

Microsoft Official Course



WS-011T00

Windows Server 2019 Administration

WS-011T00
**Windows Server 2019
Administration**

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.

© 2019 Microsoft Corporation. All rights reserved.

Microsoft and the trademarks listed at <http://www.microsoft.com/trademarks>¹ are trademarks of the Microsoft group of companies. All other trademarks are property of their respective owners.

¹ <http://www.microsoft.com/trademarks>

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.

1. "Authorized Learning Center" means a Microsoft Imagine Academy (MSIA) Program Member, Microsoft Learning Competency Member, or such other entity as Microsoft may designate from time to time.
2. "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.
3. "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.
4. "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 an MPN Member (defined below), or (iii) a Microsoft full-time employee, a Microsoft Imagine Academy (MSIA) Program Member, or a Microsoft Learn for Educators – Validated Educator.
5. "Licensed Content" means the content accompanying this agreement which may include the Microsoft Instructor-Led Courseware or Trainer Content.
6. "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.
7. "Microsoft Instructor-Led Courseware" means the Microsoft-branded instructor-led training course that educates IT professionals, developers, students at an academic institution, and other learners on Microsoft technologies. A Microsoft Instructor-Led Courseware title may be branded as MOC, Microsoft Dynamics, or Microsoft Business Group courseware.
8. "Microsoft Imagine Academy (MSIA) Program Member" means an active member of the Microsoft Imagine Academy Program.
9. "Microsoft Learn for Educators – Validated Educator" means an educator who has been validated through the Microsoft Learn for Educators program as an active educator at a college, university, community college, polytechnic or K-12 institution.
10. "Microsoft Learning Competency Member" means an active member of the Microsoft Partner Network program in good standing that currently holds the Learning Competency status.
11. "MOC" means the "Official Microsoft Learning Product" instructor-led courseware known as Microsoft Official Course that educates IT professionals, developers, students at an academic institution, and other learners on Microsoft technologies.
12. "MPN Member" means an active Microsoft Partner Network program member in good standing.

13. "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.
14. "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.
15. "Trainer" means (i) an academically accredited educator engaged by a Microsoft Imagine Academy Program Member to teach an Authorized Training Session, (ii) an academically accredited educator validated as a Microsoft Learn for Educators – Validated Educator, and/or (iii) a MCT.
16. "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.
 1. **If you are a Microsoft Imagine Academy (MSIA) Program Member:**
 1. 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.
 2. 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.
 3. For each license you acquire, you must comply with the following:
 1. you will only provide access to the Licensed Content to those individuals who have acquired a valid license to the Licensed Content,
 2. 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,
 3. 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,

4. 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,
 5. 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,
 6. you will only deliver a maximum of 15 hours of training per week for each Authorized Training Session that uses a MOC title, and
 7. you acknowledge that Trainers that are not MCTs will not have access to all of the trainer resources for the Microsoft Instructor-Led Courseware.
2. **If you are a Microsoft Learning Competency Member:**
1. Each license you acquire 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.
 2. For each license you acquire on behalf of an End User or MCT, 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) MCT with the unique redemption code and instructions on how they can access one (1) Trainer Content.
 3. For each license you acquire, you must comply with the following:
 1. you will only provide access to the Licensed Content to those individuals who have acquired a valid license to the Licensed Content,
 2. 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,
 3. 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,

4. you will ensure that each MCT teaching an Authorized Training Session has their own valid licensed copy of the Trainer Content that is the subject of the Authorized Training Session,
5. 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,
6. you will only provide access to the Microsoft Instructor-Led Courseware to End Users, and
7. you will only provide access to the Trainer Content to MCTs.

3. If you are a MPN Member:

1. 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.
2. 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.
3. For each license you acquire, you must comply with the following:
 1. you will only provide access to the Licensed Content to those individuals who have acquired a valid license to the Licensed Content,
 2. 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,
 3. 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,
 4. 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,

5. 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,
6. 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,
7. you will only provide access to the Microsoft Instructor-Led Courseware to End Users, and
8. you will only provide access to the Trainer Content to Trainers.

4. 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.

5. If you are a Trainer.

1. 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.
 2. If you are an MCT, 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.
 3. 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:
1. **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.
 2. **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.
 3. **Pre-release Term.** If you are an Microsoft Imagine Academy Program Member, Microsoft Learning Competency Member, MPN Member, Microsoft Learn for Educators – Validated Educator, 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 provided "as is", we are not obligated to 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.**
 1. 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.
 2. Outside the United States. If you acquired the Licensed Content in any other country, the laws of that country apply.
12. **LEGAL EFFECT.** This agreement describes certain legal rights. You may have other rights under the laws of your country. You may also have rights with respect to the party from whom you acquired the Licensed Content. This agreement does not change your rights under the laws of your country if the laws of your country do not permit it to do so.
13. **DISCLAIMER OF WARRANTY. THE LICENSED CONTENT IS LICENSED "AS-IS" AND "AS AVAILABLE." YOU BEAR THE RISK OF USING IT. MICROSOFT AND ITS RESPECTIVE AFFILIATES GIVES NO EXPRESS WARRANTIES, GUARANTEES, OR CONDITIONS. YOU MAY HAVE ADDITIONAL CONSUMER RIGHTS UNDER YOUR LOCAL LAWS WHICH THIS AGREEMENT CANNOT CHANGE. TO THE EXTENT PERMITTED UNDER YOUR LOCAL LAWS, MICROSOFT AND ITS RESPECTIVE AFFILIATES EXCLUDES ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT.**
14. **LIMITATION ON AND EXCLUSION OF REMEDIES AND DAMAGES. YOU CAN RECOVER FROM MICROSOFT, ITS RESPECTIVE AFFILIATES AND ITS SUPPLIERS ONLY DIRECT DAMAGES UP TO US\$5.00. YOU CANNOT RECOVER ANY OTHER DAMAGES, INCLUDING CONSEQUENTIAL, LOST PROFITS, SPECIAL, INDIRECT OR INCIDENTAL DAMAGES.**

This limitation applies to

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

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

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

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

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

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

Cette limitation concerne:

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

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

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

Revised April 2019



Contents

■	Module 0 Course introduction	1
	About this course	1
■	Module 1 Windows Server administration	3
	Overview of Windows Server administration principles and tools	3
	Introducing Windows Server 2019	15
	Overview of Windows Server Core	25
	Module 01 lab and review	31
■	Module 2 Identity services in Windows Server	37
	Overview of AD DS	37
	Deploying Windows Server domain controllers	56
	Overview of Azure AD	70
	Implementing Group Policy	81
	Overview of AD CS	93
	Module 02 lab and review	103
■	Module 3 Network infrastructure services in Windows Server	107
	Deploying and managing DHCP	107
	Deploying and managing DNS services	120
	Deploying and managing IPAM	138
	Module 03 lab and review	149
■	Module 4 File servers and storage management in Windows Server	157
	Volumes and file systems in Windows Server	157
	Implementing sharing in Windows Server	168
	Implementing Storage Spaces in Windows Server	175
	Implementing Data Deduplication	188
	Implementing iSCSI	199
	Deploying DFS	207
	Module 04 lab and review	220
■	Module 5 Hyper-V virtualization and containers in Windows Server	227
	Hyper-V in Windows Server	227
	Configuring VMs	236
	Securing virtualization in Windows Server	253
	Containers in Windows Server	260
	Overview of Kubernetes	270

	Module 05 lab and review	273
■	Module 6 High availability in Windows Server	281
	Planning for failover clustering implementation	281
	Creating and configuring failover clusters	297
	Overview of stretch clusters	315
	High availability and disaster recovery solutions with Hyper-V VMs	323
	Module 06 lab and review	329
■	Module 7 Disaster recovery in Windows Server	339
	Hyper-V Replica	339
	Backup and restore infrastructure in Windows Server	350
	Module 07 lab and review	357
■	Module 8 Windows Server security	361
	Credentials and privileged access protection in Windows Server	361
	Hardening Windows Server	374
	Just Enough Administration in Windows Server	382
	Securing and analyzing SMB traffic	390
	Windows Server Update Management	394
	Module 08 lab and review	401
■	Module 9 RDS in Windows Server	411
	Overview of RDS	411
	Configuring a session-based desktop deployment	438
	Overview of personal and pooled virtual desktops	454
	Module 09 lab and review	465
■	Module 10 Remote Access and web services in Windows Server	471
	Overview of RAS in Windows Server	471
	Implementing VPNs	480
	Implementing NPS	492
	Implementing Always On VPN	506
	Implementing Web Server in Windows Server	514
	Module 10 lab and review	527
■	Module 11 Server and performance monitoring in Windows Server	537
	Overview of Windows Server monitoring tools	537
	Using Performance Monitor	547
	Monitoring event logs for troubleshooting	557
	Module 11 lab and review	561
■	Module 12 Upgrade and migration in Windows Server	565
	AD DS migration	565
	Storage Migration Service	570
	Windows Server Migration Tools	576
	Module 12 lab and review	580



Module 0 Course introduction

About this course

About this course

Welcome to the *Windows Server 2019 Administration* course, which is part of a set of three courses that cover Windows Server administration and transitioning to cloud-only and hybrid models. This course focuses on Windows Server 2019 administration, when servers are deployed on-premises. Course *WS-012T: Windows Server 2019 Hybrid and Azure IaaS* covers the use of Microsoft Azure in managing server workloads in hybrid environments and helping secure virtual machines that are running on Windows Server 2019. Course *WS-013T: Azure Stack HCI* covers the running of virtual workloads on Microsoft platforms at medium to large scale by using software-defined datacenters and hyperconverged principles.

Level: Intermediate

Audience

This course is for IT professionals who have some experiencing working with Windows Server and want to learn how to administer Windows Server 2019. The audience for this course also includes current Windows Server administrators who have worked with older Windows Server versions and who want to update their skills in Windows Server 2019. Service-desk professionals who want to transition to server maintenance or pass exams relating to Windows Server also will find this course useful.

Prerequisites

This course assumes you have skills and experience with the following technologies and concepts:

- Active Directory Domain Services (AD DS) in Windows Server 2012 or Windows Server 2016.
- Microsoft Hyper-V and basic server virtualization.
- Windows client operating systems such as Windows 8, Windows 8.1, or Windows 10.
- Windows PowerShell.

- Windows Server 2012 or Windows Server 2016 configuration and maintenance.
- Basic security best practices.
- Core networking technologies such as IP addressing, name resolution, and Dynamic Host Configuration Protocol (DHCP).

Labs and demonstrations

You'll perform labs and the demonstrations on a virtual lab environment from an authorized lab hoster.

Course syllabus

The course content includes a mix of content, demonstrations, hands-on labs, and reference links.

	Module Name
0	Course introduction
1	Windows Server administration
2	Identity services in Windows Server
3	Network infrastructure services in Windows Server
4	File servers and storage management in Windows Server
5	Hyper-V virtualization and containers in Windows Server
6	High availability in Windows Server
7	Disaster recovery in Windows Server
8	Windows Server security
9	RDS in Windows Server
10	Remote access and web services in Windows Server
11	Server and performance monitoring in Windows Server
12	Upgrade and migration in Windows Server

Course resources

There are many resources that can help you learn about Windows Server. We recommend that you bookmark the following websites:

- **Microsoft Learn:**¹ Free role-based learning paths and hands-on experiences for practice.
- **Windows Server documentation:**² Articles and how-to guides about using Windows Server.

¹ <https://aka.ms/Microsoft-learn-home-page>

² <https://aka.ms/windows--server>



Module 1 Windows Server administration

Overview of Windows Server administration principles and tools

Lesson overview

In this lesson, you will learn about Windows Server administration best practices and the tools available for managing Windows Servers. The practice of least privilege has always been the cornerstone of security in the computing environment. Microsoft has supplied many tools and guidelines to allow management of the environment while reducing the exposure of systems and data.

