$w0rd**, and then press Enter.
10. To install Hyper-V, type the following command at the Windows PowerShell command prompt, and then press Enter:

```
Install-WindowsFeature -Name Hyper-V,Hyper-V-Tools,Hyper-V-Powershell -Restart
```

11. Type the following command at the Windows PowerShell command prompt, and then press Enter:

```
Exit
```



Note: If this generates the following error “Command ‘Exit’ was not run as the session in which it was intended to run was either closed or broken” Ignore this error.

12. Repeat steps 8–11 for **20743A-LON-NVHOST4**.
13. To configure networking, open a Windows PowerShell prompt and then type the following command: (Note that the path may differ on your host machine.)

```
Invoke-Command -VMName 20743A-LON-NVHOST3 -FilePath "D:\Program Files\Microsoft Learning\20743\Drives\CreateVirtualSwitches.ps1"
```

14. In the **Credentials request** dialog box, in the **Username** text box, type **Adatum\Administrator**, and in the **Password** text box, type **Pa\$\$w0rd**, and then press Enter.
15. Repeat steps 13–14 for **20743A-LON-NVHOST4**.
16. Switch to the Hyper-V Manager console, right-click **20743A-LON-NVHOST3**, click **Connect**, and then sign in as **Adatum\Administrator** with the password of **Pa\$\$w0rd**.
17. Repeat step 16 for **20743A-LON-NVHOST4**.

Results: After completing this exercise, you should have successfully enabled nested virtualization on **20743A-LON-NVHOST3** and **20743A-LON-NVHOST4**.

Exercise 2: Configuring Hyper-V Replica

► Task 1: Import the LON-CORE virtual machine on LON-NVHOST3

1. On **LON-NVHOST3**, open **Server Manager**, and then open the **Hyper-V Manager** console.
2. Click **LON-NVHOST3**, and then in the **Actions** pane, click **Import Virtual Machine**.
3. In the **Import Virtual Machine Wizard**, on the **Before You Begin** page, click **Next**.
4. On the **Locate Folder** page, click **Browse**.
5. Browse to folder **C:\Program Files\Microsoft Learning\20743\Drives\20743A-LON-CORE**, click **Select Folder**, and then click **Next**.



Note: The drive letter might be different based upon the number of drives on the physical host machine.

6. On the **Select Virtual Machine** page, click **20743A-LON-CORE**, and then click **Next**.
7. On the **Choose Import Type** page, click **Next**.
8. On the **Summary** page, click **Finish**.

► **Task 2: Configure a replica on both host machines**

1. On **LON-NVHOST4**, click **Start**, click the **Server Manager** icon, click **Tools**, and then click **Hyper-V Manager**.
2. In Hyper-V Manager, right-click **LON-NVHOST4**, and then click **Hyper-V Settings**.
3. In the **Hyper-V Settings for LON-NVHOST4** window, click **Replication Configuration**.
4. In the **Replication Configuration** pane, click **Enable this computer as a Replica server**.
5. In the **Authentication and ports** section, click **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 (C)**, and then click **New folder**.
8. In the **Folder Name** text box, type **VMReplica**, and then press Enter.
9. Select the **C:\VMReplica** folder, and then click **Select Folder**. (The drive letter might be different based upon the number of drives on the physical host machine.)
10. In the **Hyper-V Settings for LON-VNHOST4** window, click **OK**.
11. In the **Settings** window, read the notice, and then click **OK**.
12. Click **Start**, and then click the **Control Panel** icon.
13. In the **Control Panel** window, click **System and Security**, and then click **Windows Firewall**.
14. In the **Windows Firewall** window, click **Advanced settings**.
15. Click the **Inbound Rules** setting.
16. In the right pane, in the list, find and right-click the **Hyper-V Replica HTTP Listener (TCP-In)** rule, and then click **Enable Rule**.
17. Close the **Windows Firewall with the Advanced Security** console, and then close the **Windows Firewall** window.
18. Switch to **LON-NVHOST3**, and then repeat steps 1 through 17.

► **Task 3: Configure replication for the LON-CORE virtual machine**

1. On **LON-NVHOST3**, on the taskbar, click the **Server Manager** icon, click **Tools**, and then click **Hyper-V Manager**.
2. In Hyper-V Manager, click **LON-NVHOST3**, right-click **20743A-LON-CORE**, and then click **Enable Replication**.
3. On the **Before You Begin** page, click **Next**.
4. On the **Specify Replica Server** page, click **Browse**.
5. In the **Select Computer** window, type **LON-NVHOST4**, click **Check Names**, click **OK**, and then click **Next**.