After completing this lesson, you will be able to:

- Describe the concept of least privilege.
- Describe delegated privileges.
- Explain how to delegate privileges.
- Describe privileged access workstations.
- Describe jump servers.
- Describe Windows Server Admin Center.
- Describe Server Manager.
- Describe how to use Remote Server Administration Tools (RSAT).
- Describe how to use PowerShell to manage servers.
- Explain how to manage servers remotely.

Overview of Windows Admin Center

Managing and administrating the IT environment involves using different tools across multiple consoles. Windows Admin Center combines those tools into a single console that can easily be deployed and accessed through a web interface.

Windows Admin Center is a modular web application comprised of the following four modules:

- Server manager. Manages servers that run Windows Server 2008 R2 and newer (limited functionality for 2008 R2). If you want to manage servers other than the local server, you must add those other servers to the console.
- Failover clusters
- Hyper-converged clusters
- Windows 10 clients

Windows Admin Center has two main components:

- Gateway. The Gateway manages servers through remote PowerShell and Windows Management Instrumentation (WMI) over Windows Remote Management (WINRM).
- Web server. The Web server component observes HTTPS requests and serves the user interface to the web browser on the management station. This is not a full install of Internet Information Services (IIS), but a mini Web server for this specific purpose.

Note: Because Windows Admin Center is a web-based tool that uses HTTPS, it requires a X.509 certificate to provide SSL encryption. The installation wizard gives you the option to either use a self-signed certificate or provide your own SSL certificate. This certificate expires 60 days after it is created.

Benefits of Windows Admin Center

The following table describes the benefits of Windows Admin Center:

Table 1: Benefits of Windows Admin Center

Benefit	Description
Familiar functionality	It uses the familiar admin tools from Microsoft Management Consoles.
Easy to install and use	You can download and install it on Windows 10 or Windows Server through a single Windows Installer (MSI) and access it from a web browser.
Compliments existing solutions	It does not replace but compliments existing solutions such as Remote Server Administration Tools, System Center, and Azure Operation Management Suite.
Manage from the internet	It can be securely published to the public internet so you can connect to and manage servers from anywhere.
Enhanced security	Role-based access control lets you fine-tune which administrators have access to which management features. Gateway authentication provides support for local groups, Active Directory groups, and Azure Active Directory groups.
Azure integration	You can easily get to the proper tool within Windows Admin Center, then launch it to the Azure portal for full management of Azure services.

Benefit	Description
Extensibility	A Software Development Kit (SDK) will allow Microsoft and other partners to develop new tools and solutions for more products.
No external dependencies	Windows Admin Center doesn't require internet access or Microsoft Azure. There is no requirement for IIS or SQL server and there are no agents to deploy. The only dependency is to the requirement of Windows Management Framework 5.1 on managed servers.

Supported platforms and browsers

You can install Windows Admin Center on Windows 10 version 1709 or newer, or Windows Server 2016 or newer. Windows Admin Center is not supported on domain controllers and will return an error if you try to install it. The Windows browser versions of Microsoft Edge on Windows 10 and Google Chrome are tested and supported on Windows 10. Other modern web browsers have not been tested and are not officially supported. Internet Explorer is not supported and will return an error if you attempt to launch Windows Admin Center.

Demonstration: Use Windows Server Admin Center

In this demonstration, you will learn how to install the Windows Server Admin Center and add a server for remote administration and perform various administrative tasks.

Demonstration steps

Install Windows Server Admin Center

1. Open File Explorer and browse to **C:\Labfiles\Mod01**.
2. Double-click or select **WindowsAdminCenter1910.2.msi**, and then select Enter. Install Windows Admin Center by accepting all the defaults.

Add servers for remote administration

1. Open Microsoft Edge and go to <https://Sea-Adm1>.
2. In Windows Admin Center, add **SEA-DC1**.
3. To connect to **SEA-DC1**, select the link to SEA-DC1.Contoso.com.

Browse through the various admin sections

- In Windows Server Admin Center browse through the following panes:
 - **Overview**
 - **Certificates**

- **Files**
- **Events**
- **Devices**
- **Performance Monitor**
- **Processes**
- **Roles & Features**
- **Scheduled Tasks**
- **PowerShell**

Server Manager

Server Manager is the built-in management console that most server administrators are familiar with. You can use the current version to manage the local server and remotely manage up to 100 servers. However, this number will depend on the amount of data that you request from managed servers and the hardware and network resources available to the system running **Server Manager**. In the **Server Manager** console, you must manually add remote servers that you want to manage. IT administrators often use **Server Manager** to remotely manage server core installations.

The **Server Manager** console comes with the Remote Server Administration Tools for Windows 10. However, you can only use it to manage remote servers. You can't use **Server Manager** to manage client operating systems.

Server Manager initially opens to a dashboard which provides quick access to:

- Adding roles and features.
- Adding other servers to manage.
- Creating a server group.
- Connecting this server to cloud services.

The dashboard also has links to web-based articles about new features in **Server Manager** and links to learn more about Microsoft solutions.

Server Manager has a section for properties of the local server. Here, you can perform types of initial configuration that are similar to the types possible with the sconfig tool. These include:

- Computer name and domain membership
- Windows Firewall settings
- Remote Desktop
- Network settings
- Windows Update settings
- Time zone
- Windows activation

This section also provides basic information about the hardware, such as:

- O/S version
- Processor information
- Amount of RAM

- Total disk space

There are also sections for:

- Querying specific event logs for various event severity levels over a specific time period.
- Monitoring the status of services and stopping and starting services.
- Best practices analysis to determine if the roles are functioning properly on your servers.
- A display of Performance Monitor that allows you to set alert thresholds on CPU and memory.
- Listing the installed roles and features with the ability to add and remove them.

The navigation pane will have a link to other roles installed on the server, which will provide information about specific roles such as events relating to that role. In some cases, you will observe a sub-menu that allows you to configure aspects about the role, such as File and Storage Services and Remote Desktop Services.

Remote Server Administration Tools

Remote Server Administration Tools (RSAT) are a group of management tools that enables IT administrators to remotely manage roles and features in Windows Server from a computer that is running Windows 10, Windows 8.1, and Windows 8.

When you install RSAT on Windows 8, Windows 8.1, or Windows 10, all the tools are enabled by default. You can later choose to disable the tools by using **Turn Windows features on or off** in Control Panel. RSAT for Windows 10 consists of the full complement of available management tools including the following:

Table 1: Management tools in RSAT

Tool	Description
Active Directory Certificate Services (AD CS) Tools	AD CS Tools include Certification Authority, Certificate Templates, Enterprise PKI, and Online Responder Management snap-ins.
Active Directory Domain Services (AD DS) Tools and Active Directory Lightweight Directory Services (AD LDS) Tools	AD DS Tools and AD LDS Tools include Active Directory Administrative Center, Active Directory Domains and Trusts, Active Directory Sites and Services, Active Directory Users and Computers, ADSI Edit, Active Directory module for Windows PowerShell, and tools such as DCPromo.exe, LDP.exe, NetDom.exe, NTDSUtil.exe, RepAdmin.exe, DCdiag.exe, DSACLs.exe, DSAdd.exe, DSDBUtil.exe, DSMgmt.exe, DSMod.exe, DSMove.exe, DSQuery.exe, DSRm.exe, GPFixup.exe, KSetup.exe, KtPass.exe, NITest.exe, NSLookup.exe, and W32tm.exe.
Best Practices Analyzer	Best Practices Analyzer cmdlets for Windows PowerShell
BitLocker Drive Encryption Administration Utilities	Manage-bde, Windows PowerShell cmdlets for BitLocker, and BitLocker Recovery Password Viewer for Active Directory

Tool	Description
DHCP Server Tools	DHCP Server Tools include the DHCP Management Console, the DHCP Server cmdlet module for Windows PowerShell, and the Netsh command line tool
DirectAccess, Routing and Remote Access	Routing and Remote Access management console, Connection Manager Administration Kit console, Remote Access provider for Windows PowerShell, and Web Application Proxy
DNS Server Tools	DNS Server Tools include the DNS Manager snap-in, the DNS module for Windows PowerShell, and the Ddnscmd.exe command line tool.
Failover Clustering Tools	Failover Clustering Tools include Failover Cluster Manager, Failover Clusters (Windows PowerShell cmdlets), MSClus, Cluster.exe, Cluster-Aware Updating management console, and Cluster-Aware Updating cmdlets for Windows PowerShell.
File Services Tools	File Services Tools include the following: Share and Storage Management Tools, Distributed File System Tools, File Server Resource Manager Tools, Services for NFS Administration Tools, iSCSI management cmdlets for Windows PowerShell;
Distributed File System Tools	Distributed File System Tools include the DFS Management snap-in, and the Dfsradmin.exe, Dfsrdiag.exe, Dfscmd.exe, Dfsdiag.exe, and Dfsutil.exe command line tools and PowerShell modules for Distributed File System Name Space (DFSN) and Distributed File System Replication (DFSR).
File Server Resource Manager Tools	These include the File Server Resource Manager snap-in and the Dirquota.exe, Filescrn.exe, and Storrep.exe command line tools.
Group Policy Management Tools	Group Policy Management Tools include Group Policy Management Console, Group Policy Management Editor, and Group Policy Starter GPO Editor.
Network Load Balancing Tools	Network Load Balancing Tools include the Network Load Balancing Manager, Network Load Balancing Windows PowerShell cmdlets, and the NLB.exe and WLBS.exe command line tools.
Remote Desktop Services Tools	Remote Desktop Services Tools include the Remote Desktop snap-ins, RD Gateway Manager, tsgateway.msc, RD Licensing Manager, licmgr.exe, RD Licensing Diagnoser, and lsdiag.msc. Use Server Manager to administer all other RDS role services except RD Gateway and RD Licensing.
Server Manager	Server Manager includes the Server Manager console.

Tool	Description
SMTP Server Tools	SMTP Server Tools include the Simple Mail Transfer Protocol (SMTP) snap-in
Windows System Resource Manager Tools	Windows System Resource Manager Tools include the Windows System Resource Manager snap-in and the Wsrmc.exe command line tool.
Volume Activation	Manages volume activation through the vmw.exe file.
Windows Server Update Services Tools	Windows Server Update Services Tools include the Windows Server Update Services snap-in, WSUS, msc, and PowerShell cmdlets

Windows PowerShell

Windows PowerShell is a command line shell and scripting language that allows task automation and configuration management. Windows PowerShell cmdlets execute at a Windows PowerShell command prompt or combine into Windows PowerShell scripts. PowerShell 5.1 is included natively in Windows Server 2016 and Windows Server 2019.

Cmdlets

PowerShell uses cmdlets to perform tasks. A cmdlet is a small command that performs a specific function. You can combine multiple cmdlets to perform multiple tasks either as command line entries or to run as a script. Cmdlets employ a verb/noun naming pattern joined by a hyphen. This makes each cmdlet more literal and easier to interpret and remember. For example, in the cmdlet **Get-service**, Get is the action and service is the object the action will be performed on. This command will return a listing of all services installed on the computer and their status.

You can further granularize most cmdlets by adding parameters to fine tune the results of the cmdlet. For example, if you are interested in a specific service, you can append the *-Name* parameter with the name of the service to return information about that specific service. For example, **Get-service -Name Spooler** will return information about the status of the Print Spooler service.

Multiple cmdlets can be piped together by using the vertical line (|) character. This will help you string together cmdlets to format, filter, sort, and refine the results. The output of the first cmdlet is piped as input to the next cmdlet for further processing. For example, **Get-service -Name Spooler|restart-service** will retrieve the Spooler service object and then perform the command to restart the Print Spooler service.

For repetitive tasks, you can save these cmdlets into a script and run them manually or schedule them to run regularly. You can create a script easily by entering the commands into a text editor such as Notepad and saving the file with a PS1 extension. You can manually run the script by entering the script name in the PowerShell command shell or schedule with Task Scheduler.

Modules

Many products such as Microsoft SharePoint and Hyper-V have their own set of cmdlets specific to that product and some even have their own command shell that automatically loads the cmdlets for that app, such as Microsoft Exchange. These application-specific cmdlets are packaged together and installed as modules so that all the proper commands for that application are available. Usually, these modules become available to the PowerShell environment by installing the application. The PowerShell module for

that app is installed as part of the installation. Occasionally, you need to load these modules into the command shell by using the **Install-Module** cmdlet.

PowerShell Integrated Scripting Environment (ISE)

PowerShell ISE is a GUI-based tool that allows you to run commands and create, change, debug, and test scripts. The GUI is a tabbed interface, much like a browser, which allows you to work on multiple PowerShell projects, each in an isolated tab. The screen is split into three main sections. The top pane is a text editor where you enter your commands and the bottom pane is a command shell that displays the results so you can test the script while in development. A third pane occupies the right side of the screen and displays a listing of all the available cmdlets. After you are satisfied with the script that you create, you can use the save command from the menu to save the script.

PowerShell Direct

Many administrators choose to run some of their servers in virtualized environments. To enable a simpler administration of Windows Server Hyper-V VMs, Windows 10 and Windows Server 2019 both support a feature called PowerShell Direct.

PowerShell Direct enables you to run a Windows PowerShell cmdlet or script inside a VM from the host operating system regardless of network, firewall, and remote management configurations.

Demonstration: Manage servers remotely

In this demonstration, you will learn how to perform Windows PowerShell remote management. You will use PowerShell remote from the management server to check the status of the Internet Information Services (IIS) Admin service, and then restart the IIS Admin service on **SEA-DC1**. Then you will get a listing of all services and add it to a text file on the management computer.

Demonstration steps

1. Switch to **WS-011T00A-SEA-ADM1-B** and sign in as **Administrator**.
2. Launch PowerShell in an elevated admin session.
3. Run the cmdlet `Enter-PSSession -ComputerName SEA-DC1`.
4. Run the cmdlet `Get-Service -Name IISAdmin`. Observe the results.
5. Run the cmdlet `Get-Service -Name IISAdmin|Restart-Service`. Observe the results.
6. Run the cmdlet `Get-Service|Out-File \\SEA-ADM1\C$\ServiceStatus.txt`.
7. Use File Explorer to check if **ServiceStatus.txt** was created, and then open the file.
8. Close all open windows.

Overview of the least-privilege administration concept

Most security breaches or data loss incidents are the result of human error, malicious activity, or a combination of both. For example, a user is logged on with an account that has Enterprise Admin rights and opens an email attachment that runs malicious code. That code will have full admin rights across the enterprise because the user that ran it had full admin rights.

Least privilege is the concept of restricting access rights for users, service accounts, and computing processes to only those resources absolutely required to perform their job roles. Although the concept is easy to understand, it can be complex to implement, and in many cases, it's simply not adhered to. The principle states that all users should sign in with a user account that has the minimum permissions necessary to complete the current task and nothing more. Doing so supplies protection against malicious code, among other attacks. This principle applies to computers and the users of those computers.

Additional reading: For more information, go to **Implementing Least-Privilege Administrative Models**¹.

Delegated privileges

Accounts that are members of high privilege groups such as Enterprise Admins and Domain Admins have full access to all systems and data. As such, those accounts must be closely guarded, but there will be users who need certain admin rights to perform their duties. For example, help desk staff must be able to reset passwords and unlock accounts for ordinary users, while some IT staff will be responsible for installing applications on clients or servers, or performing backups.

Delegated privilege supplies a way to grant limited authority to certain users or groups. Also, Active Directory and member servers have built-in groups that have predetermined privileges assigned. For example, Backup Operators and Account Operators have designated rights assigned to them.

Additional reading: For more information about Active Directory security groups, go to **Active Directory Security Groups**².

If the built-in security groups do not meet your needs, you can delegate more granular privileges to users or groups by using the **Delegation of Control Wizard**. The wizard allows you to assign permissions at the site, domain, or organization unit level. The wizard has the following pre-defined tasks that you can assign:

- Create, delete, and manage user accounts
- Reset user passwords and force password change at next sign in
- Read all user information
- Create, delete, and manage groups
- Modify the membership of a group
- Join a computer to the domain (only available at the domain level)
- Manage Group Policy links
- Generate Resultant Set of Policy (Planning)
- Generate Resultant Set of Policy (Logging)
- Create, delete, and manage inetOrgPerson accounts
- Reset inetOrgPerson passwords and force password change at next logon
- Read all inetOrgPerson information

You can also combine permissions to create and assign custom tasks.

¹ <https://aka.ms/implementing-least-privilege-administrative-models>

² <https://aka.ms/active-directory-security-groups>

Demonstration: Delegate privileges

In this demonstration, you will learn how to use the **Delegation of Control Wizard**. You will create a group for sales managers and add a user from the **Managers** organizational unit (OU). You will use the **Delegation of Control Wizard** to grant permission to reset passwords for users in the **Sales** OU. Then, you will test the delegation.

Demonstration steps

Create the Sales Managers group and add a user

1. Connect to **WS-011T00A-SEA-ADM1-B**.
2. From Windows Administrative Tools, launch **Active Directory Users and Computers**.
3. Create a new global security group named **Sales Managers**.
4. Add **Ajay Manchepalli** to the **Sales Managers** group.

Delegate the permission to reset the password for users in the Sales OU to the Sales Managers group

1. Launch the **Delegation of Control Wizard** for the Sales OU.
2. Assign the **Reset user passwords and force password change at next logon** to **Sales Managers**.
3. Sign out.

Test the delegation

1. Sign into **WS-011T00A-SEA-ADM1-B** as **Ajay** with a password of **Pa55w.rd**.
2. Open the **Sales** OU.
3. Reset the password for one of the users. It will succeed.
4. Open the **Research** OU.
5. Attempt to reset the password for one of the users. You will be denied access.

Privileged Access Workstations

Malicious hackers will focus on workstations that are regularly used by administrators with high-level access to the infrastructure. A Privileged Access Workstation (PAW) is a computer that you can use for performing administration tasks, such as administration of identity systems, cloud services, and other sensitive functions. This computer will be protected from the internet and locked down so that only the required administration apps can run.

You should never use this workstation for web browsing, email, and other common end user apps, and it should have strict app whitelisting. You shouldn't allow connection to Wi-Fi networks or to external USB devices. A PAW should implement security features such as multifactor authentication. You must configure privileged servers to not accept connections from a non-privileged workstation.

Microsoft recommends using Windows 10 Enterprise because it supports security features that are not available in other editions, such as **Credential Guard** and **Device Guard**. Microsoft recommends using one of the following hardware profiles:

- **Dedicated hardware.** Separate dedicated devices for user tasks vs. administrative tasks. The admin machine must support hardware security mechanisms such as a Trusted Platform Module (TPM).
- **Simultaneous use.** Single device that can run user tasks and administrative tasks concurrently by running two operating systems, where one is a user system and the other is an admin system. You can accomplish this by running a separate operating system in a VM for daily use.

Additional reading: For more information on Privileged Access Workstations, go to **Privileged Access Workstations**³.

Jump servers

A jump server is a hardened server used to access and manage devices in a different security zone, such as between an internal network and a perimeter network. The jump server can function as the single point of contact and management. Jump servers do not typically have any sensitive data, but user credentials will be stored in the memory and malicious hackers can target it. For that reason, jump servers need to be hardened. A jump server would typically be accessed by a Privileged Access Workstation to ensure secure access.

Securing the configuration

This server will run on dedicated hardware that supports both hardware and software-based security features such as:

- **Credential Guard** to encrypt the domain credentials in memory.
- **Remote Credential Guard** to prevent remote credentials from being sent to the jump server, instead using Kerberos version 5 single sign-on tickets.
- **Device Guard** that uses Hypervisor Enforced Code Integrity (HVCI) to use virtualization-based security to enforce kernel mode components to follow with the code integrity policy.
- **Device Guard** that uses Config Code Integrity to allow admins to create a custom code integrity policy and specify trusted software.

Test your knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

What cmdlet can be run on a remote Windows computer to allow PowerShell remote management?

- Enable-PSSession
- Enable-PSRemoting
- Enable-PSSessionConfiguration

³ <https://aka.ms/privileged-access-workstations>

Question 2

True or False: The Windows Admin Center is supported on Internet Explorer 11.

- True
- False

Introducing Windows Server 2019

Lesson overview

In this lesson, you will learn about the Windows Server 2019 editions and their capabilities. You will learn about the hardware requirements and various deployment options. You will be able to describe deployment accelerators, servicing channels, and licensing models for Windows Server. Finally, you will learn about the new features in Windows Server 2019.

Lesson objectives

After completing this lesson, you will be able to:

- Describe the different editions of Windows Server 2019.
- Identify hardware requirements for Windows Server 2019.
- Describe the deployment options.
- Describe deployment accelerators.
- Identify the servicing channels for Windows Server.
- Describe the licensing and activation models for Windows Server.
- Describe the new features in Windows Server 2019.

Windows Server 2019 editions

You can choose one of the four editions of Windows Server 2019. These editions allow organizations to select a version of Windows Server 2019 that best meets their needs, rather than pay for features they do not require. When deploying a server for a specific role, system administrators can save by selecting the appropriate edition. Windows Server 2019 is released in the following four editions:

- Windows Server 2019 Essentials
- Windows Server 2019 Standard
- Windows Server 2019 Datacenter
- Microsoft Hyper-V Server 2019

Each edition supports unique features. The following table describes the Windows Server 2019 editions:

Table 1: Windows Server editions

Edition	Description
Windows Server 2019 Essentials	Like its predecessor, Windows Server 2019 Essentials edition is designed for small businesses. This edition allows up to 25 users and 50 devices. Users do not need Client Access Licenses (CALs) to connect to the server, but you can't increase the 25-user limit. It supports two processor cores and up to 64 gigabytes (GB) of random-access memory (RAM). It includes added support for Microsoft Azure Active Directory (Azure AD) through Azure AD Connect. If configured as a domain controller, it must be the only domain controller, must run all Flexible Single Master Operations (FSMO) roles, and can't have two-way trusts with other Active Directory domains. Microsoft recommends that small businesses move to Microsoft 365 instead of deploying Windows Server 2019 Essentials. The Windows Server Essentials Experience role has been deprecated from Windows Server Essentials 2019.
Windows Server 2019 Standard edition	Windows Server 2019 Standard edition is designed for physical server environments with little or no virtualization. It supplies most of the roles and features available for the Windows Server operating system. This edition supports up to 64 sockets and up to 4 terabytes (TB) of RAM. Nano Server is available only as a container base OS image. You need to run it on a container host as a container. You can't install a bare-metal Windows Server 2019 Nano Server. It includes licenses for up to two VMs. You can run two VMs on one physical host by using one standard license if the physical host is only used for hosting and managing the VMs. If the physical host is used to run other services such as DNS, you can only run one VM by using a standard license.
Windows Server 2019 Datacenter edition	Windows Server 2019 Datacenter edition is designed for highly virtualized infrastructures, including private cloud and hybrid cloud environments. It supplies all the roles and features available for the Windows Server operating system. This edition supports up to 64 sockets, up to 640 processor cores, and up to 4 TB of RAM. It includes unlimited VM licenses based on Windows Server for VMs that run on the same hardware. It also includes features such as Storage Spaces Direct and Storage Replica, along with Shielded VMs and features for software-defined datacenter scenarios.