6. On the **Specify Connection Parameters** page, review the settings, ensure that **Use Kerberos authentication (HTTP)** is selected, and then click **Next**.
7. On the **Choose Replication VHDs** page, ensure that **20743A-LON-CORE.vhd** is selected, and then click **Next**.
8. On the **Configure Replication Frequency** page, in the drop-down list box, click **30 seconds**, and then click **Next**.
9. On the **Configure Additional Recovery Points** page, click **Maintain only the latest recovery point**, and then click **Next**.
10. On the **Choose Initial Replication Method** page, click **Send initial copy over the network**, click **Start replication immediately**, and then click **Next**.
11. On the **Completing the Enable Replication Wizard** page, click **Finish**.
12. Wait five to ten minutes. You can monitor the progress of initial replication in the **Status** column in the **Hyper-V Manager** console.
13. When initial replication completes (progress reaches 100 percent), ensure that **20743A-LON-CORE** now displays on **LON-NVHOST4** in the **Hyper-V Manager** console.

► **Task 4: Validate a planned failover to the replica site**

1. On **LON-NVHOST4**, in Hyper-V Manager, right-click **20743A-LON-CORE**, click **Replication**, and then click **View Replication Health**.
2. Review the content of the window that appears, ensure that there are no errors, and then click **Close**.
3. Switch to **LON-NVHOST3**.
4. On **LON-NVHOST3**, on the taskbar, click the **Server Manager** icon, click **Tools**, click **Hyper-V Manager**, and then verify that **20743A-LON-CORE** is turned off.
5. Right-click **20743A-LON-CORE**, click **Replication**, and then click **Planned Failover**.
6. In the **Planned Failover** window, select the **Reverse the replication direction after failover** option, ensure that **Start the Replica virtual machine after failover** is selected, and then click **Fail Over**.
7. In the **Planned Failover** window, click **Close**.
8. On **LON-NVHOST4**, in the **Hyper-V Manager** console, ensure that **20743A-LON-CORE** is running.
9. On **LON-NVHOST3**, right-click **20743A-LON-CORE**, point to **Replication**, and then click **Remove replication**.
10. In the **Remove replication** dialog box, click **Remove Replication**.
11. On **LON-NVHOST4**, right-click **20743A-LON-CORE**, and then click **Shut Down**.
12. In the **Shut Down Machine** dialog box, click **Shut Down**.

Results: After completing this exercise, you should have successfully configured Hyper-V Replica.

Exercise 3: Configuring a failover cluster for Hyper-V

► Task 1: Connect to the iSCSI target from both host machines

1. On **LON-NVHOST3**, click **Start**, click the **Server Manager** icon, click **Tools**, and then click **iSCSI Initiator**.
2. At the **Microsoft iSCSI** prompt, click **Yes**.
3. Click the **Discovery** tab.
4. On the **Discovery** tab, click **Discover Portal**.
5. In the **IP address or DNS name** text box, type **172.16.0.10**, and then click **OK**.
6. Click the **Targets** tab, and then click **Refresh**.
7. In the **Discovered targets** list, click **iqn.1991-05.com.microsoft:lon-dc1-target1-target**, and then click **Connect**.
8. Click **Add this connection to the list of Favorite Targets**, and then click **OK**.
9. To close **iSCSI Initiator Properties** dialog box, click **OK**.
10. Switch to **LON-NVHOST4**.
11. On **LON-NVHOST4**, open **Server Manager**, click **Tools**, and then click **iSCSI Initiator**.
12. In the **Microsoft iSCSI** dialog box, click **Yes**.
13. In the **iSCSI Initiator** dialog box, click the **Discovery** tab.
14. On the **Discovery** tab, click **Discover Portal**.
15. In the **IP address or DNS name** text box, type **172.16.0.10**, and then click **OK**.
16. Click the **Targets** tab, and then click **Refresh**.
17. In the **Discovered targets** list, click **iqn.1991-05.com.microsoft:lon-dc1-target1-target**, and then click **Connect**.
18. Click **Add this connection to the list of Favorite Targets**, and then click **OK**.
19. To close the **iSCSI Initiator Properties** dialog box, click **OK**.
20. On **LON-NVHOST4**, in **Server Manager**, click **Tools**, and then click **Computer Management**.
21. Expand **Storage**, and then click **Disk Management**.
22. Right-click **Disk 1**, and then click **Online**. (This is the first disk that is 20 GB in size. Your disk number might be different.)
23. Right-click **Disk 1**, and then click **Initialize Disk**.
24. In the **Initialize Disk** dialog box, click **OK**.
25. Right-click the unallocated space next to **Disk 1**, and then click **New Simple Volume**.
26. On the **Welcome** page, click **Next**.
27. On the **Specify Volume Size** page, click **Next**.
28. On the **Assign Drive Letter or Path** page, click **Next**.
29. On the **Format Partition** page, in the **Volume label** text box, type **ClusterDisk**.
30. Select the **Perform a quick format** check box, click **Next**, and then click **Finish**.