Edition	Description
Hyper-V Server 2019	Acts as a standalone virtualization server for VMs, including all the new features around virtualization in Windows Server. The host operating system has no licensing cost, but you must license VMs separately. This edition supports up to 64 sockets and up to 4 TB of RAM. It supports domain joining, but it does not support Windows Server roles other than limited file service features. This edition has no GUI but does have a UI that displays a menu of configuration tasks. You can manage this edition remotely by using remote management tools.

Hardware requirements for Windows Server 2019

The hardware requirements for Windows Server 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. The following table shows the absolute minimum required for a Server Core installation on a physical machine.

Table 1: Server Core installation requirements

Component	Requirement
Processor architecture	64-bit
Processor speed	1.4 Gigahertz (GHz)
RAM	512 MB. Note that VMs require at least 800 MB of RAM during installation. You can reduce it to 512 MB after the installation is complete.
Hard drive space	32 GB

Virtualized deployments of Windows Server require the same hardware specifications for physical deployments. However, during installation, you will need to allocate extra memory to the VM, which you can then deallocate after installation, or you will need to create an installation partition during the boot process.

Desktop Experience

To install Windows Server with Desktop Experience, you need a minimum of 4 GB hard drive space.

Other hardware requirements

In addition to the previously referenced requirements, there are a variety of other hardware requirements to consider, depending on your specific organizational needs and installation scenarios:

- Greater disk space is needed for network installations or for computers with more than 16 GB of RAM.
- Storage and network adapters must be PCI Express compliant.

Overview of deployment options

There are several ways to move your server infrastructure to Windows Server 2019. You can perform a clean install of the operating system to new hardware or a VM and migrate roles and applications to the new server installation. However, Windows Server 2019 now has the option to perform an upgrade of the existing server operating system. This allows you to keep your configurations, server services, applications, and server roles.

Cluster operating system rolling upgrades allow you to upgrade the operating systems of cluster nodes without having to stop the Hypervisor or any Scale-Out File Server Workloads.

Clean install

A clean install to new or existing hardware, or a VM is the easiest way to install Windows Server 2019. The installation steps have not changed from Windows Server 2016 and follows these basic steps:

1. Boot the machine or VM from the Windows Server 2019 media.
2. Choose the installation language, time and currency format, and keyboard layout.
3. Choose the architecture, either Standard or Datacenter, with or without Desktop Experience.
4. Accept the license.
5. Choose custom installation.
6. Choose the volume that will host the installation.

The installation will copy the files and install the operating system. After the installation is complete, you will create a password for the local Administrator account. Further configuration will occur after the initial sign-in into the Administrator account.

In-place upgrade

An in-place upgrade allows you to upgrade your server operating system and keep all the server roles, applications, and data intact. You can upgrade from Standard to Datacenter, but you must have the proper licensing. You can only upgrade Windows Server 2012 R2 and newer to Windows Server 2019.

Note: If you switch from Desktop Experience to Core, you can't preserve files, apps, or settings.

The steps to upgrade existing operating system are:

1. Insert the disk or mount the ISO of Windows Server 2019 media and then run Setup.exe.
2. Respond to the prompt to download updates, drivers, and optional features.
3. Choose the architecture, either Standard or Datacenter, with or without Desktop Experience.
4. Accept the license.
5. Choose what to keep, personal files and apps, or nothing.

The upgrade will take some time to complete and then the server will restart.

Deployment accelerators

Organizations should consider using software tools to help them plan their upgrade and migration to Windows Server 2019. Along with guidance content to help you design and plan your Windows Server deployment, Microsoft also supplies solution accelerators to aid in the process.

Solution accelerators are free, scenario-based guides and automations designed to help you with

planning, deploying, and operating Microsoft products and technologies. Solution Accelerator scenarios focus on security and compliance, management and infrastructure, and communication and collaboration.

Microsoft Deployment Toolkit (MDT)

Microsoft Deployment Toolkit (MDT) is both a process and a lightweight tool for automated server (and desktop) deployments. It's used for deploying standardized images. MDT is based on a variety of Microsoft technologies including Preboot Execution Environment (PXE), Windows Deployment Services (WDS), and Microsoft Endpoint Configuration Manager. 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.

Microsoft Assessment and Planning Toolkit (MAP)

The Microsoft Assessment and Planning Toolkit (MAP) is an agentless 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 2019, Windows 10, and for other scenarios, such as:

- Assessing the environment for Microsoft 365 and Office 2019.
- Sizing your desktop virtualization needs for Virtual Desktop Infrastructure (VDI).
- Migrating to Microsoft Azure VM.
- Virtualizing Linux servers to Hyper-V.
- Setting up SQL Server platforms in the cloud.
- Planning Hyper-V servers.
- Assessing Microsoft Volume licensing compliance and positioning for Server and Cloud Enrollment.

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 can run Windows Server 2019, which machines can run Windows Server 2019 with minimum system requirements, and which machines are not capable of running Windows Server 2019. MAP also recommends specific upgrades that ensure computers can run Windows Server 2019.
- Generate a report or proposal based on the Windows Server 2019 Readiness Assessment. The report or proposal is a document that contains an executive overview, assessment results, next steps, and a worksheet summarizing Windows Server 2019 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 can also choose which current physical servers are the best candidates for virtualization and the hosts on which you should place those VMs.

Additional reading: For more information, go to **Microsoft Assessment and Planning Toolkit**⁴.

⁴ <https://aka.ms/assessment-planning-toolkit>

Servicing channels for Windows Server

Servicing channels allow you to choose if new features and functionality will be delivered regularly during the production lifespan of the server, or if you will choose when to move to a new server version. Windows Server supports two release channels, Long Term Servicing Channel (LTSC) and the Semi-Annual Channel (SAC).

Long-Term Servicing Channel

LTSC dictates that a major version of Windows Server will release every two or three years. This includes five years of mainstream support and five years of extended support from the release date. Normal security updates and Windows updates will continue to deliver on a regular basis as in the past, but without most new features or functionality. LTSC releases do get some new feature functionality. For example, Storage Migration Services is available on Windows Server 2019 and has been updated for workloads such as migrating Linux NAS appliances.

LTSC is the traditional “numbered” version of Windows Server. Windows Server 2016 and Windows Server 2019 are LTSC releases. For most server requirements, the LTSC will be the best choice.

Semi-Annual Channel

SAC only releases as Server Core or Nano Server container images, so it's restricted in the roles and features that you can install. New features will be delivered semi-annually, once in the second quarter and once in the fourth quarter. SAC is limited to software assurance and cloud customers. These releases will be supported for 18 months from the initial date of release. Normal security updates and Windows updates will continue to be delivered on a regular basis. Features that are included in SAC will be rolled up and delivered to the LTSC on the next major release. SAC releases can be identified by their version number, which is a combination of the year and month the features were released. For example, version 1903 means the feature was released in the third month of 2019. You cannot do an in-place upgrade from an LTSC release to a SAC release.

Note: SAC releases will always be a clean installation. SAC implies you have a CI/CD type pipeline where you'd just deploy the newer OS image in the same way that you'd deploy a new container when the base container image got updates.

Choosing a servicing channel

LTSC is recommended for general purpose file servers, Microsoft and non-Microsoft workloads, traditional apps, infrastructure roles, software-defined Datacenter, and hyper-converged infrastructure.

Semi-Annual Channel is recommended for containerized applications, container hosts, and application scenarios benefiting from adaption of new features.

Additional reading: For more information on servicing channels, go to **Windows Server servicing channels: LTSC and SAC**⁵.

⁵ <https://aka.ms/servicing-channels-19>

Licensing and activation models for Windows Server

Licensing for Windows Server Essentials is per server. It includes Client Access Licenses for 25 users and is limited to 2 sockets. You can't buy licensing for more than this limit. The licensing model for Windows Server Standard and Datacenter changed with Windows Server 2016 and continues through the 2019 version. Licensing for Windows Server Standard and Datacenter is now based on the number of cores, not processors.

Core-based licensing requirements

Each Windows Server has the following minimum license requirement:

- All physical cores must be licensed
- There must be eight core licenses per processor

This minimum applies to both Datacenter and Standard version. Licenses can be bought in 2-core packs or 16-core packs. Servers that have more processors require you to buy more core licenses. For example, a server with four processors will require 32 core licenses because each processor needs eight core licenses.

Note: Standard edition includes licensing for two VMs or Hyper-V containers. If more VMs are added, you must buy more core licenses.

Note: There are special licensing rules for VMs running in failover clusters. Licensing must fully support all the VMs to run on each host. For example, if you have a two-node failover cluster with two VMs on each host, then each host must be licensed to support all four VMs potentially running on a single host.

Client Access Licenses (CALs)

Each user or device that connects to the server for any purpose must have a Client Access License (CAL). There are user CALs and device CALs. CALs allow users or devices to connect to any of the servers in the organization. In a typical organization where each user has an assigned workstation, it makes sense to buy a user CAL for each user. If your organization has many shared workstations, it might make more sense to purchase a device CALs.

Note: Remote desktop connections are not included in user CALs. You need to buy remote desktop CALs separately.

Windows server activation

There are multiple ways to activate Windows Server. Often, if you buy a server from an original equipment manufacturer (OEM), it will come with the operating system pre-installed and activated.

To ensure that your organization has the proper licenses, and to receive notices for product updates, you must activate every copy of Windows Server that you install. Windows Server 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, you can't customize your operating system. There are two general activation strategies:

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

Manual activation

When you perform a manual activation, you must enter the product key. You can perform manual activation by using the retail product key or the multiple activation key (MAK). You can use a retail product key to activate only a single computer. However, a MAK 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. OEM keys 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 10. 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 supplies an option to automatically activate a large number of computers without having to enter product keys manually on each system.

There are several technologies available that you can use to automate activating Windows Server licenses:

- **Key Management Services (KMS)**. This 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 Microsoft 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 5 servers and 25 clients. KMS is the default key for volume activation.
- **Active Directory-Based Activation**. This is a service that lets you use Active Directory Domain Services (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.
- **Volume Activation Tools** console. This console is used to install, activate, and manage volume license activation keys in AD DS or KMS.
- **Volume Activation Management Tool (VAMT)**. This is a no cost tool that you can use to manage volume activation by 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.
- **Automatic Virtual Machine Activation (AVMA)**. AVMA lets you install VMs on a virtualization server with no product key, even in disconnected environments. AVMA binds the VM activation to the licensed virtualization server and activates the VM when it starts up. AVMA is supported on Windows Server 2019 Datacenter.

Additional reading: For more information about activating Windows operating systems, go to **Review and Select Activation Methods**⁶.

What's new in Windows Server 2019?

With each updated version of Windows Server, Microsoft introduces new and innovative technologies to improve administration or add needed functionality. Windows Server 2019 is designed to easily link your

⁶ <https://aka.ms/review-select-activation-methods>

on-premises infrastructure with Microsoft Azure. It includes many improvements to existing and new features, including:

Table 1: Features of Windows Server 2019

Feature	Description
Deduplication for ReFS volumes	Windows Server 2019 fully supports deduplication of the Resilient File System (ReFS) file system. This can save large amounts of storage space when used for Hyper-V machine storage.
Storage Class Memory support	Storage created from flash-based non-volatile media that is connected to a dual in-line memory module (DIMM) slot much like traditional dynamic random-access memory (DRAM). This concept moves the storage closer to the CPU to improve performance.
Cluster sets	Allows you to create large scale-out clusters. A cluster set is a group of multiple failover clusters that is loosely coupled to a single master endpoint which distributes requests.
Storage Migration Services	Allows you to inventory and migrate data, security, and configurations from legacy systems to Windows Server 2019 or Azure.
System Insights	Provides local predictive analytics capabilities native to Windows Server. It focuses on capacity forecasting, and predicting future usage for computing, networking, and storage, which allows you to proactively manage your environment.
Storage Replica for Standard edition	Previously only available for Datacenter, it's now included with Standard edition, with some limitations:
	Servers must run 2019
	Only single volumes can be replicated
	Volume size is limited to 2 TB
Windows Defender Advanced Threat Protection and Windows Defender Exploit Guard	Previously only available for Windows 10 platforms, it is a new set of host intrusion prevention such as attack detection and zero-day exploits. It is a single solution to detect and respond to advanced threats.
Shielded VMs for Linux	Protects Linux VMs from attacks and rogue administrators.
Azure Stack Hyper-Converged Infrastructure (HCI)	HCI is a fully software-defined platform based on Hyper-V. You can dynamically add or remove host servers from the Windows Server 2019 Hyper-V HCI cluster to increase or decrease capacity.

Feature	Description
Server Core App Compatibility Feature on Demand (FOD)	An optional feature package that you can add to Windows Server 2019 Server Core installations. It improves the app compatibility of the Windows Server Core by including a subset of binaries and packages from Windows Server with Desktop Experience, without adding the Desktop Experience graphical environment.

Test your knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

You are the administrator of a small company of 50 users. Most of your business applications are cloud based. You're going to set up two Windows Servers, one as a domain controller and one as a file and print server. Which edition of Windows Server will best suit your needs?

- Standard
- Essentials
- Hyper-V
- Datacenter

Question 2

Which tool can help you inventory your organization's IT infrastructure?

- Microsoft Deployment Toolkit
- Microsoft Assessment and Planning Toolkit

Overview of Windows Server Core

Lesson overview

In this lesson, you will learn about the differences between Server Core and Windows Server with Desktop Experience. The Server Core option is a minimal installation option that is available when you are deploying the Standard or Datacenter edition of Windows Server. You must know how to enable and perform the remote management of your server infrastructure because Server Core provides no graphical management tools. You'll learn about the installation options and the tools used to configure and manage Windows Server Core.

After completing this lesson, you will be able to:

- Describe the differences between Server Core and Windows Server with Desktop Experience.
- Describe how to perform the installation and post installation tasks.
- Describe how to install features on demand.
- Describe how to Use the sconfig tool in Server Core.
- Explain how to configure Server Core.

Server Core vs. Windows Server with Desktop Experience

When you install Windows Server 2019, you need to choose between installing the server with or without the Desktop Experience. This is an important decision because you can't add or remove the Desktop Experience after you install the server.

Server Core is an installation of Windows Server without the Desktop Experience. Server Core is available for both Standard and Datacenter editions, but it isn't available for Windows Server 2019 Essentials, and the free version of Hyper-V server is only available as a Server Core installation.

You can administer and configure Server Core on the server itself through PowerShell, the command line, or through the text-based tool called Sconfig. Remote administration is the normal method of managing the server by using several tools such as PowerShell Remoting, the Remote Server Administration Tool (RSAT), and the Windows Admin Center.

Note: There are some GUI-based tools available in Server Core. For example, Regedit, Notepad, Msinfo32, and Task Manager (Taskmgr) will launch from the command prompt in their traditional GUI.

Server Core has advantages over Windows Server with Desktop Experience and is the recommended installation for most scenarios, but it might not be suitable in every case. The following table lists the major advantages and disadvantages:

Table 1: Advantages and Disadvantages of Server Core installation

Advantages	Disadvantages
Small footprint which uses fewer server resources and less disk space, as little as 5 GB for a basic installation	You can't install several applications on Server Core. The applications include Microsoft Server VM Manager 2019, System Center Data Protection Manager 2019, SharePoint Server 2019, Project Server 2019, and Exchange versions prior to Exchange 2019.

Advantages	Disadvantages
Because Server Core installs fewer components, there are fewer software updates. This reduces the number of monthly restarts required and the time required for you to service Server Core.	Several roles and role services are not available, including Remote Desktop Services Session Host, Web Access, and Gateway service; Fax Server; SMTP Server; and Windows PowerShell ISE
The small attack surface makes Server Core much less vulnerable to exploits	You can't install many vendor lines of business applications on Server Core. However, the App Compatibility Feature on Demand can help mitigate that in some cases.

Choosing your installation

Choosing your installation will depend on the workload you need the server to perform. For network infrastructure roles, such as Hyper-V, Active Directory, File Server, and Web Server, Server Core is the better option.

Additional reading: For more information about server roles and features not available for Server Core, go to **Roles, Role Services, and Features not in Windows Server - Server Core**⁷.

Additional reading: For more information about available server roles and features for Server Core, go to **Roles, Role Services, and Features included in Windows Server - Server Core**⁸.

If you have requirements for a line of business apps that require a GUI or the presence of certain binaries, then Windows Server with Desktop Experience will be the right choice.

Server Core installation and post-installation tasks

The installation of Server Core is straightforward and the same whether you are installing with Desktop Experience or not. There are some tasks you should perform before installing the operating system:

- Disconnect any uninterruptible power supply (UPS) that is connected to the destination computer with a serial cable. This is because setup attempts to detect any devices connected to serial ports, and UPS equipment can cause problems with this process.
- Back up your server if this is an upgrade install.
- Disable or remove virus protection software that might be installed on the target computer if this is an upgrade.
- Ensure you have any mass storage driver files provided by the manufacturer on a disk, flash drive, or other portable media so that the driver files can be provided during setup. Most modern servers provide a disk or built-in wizard with appropriate drivers to guide you through the installation for that specific hardware.

Typically, before you install the operating system, you will use the vendor-provided guidance to do the initial hardware configurations, which includes the following tasks:

- Update BIOS, firmware, and drivers
- Configure disk arrays
- Configure out of band management

⁷ <https://aka.ms/server-core-removed-roles>

⁸ <https://aka.ms/server-core-roles-and-services>

- Configure network settings

After those tasks are complete, you can install the operating system by performing the following steps:

1. Connect to the installation source. Options for this include:
 - Insert a DVD-ROM containing the installation files, and boot from the DVD-ROM.
 - Connect a specially prepared USB drive that hosts the installation files.
 - Perform a Preboot Execution Environment (PXE) boot and connect to a Windows Deployment Services server.
2. On the first page of **Windows Setup Wizard**, select the following locale-based information:
 - Language to install
 - Time and currency format
 - Keyboard or input method
3. On the second page of **Windows Setup Wizard**, select **Install now**.
4. In **Windows Setup Wizard**, on the **Select The Operating System You Want To Install** page, choose from the available operating system installation options. The default option is **Server Core Installation**.
5. On the **License Terms** page, review the terms of the operating system license. You must choose to accept the license terms before you can go ahead with the installation process.
6. On the **Which Type Of Installation Do You Want** page, you have the following options:
 - **Upgrade**. Select this option if you have an existing installation of Windows Server that you want to upgrade to Windows Server 2019. You should launch upgrades from within the previous version of Windows Server rather than booting from the installation source.
 - **Custom**. Select this option if you want to perform a new installation.
7. On the **Where do you want to install Windows** page, choose an available disk on which to install Windows Server. You can also choose to repartition and reformat disks from this page. When you select **Next**, the installation process will copy the files and reboot the computer several times.
8. On the **Settings** page, provide a password for the local Administrator account.

Install features on demand

In the past, if you tried to install an application that had dependencies on certain binaries and packages from the Desktop Experience that were not present in Windows Server Core, the install would fail.

Microsoft is striving to improve the Windows Server Core experience by releasing the Server Core App Compatibility feature-on-demand (FOD), making it possible to install these applications.

The FOD does not come pre-installed. You must download and install it. You can obtain it through Windows Update if your server connects directly to the internet or you can download the ISO image file from the Microsoft Volume License Service Center.

Some operating system components that become available after installing the FOD include:

- Event Viewer
- Performance Monitor
- Resource Monitor
- Device Manager

- Microsoft Management Console
- File Explorer
- Internet Explorer
- Windows PowerShell ISE
- Failover Cluster Manager

Installing the FOD

There are two ways to install the FOD.

The simplest way to install the FOD is through Windows Update by using PowerShell. Launch an elevated PowerShell session and run the following command:

```
Add-WindowsCapability -Online -Name ServerCore.AppCompatibility~~~~0.0.1.0
```

Then restart the server.

If connecting to Windows Update is not an option, then use the following method:

1. Save the ISO that you downloaded to a network share.
2. Connect the Server Core to the network location and copy the ISO to a local folder.
3. Mount the ISO from an elevated PowerShell session by using the following command:

```
Mount-DiskImage -ImagePath drive_letter:\folder_where_ISO_is_saved\ISO_file-name.iso
```

4. Run the following command:

```
Add-WindowsCapability -Online -Name ServerCore.AppCompatibility~~~~0.0.1.0  
-Source <Mounted_Server_FOD_Drive> -LimitAccess
```

5. Restart the server

Use the sconfig tool in Server Core

Server Core has no GUI, so after the initial installation, you are presented with only a command prompt.

Sconfig is a text-based tool that allows you to do the basic configuration of Server Core to prepare it for use in your production environment. Sconfig is included in Windows Server Desktop Experience and Server Core.

You typically use sconfig to perform the initial configuration directly after the installation completes, but you can run it at any time to change the settings as required.

Sconfig provides 15 different options as outlined in the following table:

Table 1: Sconfig options

Option	Description
Domain/Workgroup	Join the domain or workgroup of choice
Computer Name	Set the computer name
Add Local Administrator	Add additional users to the local Administrators group

Option	Description
Configure Remote Management	Remote management is enabled by default. This setting allows you to enable or disable remote management and configure the server to respond to a ping.
Windows Update Settings	Configure the server to use automatic, download only or manual updates.
Download and Install Updates	Perform an immediate search for all updates or only recommended updates.
Network Settings	Configure the IP address to be assigned automatically by a Dynamic Host Configuration Protocol (DHCP) Server or you can assign a static IP address manually. This option also allows you to configure Domain Name System (DNS) Server settings for the server.
Date and Time	Brings up the GUI for changing the date, time, and time zone. It also has tabs to add additional clocks and choose an Internet time server to sync with.
Telemetry Settings	Allows Windows to periodically collect and upload statistical information about the server and upload it to Microsoft.
Windows Activation	Provides three options—Display license info, Activate Windows, and Install product key
Log Off User	Logs off the current user
Restart Server	Restarts the server
Shut Down Server	Shuts down the server
Exit to Command Line	Returns to the command prompt

Demonstration: Configure Server Core

In this demonstration, you will learn how to use the `sconfig` tool to perform basic configuration tasks. You will configure date and time and demonstrate the network settings configuration. Before you begin the demonstration, you must complete the following steps:

Demonstration steps

1. Connect to **WS-011T00A-SEA-DC1-B** and sign in as **Administrator** by using the password **Pa55w.rd**.
2. Run `sconfig`.
3. Briefly discuss the various options, set the time zone to your time zone, and then return to the main menu.
4. Describe the network settings, and then return to the main menu.
5. Return to the main menu.
6. Leave the VM running for the next demonstration.

Test your knowledge

Use the following question to check what you've learned in this lesson.

Question 1

Which of the following roles or role services can run on Server Core? Select two.