31. Repeat steps 22 through 30 for Disk 2 and Disk 3. In step 29, use the following settings:
 - o Disk 2 name: **ClusterVMs**
 - o Disk 3 name: **Quorum**
32. Switch back to **LON-NVHOST3**.
33. On **LON-NVHOST3**, in **Server Manager**, click **Tools**, and then click **Computer Management**.
34. Expand **Storage**, and then click **Disk Management**.
35. Right-click **Disk Management**, and then click **Refresh**.
36. Right-click **Disk 1**, and then click **Online**.
37. Right-click **Disk 2**, and then click **Online**.
38. Right-click **Disk 3**, and then click **Online**.



Note: Disk numbers might vary based on the number of physical disks in the host computer. Choose the disks that are 20 GB in size.

► Task 2: Configure failover clustering on both host machines

1. On **LON-NVHOST3**, click **Start**, and then click the **Server Manager** icon.
2. In **Server Manager**, on the **Dashboard**, click **Add roles and features**.
3. On the **Before you begin** page, click **Next**.
4. On the **Select installation type** page, click **Next**.
5. On the **Select destination server** page, ensure that **Select server from the server pool** is selected, and then click **Next**.
6. On the **Select server roles** page, click **Next**.
7. On the **Select features** page, in the **Features** list, click **Failover Clustering**. At the **Add features that are required for failover clustering** prompt, click **Add Features**, and then click **Next**.
8. On the **Confirm installation selections** page, click **Install**.
9. When installation is complete, click **Close**.
10. Switch to **LON-NVHOST4**.
11. Repeat steps 1 through 9 on **LON-NVHOST4**.
12. Switch back to **LON-NVHOST3**.
13. On **LON-NVHOST3**, in **Server Manager**, click **Tools**, and then click **Failover Cluster Manager**.
14. In Failover Cluster Manager, in the center pane, under **Management**, click **Create Cluster**.
15. On the **Before You Begin** page of the **Create Cluster Wizard**, read the information, and then click **Next**.
16. On the **Select Servers** page, in the **Enter server name** text box, type **LON-NVHOST3**, and then click **Add**. Again in the **Enter server name** text box, type **LON-NVHOST4**, and then click **Add**.
17. Verify the entries, and then click **Next**.
18. On the **Validation Warning** page, click **No. I don't require support from Microsoft for this cluster**, and then click **Next**.

19. On the **Access Point for Administering the Cluster** page, in the **Cluster Name** text box, type **VMCluster**.
20. In the **Address** text box, type **172.16.0.126**, and then click **Next**.
21. In the **Confirmation** dialog box, verify the information, clear the **Add all eligible storage to the cluster** option check box, and then click **Next**.
22. On the **Summary** page, click **Finish**.

► **Task 3: Configure disks for a failover cluster**

1. On **LON-NVHOST3**, in the **Failover Cluster Manager** console, expand **VMCluster.Adatum.com**, expand **Storage**, right-click **Disks**, and then click **Add Disk**.
2. In the **Add Disks to Cluster** dialog box, verify that all disks are selected, and then click **OK**.
3. Verify that all disks display as available for cluster storage in **Failover Cluster Manager**.
4. Click **Cluster Disk 1**, right-click that disk, and then click **Add to Cluster Shared Volumes**.
5. Right-click **VMCluster.adatum.com**, click **More Actions**, click **Configure Cluster Quorum Settings**, and then 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**.

Results: After completing this exercise, you should have successfully configured the failover clustering infrastructure for Hyper-V.

Exercise 4: Configuring a highly available virtual machine

► **Task 1: Move virtual machine storage to the iSCSI target**

1. Ensure that **LON-NVHOST3** is the owner of the disk that you just assigned to Cluster Shared Volume. You can read owner value in the **Owner node** column. If that is not the case, then move the disk to **LON-NVHOST3** before proceeding to step 2.