- SMTP server
- Web Server IIS
- Remote Desktop Gateway
- Active Directory Certificate Services

Module 01 lab and review

Lab: Deploying and configuring Windows Server

Scenario

Contoso, Ltd. wants to implement several new servers in their environment, and they have decided to use Server Core. They also want to implement Windows Admin Center for remote management of both these servers and other servers in the organization.

Objectives

- Deploy and configure Server Core
- Implement and configure Windows Admin Center

Estimated time: 45 minutes

Module review

Use the following questions to check what you've learned in this module.

Question 1

What tool is commonly used for the initial configuration of Server Core?

- Windows Admin Center
- Windows PowerShell
- Sconfig
- Server Manager

Question 2

You have Windows Server Standard edition installed and it has DNS and DHCP and Hyper-V installed. How many VMs can you run in Hyper-V before you need to buy a license?

- One
- Two
- Unlimited
- None

Question 3

True or False: You must install an SSL certificate to use the Windows Admin center.

- True
- False

Question 4

You want the helpdesk group to only be able to add and remove users from security groups. How should you accomplish this?

- Add the helpdesk group to the Account Operators group
- Add the helpdesk group to the Server Operators group
- Use the Delegation of Control Wizard to assign the task
- Add the helpdesk group to the Domain Admins group

Answers

Question 1

What cmdlet can be run on a remote Windows computer to allow PowerShell remote management?

- Enable-PSSession
- Enable-PSRemoting
- Enable-PSSessionConfiguration

Explanation

The Enable-PSRemoting cmdlet will allow PowerShell remote management. PowerShell remote management is enabled by default on Windows Servers 2012 and newer, but not on client computers.

Question 2

True or False: The Windows Admin Center is supported on Internet Explorer 11.

- True
- False

Explanation

The Windows Admin Center is not supported on Internet Explorer and will return an error if you try to launch it.

Question 1

You are the administrator of a small company of 50 users. Most of your business applications are cloud based. You're going to set up two Windows Servers, one as a domain controller and one as a file and print server. Which edition of Windows Server will best suit your needs?

- Standard
- Essentials
- Hyper-V
- Datacenter

Explanation

The Standard edition is the best choice because its license allows two VMs to run and you need two servers. The Essentials edition does not allow that many users and Datacenter would be expensive for only two servers. Hyper-V is free but you would have to pay for two server licenses for the VMs that run on it.

Question 2

Which tool can help you inventory your organization's IT infrastructure?

- Microsoft Deployment Toolkit
- Microsoft Assessment and Planning Toolkit

Explanation

The Microsoft Assessment and Planning Toolkit is an agentless 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. The Microsoft Deployment Toolkit is used for deploying standardized images.

Question 1

Which of the following roles or role services can run on Server Core? Select two.

- SMTP server
- Web Server IIS
- Remote Desktop Gateway
- Active Directory Certificate Services

Explanation

You can install certain roles on Server Core while some roles are not available because Server Core does not have the code base required for those roles.

Question 1

What tool is commonly used for the initial configuration of Server Core?

- Windows Admin Center
- Windows PowerShell
- Sconfig
- Server Manager

Explanation

Sconfig is the best tool for the initial configuration of Server Core. It allows for IP address assignment, setting computer name, and domain membership.

Question 2

You have Windows Server Standard edition installed and it has DNS and DHCP and Hyper-V installed. How many VMs can you run in Hyper-V before you need to buy a license?

- One
- Two
- Unlimited
- None

Explanation

You can run one VM before you must buy a license because you are using this host server for more than just a Hyper-V host.

Question 3

True or False: You must install an SSL certificate to use the Windows Admin center.

- True
- False

Explanation

True, a self-generated one is included, but it is only valid for 60 days.

Question 4

You want the helpdesk group to only be able to add and remove users from security groups. How should you accomplish this?

- Add the helpdesk group to the Account Operators group
- Add the helpdesk group to the Server Operators group
- Use the Delegation of Control Wizard to assign the task
- Add the helpdesk group to the Domain Admins group

Explanation

Use the Delegation of Control Wizard to assign the task. Although Account Operators and Domain Admins would work, it would give too much administrative rights to the helpdesk group.



Module 2 Identity services in Windows Server

Overview of AD DS

Lesson overview

The Microsoft Active Directory Domain Services (AD DS) database stores information on user identity, computers, groups, services, and resources in a hierarchical structure, called the *directory*. AD DS domain controllers also host the service that authenticates user and computer accounts when they sign in to the domain. Because AD DS stores information about all domain objects, and because all users and computers must connect to AD DS domain controllers at sign-in, AD DS is the primary way to configure and manage user and computer accounts on your network. This lesson covers the core logical components and physical components that make up an AD DS deployment.

Lesson objectives

After completing this lesson, you'll be able to:

- Describe AD DS.
- Describe the components of AD DS.
- Identify and describe AD DS forests and domains.
- Describe organizational units (OUs).
- Describe the AD DS schema.
- Explain AD DS replication.
- Describe the AD DS sign-in process.
- List and describe the available tools for AD DS management and administration.
- Use tools to manage AD DS objects.

What is AD DS?

Active Directory Domain Services (AD DS) and its related services form the foundation for enterprise networks that run Windows operating systems. The AD DS database is the central store of all the domain objects, such as user accounts, computer accounts, and groups. AD DS provides a searchable, hierarchical directory and a method for applying configuration and security settings for objects in an enterprise.

AD DS includes both logical and physical components. It is important that you understand how AD DS components work together so that you can manage your infrastructure efficiently. In addition, you can use AD DS options to perform actions such as:

- Installing, configuring, and updating apps.
- Managing the security infrastructure.
- Enabling Remote Access Service and DirectAccess.
- Issuing and managing digital certificates.

Logical components

AD DS logical components are structures that you use to implement an AD DS design that is appropriate for an organization. The following table describes the types of logical components that an AD DS database contains.

Table 1: AD DS logical components

Logical component	Description
Partition	A partition, or naming context, is a portion of the AD DS database. Although the database consists of one file named Ntds.dit, different partitions contain different data. For example, the schema partition contains a copy of the Active Directory schema. The configuration partition contains the configuration objects for the forest, and the domain partition contains the users, computers, groups, and other objects specific to the domain. You can store copies of a partition on multiple domain controllers and update them through directory replication.
Schema	A schema is the set of definitions of the object types and attributes that you use to define the objects created in AD DS.
Domain	A domain is a logical administrative container for objects such as users and computers. A domain maps to a specific partition and you can organize the domain with parent-child relationships to other domains.
Domain tree	A domain tree is a hierarchical collection of domains that share a common root domain and a contiguous Domain Name System (DNS) namespace.

Logical component	Description
Forest	A forest is a collection of domains that share a common AD DS root and schema, which have a two-way trust relationship.
Organizational unit (OU)	An OU is a container object for users, groups, and computers that provides a framework for delegating administrative rights and administration by linking Group Policy Objects (GPOs).
Container	A container is an object that provides an organizational framework for use in AD DS. You can use the default containers or you can create custom containers. You can't link GPOs to containers.

Physical components

The following table describes some of the physical components of AD DS.

Table 2: AD DS physical components

Physical component	Description
Domain controller	A domain controller contains a copy of the AD DS database. For most operations, each domain controller can process changes and replicate the changes to all the other domain controllers in the domain.
Data store	A copy of the data store exists on each domain controller. The AD DS database uses Microsoft Jet database technology and stores the directory information in the Ntds.dit file and associated log files. The C:\Windows\NTDS folder stores these files by default.
Global catalog server	A global catalog server is a domain controller that hosts the global catalog, which is a partial, read-only copy of all the objects in a multiple-domain forest. A global catalog speeds up searches for objects that might be stored on domain controllers in a different domain in the forest.
Read-only domain controller (RODC)	An RODC is a special, read only installation of AD DS. RODCs are common in branch offices where physical security is not optimal, IT support is less advanced than in the main corporate centers, or line-of-business applications need to run on a domain controller.
Site	A site is a container for AD DS objects, such as computers and services that are specific to their physical location. This is in comparison to a domain, which represents the logical structure of objects, such as users and groups, in addition to computers.

Physical component	Description
Subnet	A subnet is a portion of the network IP addresses of an organization assigned to computers in a site. A site can have more than one subnet.

AD DS objects

In addition to the high-level components and objects, Active Directory Domain Services (AD DS) contains other objects such as users, groups, and computers.

User objects

In AD DS, you must configure all users who require access to network resources with a user account. With this user account, users can authenticate to the AD DS domain and access network resources.

In Windows Server, a user account is an object that contains all the information that defines a user. A user account includes the username, user password, and group memberships. A user account also contains settings that you can configure based on your organizational requirements.

The username and password of a user account serve as the user's sign-in credentials. A user object also includes several other attributes that describe and manage the user. You can use Active Directory Users and Computers, Active Directory Administrative Center, Windows PowerShell, or the **dsadd** command-line tool to create a user object.

Group objects

Although it might be practical to assign permissions and abilities to individual user accounts in small networks, this becomes impractical and inefficient in large enterprise networks. For example, if several users need the same level of access to a folder, it is more efficient to create a group that contains the required user accounts, and then assign the required permissions to the group. As an added benefit, you can change users' file permissions by adding or removing them from groups rather than editing the file permissions directly. Before you implement groups in your organization, you must understand the scope of various Windows Server group types. In addition, you must understand how to use group types to manage access to resources or to assign management rights and responsibilities.

Group types

In a Windows Server enterprise network, there are two types of groups:

- **Security.** Security groups are security-enabled, and you use them to assign permissions to various resources. You can use security groups in permission entries in access control lists (ACLs) to help control security for resource access. If you want to use a group to manage security, it must be a security group.
- **Distribution.** Email applications typically use distribution groups, which are not security-enabled. You also can use security groups as a means of distribution for email applications.

Note: When you create a group, you choose the group type and scope. The group type determines the capabilities of the group.

Group scopes

Windows Server supports group scoping. The scope of a group determines both the range of a group's abilities or permissions and the group membership. There are four group scopes:

- **Local.** You use this type of group for standalone servers or workstations, on domain-member servers that are not domain controllers, or on domain-member workstations. Local groups are available only on the computer where they exist. The important characteristics of a local group are:
 - You can assign abilities and permissions on local resources only, meaning on the local computer.
 - Members can be from anywhere in the AD DS forest.
- **Domain-local.** You use this type of group primarily to manage access to resources or to assign management rights and responsibilities. Domain-local groups exist on domain controllers in an AD DS domain, and so, the group's scope is local to the domain in which it resides. The important characteristics of domain-local groups are:
 - You can assign abilities and permissions on domain-local resources only, which means on all computers in the local domain.
 - Members can be from anywhere in the AD DS forest.
- **Global.** You use this type of group primarily to consolidate users who have similar characteristics. For example, you might use global groups to join users who are part of a department or a geographic location. The important characteristics of global groups are:
 - You can assign abilities and permissions anywhere in the forest.
 - Members can be from the local domain only and can include users, computers, and global groups from the local domain.
- **Universal.** You use this type of group most often in multidomain networks because it combines the characteristics of both domain-local groups and global groups. Specifically, the important characteristics of universal groups are:
 - You can assign abilities and permissions anywhere in the forest similar to how you assign them for global groups.
 - Members can be from anywhere in the AD DS forest.

Computer objects

Computers, like users, are security principals, in that:

- They have an account with a sign-in name and password that Windows Server changes automatically on a periodic basis.
- They authenticate with the domain.
- They can belong to groups and have access to resources, and you can configure them by using Group Policy.

A computer account begins its lifecycle when you create the computer object and join it to your domain. After you join the computer account to your domain, day-to-day administrative tasks include:

- Configuring computer properties.
- Moving the computer between OUs.
- Managing the computer itself.

- Renaming, resetting, disabling, enabling, and eventually deleting the computer object.

Computers container

Before you create a computer object in AD DS, you must have a place to put it. When you create a domain, Windows Server create the **Computers** container by default. This container is the default location for the computer accounts when a computer joins the domain.

This container is not an organizational unit (OU). Instead, it is an object of the **Container** class. Its common name is **CN=Computers**. There are subtle but important differences between a container and an OU. You cannot create an OU within a container, so you cannot subdivide the **Computers** container. You also cannot link a Group Policy Object to a container. Therefore, we recommend that you create custom OUs to host computer objects, instead of using the **Computers** container.

AD DS forests and domains

As described briefly earlier, an Active Directory Domain Services (AD DS) forest is a collection of one or more AD DS trees that contain one or more AD DS domains. Domains in a forest share a common root, a common schema, and a global catalog. An AD DS domain is a logical administrative container for objects such as users and computers. In this topic, you'll learn more about these two important container objects in AD DS.

What is an AD DS forest?

A forest is a top-level container in AD DS. Each forest is a collection of one or more domain trees that share a common directory schema and a global catalog. A domain tree is a collection of one or more domains that share a contiguous namespace. The forest root domain is the first domain that you create in the forest. The forest root domain contains objects that do not exist in other domains in the forest. Because you always create these objects on the first domain controller, a forest can consist of as few as one domain with a single domain controller, or it can consist of several domains across multiple domain trees.

The following objects only exist in the forest root domain:

- The schema master role. This is a special, forest-wide domain controller role. Only one schema master exists in any forest. You can change the schema only on the domain controller that holds the schema master.
- The domain naming master role. This is also a special, forest-wide domain controller role. Only one domain naming master exists in any forest. Only the domain naming master can add new domain names to the directory or remove domain names from the directory.
- The **Enterprise Admins** group. By default, the **Enterprise Admins** group includes the Administrator account for the forest root domain as a member. The **Enterprise Admins** group is a member of the domain local Administrators group in every domain in the forest. This allows members of the **Enterprise Admins** group to have full control administrative rights to every domain throughout the forest.
- The **Schema Admins** group. By default, the **Schema Admins** group contains only the Administrator account from the AD DS forest root domain. Only members of the **Enterprise Admins** group or the **Domain Admins** group (in the forest root domain), can add additional members to the **Schema Admins** group. Only members of the **Schema Admins** group can change the schema.

Note: Although these objects exist initially in the root domain, you can move them to other domain controllers if required.

Security boundary

An AD DS forest is a security boundary. By default, no users from outside the forest can access any resources inside the forest. Typically, an organization creates only one forest. However, you can create multiple forests to isolate administrative permissions among different parts of the organization.

By default, all the domains in a forest automatically trust the other domains in the forest. This makes it easy to enable access to resources, such as file shares and websites, for all the users in a forest, regardless of the domain to which they belong.

Replication boundary

An AD DS forest is the replication boundary for the configuration and schema partitions in the AD DS database. As a result, all the domain controllers in the forest must share the same schema. Therefore, organizations that want to deploy applications with incompatible schemas need to deploy additional forests.

The AD DS forest is also the replication boundary for the global catalog. The global catalog makes it possible to find objects from any domain in the forest. For example, the global catalog is used whenever user principal name (UPN) sign-in credentials are used or when Microsoft Exchange Server address books are used to find users.

What is an AD DS domain?

An AD DS domain is a logical container for managing user, computer, group, and other objects. The AD DS database stores all domain objects, and each domain controller stores a copy of the database.

The AD DS database includes several types of objects. The most commonly used objects are:

- User accounts. User accounts contain information about users, including the information required to authenticate a user during the sign-in process and build the user's access token.
- Computer accounts. Each domain-joined computer has an account in AD DS. You can use computer accounts for domain-joined computers in the same way that you use user accounts for users.
- Groups. Groups organize users or computers to simplify the management of permissions and Group Policy Objects in the domain.

The AD DS domain is a replication boundary

When you make changes to any object in the domain, the domain controller where the change occurred replicates that change to all other domain controllers in the domain.

Note: The originating domain controller replicates the change to its replication partners, which in turn replicate it to their partners.

If multiple domains exist in the forest, only subsets of the changes replicate to other domains. AD DS uses a multi-master replication model that allows every domain controller to make changes to objects in the domain.

AD DS allows a single domain to contain nearly two billion objects. With this much capacity, most organizations can deploy only a single domain to ensure that all domain controllers contain all domain information. However, organizations with decentralized administrative structures or multiple locations might consider implementing multiple domains in the same forest to accommodate their administrative needs.

The AD DS domain is an administrative center

The AD DS domain contains an Administrator account and a **Domain Admins** group. By default, the Administrator account is a member of the **Domain Admins** group, and the **Domain Admins** group is a member of every local **Administrators** group of domain-joined computers. Also, by default, the **Domain Admins** group members have full control over every object in the domain. The Administrator account in the forest root domain has additional rights, as detailed earlier in this topic.

The AD DS domain provides authentication

Whenever a domain-joined computer starts or a user signs in to a domain-joined computer, AD DS authenticates it. Authentication helps to verify that the computer or user has the proper credentials for an AD DS account.

The AD DS domain provides authorization

Windows operating systems use authorization and access control technologies to allow authenticated users to access resources. Typically, authorization occurs locally at the resource level. Domain-based Dynamic Access Control enables central access rules to control the access to resources. Central access rules do not replace the current access control technology but provide an added level of control.

Trust relationships

AD DS trusts enable access to resources in a complex AD DS environment. When you deploy a single domain, you can easily grant access to resources within the domain to users and groups from the domain. When you implement multiple domains or forests, you should ensure that the appropriate trusts are in place to enable the same access to resources.

In a multiple-domain AD DS forest, two-way transitive trust relationships generate automatically between AD DS domains so that a path of trust exists between all the AD DS domains. The trusts that create automatically in the forest are all transitive trusts, which means that if domain A trusts domain B, and domain B trusts domain C, then domain A trusts domain C.

You can deploy other types of trusts. The following table describes the main trust types.

Table 1: Trusts in AD DS

Trust type	Description	Direction	Description
Parent and child	Transitive	Two-way	When you add a new AD DS domain to an existing AD DS tree, you create new parent and child trusts.
Tree-root	Transitive	Two-way	When you create a new AD DS tree in an existing AD DS forest, you automatically create a new tree-root trust.

Trust type	Description	Direction	Description
External	Nontransitive	One-way or two-way	External trusts enable resource access with a Windows NT 4.0 domain or an AD DS domain in another forest. You also can set these up to provide a framework for a migration.
Realm	Transitive or nontransitive	One-way or two-way	Realm trusts establish an authentication path between a Windows Server AD DS domain and a Kerberos version 5 (v5) protocol realm that implements by using a directory service other than AD DS.
Forest (complete or selective)	Transitive	One-way or two-way	Trusts between AD DS forests allow two forests to share resources.
Shortcut	Nontransitive	One-way or two-way	Configure shortcut trusts to reduce the time taken to authenticate between AD DS domains that are in different parts of an AD DS forest. No shortcut trusts exist by default, and an administrator must create them.

How trusts work in a forest

When you set up trusts between domains within the same forest, across forests, or with an external realm, Windows Server creates a trusted domain object to store the trusts' information, such as transitivity and type, in AD DS. Windows Server stores this trusted domain object in the **System** container in AD DS.

How trusts enable users to access resources in a forest

When a user in a domain attempts to access a shared resource in another domain in the forest, the user's computer first contacts a domain controller in its domain to request a session ticket to the resource. Because the resource is not in the user's domain, the domain controller must determine whether a trust exists with the target domain.

The domain controller can use the trust domain object to verify that the trust exists. However, to access the resource, the client computer must communicate with a domain controller in each domain along the trust path. The domain controller in the client computer's domain refers the client computer to a domain controller in the next domain along the trust path. If that is not the domain where the resource is located, that domain controller refers the client computer to a domain controller in the next domain. Eventually,

the client computer is referred to a domain controller in the domain where the resource is located, and the client is issued a session ticket to access the resource.

The trust path is the shortest path through the trust hierarchy. In a forest in which only the default trusts are configured, the trust path goes up the domain tree to the forest root domain, and then down the domain tree to the target domain. If shortcut trusts are configured, the trust path might be a single hop from the client computer domain to the domain that contains the resource.

OUs

An organizational unit (OU) is a container object within a domain that you can use to consolidate users, computers, groups, and other objects. You can link Group Policy Objects (GPOs) directly to an OU to manage the objects contained in the OU. You can also assign an OU manager and associate a COM+ partition with an OU.

You can create new OUs in Active Directory Domain Services (AD DS) by using Windows PowerShell with the Active Directory PowerShell module, and with the Active Directory Administrative Center. There are two reasons to create an OU:

- To group objects together to make it easier to manage them by applying GPOs to the whole group. When you assign GPOs to an OU, the settings apply to all the objects within the OU. GPOs are policies that administrators create to manage and configure settings for computers or users. You deploy the GPOs by linking them to OUs, domains, or sites.
- To delegate administrative control of objects within the OU. You can assign management permissions on an OU, thereby delegating control of that OU to a user or a group within AD DS, in addition to the **Domain Admins** group.

You can use OUs to represent the hierarchical, logical structures within your organization. For example, you can create OUs that represent the departments within your organization, the geographic regions within your organization, or a combination of both departmental and geographic regions. You can use OUs to manage the configuration and use of user, group, and computer accounts based on your organizational model.

Generic containers

AD DS has several built-in containers, or generic containers, such as Users and Computers. These containers store system objects or function as the default parent objects to new objects that you create. Do not confuse these generic container objects with OUs. The primary difference between OUs and containers is the management capabilities. Containers have limited management capabilities. For example, you cannot apply a GPO directly to a container.

Installing AD DS creates the Domain Controllers OU and several generic container objects by default. AD DS primarily uses some of these default objects, which are also hidden by default. The following objects are visible by default within the Active Directory Administrative Center:

- **Domain**. The top level of the domain organizational hierarchy.
- **Builtin** container. A container that stores several default groups.
- **Computers** container. The default location for new computer accounts that you create in the domain.
- **Foreign Security Principals** container. The default location for trusted objects from domains outside the AD DS forest that you add to a group in the AD DS domain.
- **Managed Service Accounts** container. The default location for managed service accounts. AD DS provides automatic password management in managed service accounts.

- **Users** container. The default location for new user accounts and groups that you create in the domain. The Users container also holds the administrator and guest accounts for the domain and for some default groups.
- Domain Controllers OU. The default location for domain controllers' computer accounts. This is the only OU that is present in a new installation of AD DS.

There are several containers that you can observe when you select **Advanced Features** on the **View** menu. The following objects are hidden by default:

- **LostAndFound**. This container holds orphaned objects.
- **Program Data**. This container holds Active Directory data for Microsoft applications, such as Active Directory Federation Services (AD FS).
- **System**. This container holds the built-in system settings.
- **NTDS Quotas**. This container holds directory service quota data.
- **TPM Devices**. This container stores the recovery information for Trusted Platform Module (TPM) devices.

Note: Containers in an AD DS domain cannot have GPOs linked to them. To link GPOs to apply configurations and restrictions, create a hierarchy of OUs and then link the GPOs to them.

Hierarchy design

The administrative needs of the organization dictate the design of an OU hierarchy. Geographic, functional, resource, or user classifications could all influence the design. Whatever the order, the hierarchy should make it possible to administer AD DS resources as effectively and flexibly as possible. For example, if you need to configure all IT administrators' computers in a certain way, you can group all the computers in an OU and then assign a GPO to manage those computers.