Note: To move the disk:

- a. Right-click the disk, and then click **Move**.
 - b. Click **Select Node**, click **LON-NVHOST3**, and then click **OK**.
2. On **LON-NVHOST3**, on the desktop, on the taskbar, click the **File Explorer** icon.

3. In File Explorer, expand drive **C:**, expand **Program Files**, expand **Microsoft Learning**, expand **20743**, expand **Drives**, expand **20743A-LON-CORE**, and then click **Virtual Hard Disks**.



Note: The drive letter might be different depending on the physical machine.

4. In the details pane, move the **20743A-LON-CORE.vhd** virtual hard disk file to the **C:\ClusterStorage\Volume1** location.

► **Task 2: Configure the virtual machine as highly available**

1. On **LON-NVHOST3**, in Failover Cluster Manager, click **Roles**, and then in the **Actions** pane, click **Virtual Machines**.
2. Click **New Virtual Machine**.
3. Select **LON-NVHOST3** as the cluster node, and then click **OK**.
4. In the **New Virtual Machine Wizard**, on the **Before You Begin** page, click **Next**.
5. On the **Specify Name and Location** page, in the **Name** text box, type **TestClusterVM**, click **Store the virtual machine in a different location**, and then click **Browse**.
6. Browse to and select **C:\ClusterStorage\Volume1**, click **Select Folder**, and then click **Next**.
7. On the **Specify Generation** page, click **Generation 1**, and then click **Next**.
8. On the **Assign Memory** page, type **512**, and then click **Next**.
9. On the **Configure Networking** page, leave the selection as **Not Connected**, and then click **Next**.
10. On the **Connect Virtual Hard Disk** page, click **Use an existing virtual hard disk**, and then click **Browse**.
11. Browse to **C:\ClusterStorage\Volume1**, click **20743A-LON-CORE.vhd**, and then click **Open**.
12. Click **Next**, and then click **Finish**. If an error appears informing you that Microsoft Management has stopped working, restart this task from step 1.
13. On the **Summary** page of the **High Availability Wizard**, click **Finish**.
14. Right-click the **TestClusterVM**, and then click **Settings**.
15. On **LON-NVHOST3**, in the **Settings for TestClusterVM** dialog box, in the navigation pane, expand **Processor**, and then click **Compatibility**.
16. In the right pane, select the **Migrate to a physical computer with a different processor version** check box, and then click **OK**.
17. Right-click **TestClusterVM**, and then click **Start**.
18. Ensure that the virtual machine starts successfully.

► **Task 3: Perform a Live Migration for the virtual machine**

1. On **LON-NVHOST4**, open **Failover Cluster Manager**.
2. Expand **VMCluster.Adatum.com**, and then click **Roles**.
3. Right-click **TestClusterVM**, click **Move**, click **Live Migration**, and then click **Select Node**.
4. Click **LON-NVHOST4**, and then click **OK**. Wait until the machine is migrated. You will see that the **Owner Node** column will change the value when migration completes.

5. Right-click **TestClusterVM**, and then click **Connect**.
6. Ensure that you can access and operate the virtual machine while it is migrating to another host.

► **Task 4: Perform a Storage Migration for the virtual machine**

1. On **LON-NVHOST4**, on the taskbar, click the **Server Manager** icon, click **Tools**, and then click **Hyper-V Manager**.
2. In Hyper-V Manager, in the central pane, click **20743A-LON-CORE**.
3. In the **Actions** pane, click **Start**, and then wait until the virtual machine is fully started.
4. Switch back to the **Hyper-V Manager** console, and then in the **Actions** pane, click **Move**.
5. On the **Before You Begin** page, click **Next**.
6. On the **Choose Move Type** page, click **Move the virtual machine's storage**, and then click **Next**.
7. On the **Choose Options for Moving Storage** page, click **Move all of the virtual machine's data to a single location**, and then click **Next**.
8. On the **Choose a new location for virtual machine** page, click **Browse**.
9. Browse to **C:**, create a new folder named **LON-CORE**, click **Select Folder**, and then click **Next**.
10. On the **Summary** page, click **Finish**.
11. While waiting for the move process to complete, connect to the virtual machine, and verify that it is fully operational.
12. Shut down all running virtual machines.

Results: After completing this exercise, you should have successfully configured the virtual machine as highly available.

► **Task 5: Prepare for the end of the course**

- Shut down the host computer.

MCT USE ONLY. STUDENT USE PROHIBITED