You also can create OUs within other OUs. For example, your organization might have multiple offices, each with its own IT administrator who is responsible for managing user and computer accounts. In addition, each office might have different departments with different computer-configuration requirements. In this situation, you can create an OU for each office, and then within each of those OUs, create an OU for the IT administrators and an OU for each of the other departments.

Although there is no limit to the number of levels in your OU structure, limit your OU structure to a depth of no more than 10 levels to ensure manageability. Most organizations use five levels or fewer to simplify administration. Note that applications that work with AD DS can impose restrictions on the OU depth within the hierarchy for the parts of the hierarchy that they use.

Suggest that students can open AD DS in Active Directory Users and Computers and display OUs while you discuss. Emphasize that the only purpose of an OU is to contain users and computers so that you can:

- Configure the objects.
- Delegate control over the objects.

AD DS schema

The Active Directory Domain Services (AD DS) schema is the component that defines all the object classes and attributes that AD DS uses to store data. All domains in a forest contain a copy of the schema that applies to that forest. Any change in the schema replicates to every domain controller in the forest via

their replication partners. However, changes originate at the schema master, which is typically the first domain controller in the forest.

AD DS stores and retrieves information from a wide variety of applications and services. It does this, in part, by standardizing how the AD DS directory stores data. By standardizing data storage, AD DS can retrieve, update, and replicate data while helping to maintain data integrity.

Objects

AD DS uses objects as units of storage. The schema defines all object types. Each time the directory manages data, the directory queries the schema for an appropriate object definition. Based on the object definition in the schema, the directory creates the object and stores the data.

Object definitions specify both the types of data that the objects can store and the syntax of the data. You can create only objects that the schema defines. Because objects store data in a rigidly defined format, AD DS can store, retrieve, and validate the data that it manages, regardless of which application supplies it.

Relationships among objects, rules, attributes, and classes

In AD DS, the schema defines the following:

- Objects that store data in the directory
- Rules that define the structure of the objects
- The structure and content of the directory itself

AD DS schema objects consist of attributes, which are grouped together into classes. Each class has rules that define which attributes are mandatory and which are optional. For example, the **user** class consists of more than 400 possible attributes, including **cn** (the common name attribute), **givenName**, **displayName**, **objectSID**, and **manager**. Of these attributes, the **cn** and **objectSID** attributes are mandatory. The **cn** attribute is a single-value Unicode string that is 1 through 64 characters long and that replicates to the global catalog.

Change the schema

Only members of the **Schema Admins** group can modify the AD DS schema. You cannot remove anything from the AD DS schema. You can only extend the AD DS schema by using AD DS schema extensions or by modifying the attributes of existing objects. For example, when you are preparing to install Microsoft Exchange Server, you must apply the Exchange Server Active Directory schema changes. These changes add or modify hundreds of classes and attributes.

You should change the schema only when necessary because the schema controls the storage of information. Any changes made to the schema affect every domain controller. Before you change the schema, you should review the changes and implement them only after you have performed testing.

This will help ensure that the changes will not adversely affect the rest of the forest or any applications that use AD DS.

The schema master is one of the operations' master roles hosted on a single domain controller in AD DS. Because it is a single master, you must use the Active Directory Schema snap-in to make changes to the schema by targeting the domain controller that holds the schema master.

Overview of AD DS replication

Within an Active Directory Domain Services (AD DS) infrastructure, standard domain controllers replicate Active Directory information by using a multi-master replication model. This means that if a change occurs on one domain controller, the change replicates to all other domain controllers in the domain, and potentially to all domain controllers throughout the entire forest.

AD DS partitions

The Active Directory data store contains information that AD DS distributes to all domain controllers throughout the forest infrastructure. Much of the information that the data store contains is distributed within a single domain. However, some information might relate to, or replicate throughout, the entire forest, regardless of the domain boundaries.

To provide replication efficiency and scalability between domain controllers, the Active Directory data is separated logically into several partitions. Each partition is a unit of replication, and each partition has its own replication topology.

The default partitions include the following types:

- **Configuration partition.** The configuration partition is created automatically when you create the first domain controller in a forest. The configuration partition contains information about the forest-wide AD DS structure, including which domains and sites exist and which domain controllers exist in each domain. The configuration partition also stores information about forest-wide services such as Dynamic Host Configuration Protocol (DHCP) authorization and certificate templates. This partition replicates to all domain controllers in the forest. It is smaller than the other partitions, and its objects do not change frequently. Therefore, replication is also infrequent.
- **Schema partition.** The schema partition contains definitions of all the objects and attributes that you can create in the data store, and the rules for creating and manipulating them. Schema information replicates to all domain controllers in the forest. Therefore, all objects must comply with the schema object and attribute definition rules. AD DS contains a default set of classes and attributes that you cannot modify. However, if you have **Schema Admins** group credentials, you can extend the schema by adding new attributes and classes to represent application-specific classes. Many applications such as Microsoft Exchange Server and Microsoft Endpoint Configuration Manager might extend the schema to provide application-specific configuration enhancements. These changes target the domain controller that contains the forest's schema master role. Only the schema master can make additions to classes and attributes. Similar to the configuration partition, the schema partition is small and needs to replicate only when changes occur to the data that is stored there. This does not happen often, except in those cases when you extend the schema.
- **Domain partition.** When you create a new domain, AD DS automatically creates and replicates an instance of the domain partition to all the domain's domain controllers. The domain partition contains information about all domain-specific objects, including users, groups, computers, OUs, and domain-related system settings. Usually, this is the largest of the AD DS partitions because it stores all the objects that the domain contains. Changes to this partition are constant because every time you create, delete, or modify an object by changing an attribute's value, AD DS automatically replicates those changes. All objects in every domain partition in a forest are stored in the global catalog with only a subset of their attribute values.
- **Application partition.** The application partition stores nondomain, application-related information that you might update frequently or that might have a specified lifetime, such as a Domain Name System (DNS) partition on domain controllers. DNS application partitions have two types: ForestDNSZones and DODomainDNSZones. They are created when you install the DNS Server role on a domain controller. An application typically is programmed to determine how it stores, categorizes, and uses applica-

tion-specific information that is stored in the Active Directory database. To prevent unnecessary replication of an application partition, you can designate which domain controllers in a forest will host the specific application's partition. Unlike a domain partition, an application partition does not store security principal objects, such as user accounts. Additionally, the global catalog does not store data that is contained in application partitions. The application partition's size and replication frequency can vary widely according to usage. Using Active Directory–integrated DNS with a large and robust DNS zone of many domain controllers, servers, and client computers will result in the frequent replication of the partition.

Note: You can use the Active Directory Services Interfaces Editor (ADSI Edit) to connect to the partitions and to review them.

Characteristics of AD DS replication

An effective AD DS replication design ensures that each partition on a domain controller is consistent with the replicas of that partition that are hosted on other domain controllers. Typically, not all domain controllers have the same information in their replicas at any particular moment because changes constantly occur to the partition. However, AD DS replication ensures that all changes to a partition transfer to all replicas of the partition. AD DS replication balances accuracy, or integrity, and consistency, or convergence, with performance. This keeps replication traffic to a reasonable level.

The key characteristics of AD DS replication are:

- Multi-master replication. Any domain controller except a read-only domain controller (RODC) can initiate and commit a change to AD DS. This provides fault tolerance and eliminates dependency on a single domain controller to maintain the directory store's operations.
- Pull replication. A domain controller requests, or pulls, changes from other domain controllers. A domain controller can notify its replication partners that it has changes to the directory or poll its partners to check if they have changes to the directory. However, the target domain controller requests and pulls the changes itself.
- Store-and-forward replication. A domain controller can pull changes from one replication partner and then make those changes available to another replication partner. For example, domain controller B can pull changes initiated by domain controller A. Then, domain controller C can pull the changes from domain controller B. This helps balance the replication load for domains that contain several domain controllers.
- Data store partitioning. A domain's domain controllers host the domain-naming context for their domains, which helps minimize replication, particularly in multiple domain forests. The domain controllers also host copies of schema and configuration partitions, which replicate forest wide. However, changes in configuration and schema partitions are much less frequent than in the domain partition. By default, other data, including application directory partitions and the partial attribute set (the global catalog), do not replicate to every domain controller in the forest. You can enable replication to be universal by configuring all the domain controllers in a forest as global catalog servers.
- Automatic generation of an efficient and robust replication topology. By default, AD DS configures an effective, multidirectional replication topology so that the loss of one domain controller does not impede replication. AD DS automatically updates this topology as you add, remove, or move domain controllers between sites.
- Attribute-level replication. When an object's attribute changes, only that attribute and minimal metadata describing that attribute replicates. The entire object does not replicate, except on its initial creation. For multiple valued attributes, such as account names in the Member of attribute of a group account, only changes to actual names replicate, and not the entire list of names.

- Distinct control of intersite replication. You can control replication between sites.
- Collision detection and management. There are only a few situations in which replication conflicts occur. Conflicts occur when:
 - You create objects with the same fully qualified domain name (FQDN) at two domain controllers within the same replication cycle.
 - On one domain controller, you delete an organizational unit (OU) and on another domain controller, you move an object to that OU within the same replication cycle.
 - You modify the same attribute of the same object on two domain controllers within the same replication cycle.
- AD DS replication always resolves the conflicts, based on metadata that replicates with the change.

AD DS sign-in process

When a computer starts, it authenticates with Active Directory Domain Services (AD DS). It searches for a domain controller by using a Domain Name System (DNS) lookup. When a user attempts to sign in to that computer, the computer attempts to contact the same domain controller it previously used to authenticate. If it fails, the computer searches for another domain controller to authenticate the user by using DNS lookup. The computer sends the user's name and password to the domain controller for authentication. The Local Security Authority (LSA) on the domain controller manages the actual authentication process.

If the sign in succeeds, the LSA builds an access token for the user that contains the security IDs (SIDs) for the user and any groups in which the user is a member. The token provides the access credentials for any process that the user initiates. For example, after signing in to AD DS, if a user attempts to open a Microsoft Word file, Word uses the credentials in the user's access token to verify the level of the user's permissions for that file.

Note: An SID is a unique string in the form S R X Y1 Y2 Yn 1 Yn. For example, a user SID can be S-1-5-21-322346712-1256085132-1900709958-500.

The following table explains the parts of this SID.

Table 1: Components of the SID

Component	Definition	In the example
S	Indicates that the string is an SID	S
R	Revision level	1
X	Identifier authority value	5 (NT Authority)
Y1-Y2-Yn-1	Domain identifier	21-322346712-1256085132-1900709958
Yn	Relative ID (RID)	500

Every user and computer account and every group that you create has a unique SID. The SIDs differ from each other only because of the unique relative ID (RID). The SID in the example is a well-known SID for the domain administrator account. Default accounts and groups use well-known SIDs. The domain administrator account's SID always ends with 500.

Although the sign-in process appears to the user as a single event, it has two parts:

- The user provides credentials, usually a user account name and password, which are checked against the AD DS database. If the user account name and password match the information stored in the AD DS database, the user becomes an authenticated user and the domain controller issues the user a

ticket-granting ticket (TGT). At this point, the user does not have access to any resources on the network.

- A secondary process in the background sends the TGT to the domain controller and requests access to the local computer. The domain controller issues a service ticket to the user, who can then interact with the local computer. At this point in the process, the user has authenticated to AD DS and signed in to the local computer.

When a user later attempts to connect to another computer on the network, the secondary process runs again, and sends the TGT to the nearest domain controller. When the domain controller returns a service ticket, the user can access the computer on the network, which generates a logon event at that computer.

Note: Remember that a domain-joined computer also signs in to AD DS when it starts. You do not notice the transaction when the computer uses its computer account name and password to sign in to AD DS. After authentication, the computer becomes a member of the **Authenticated Users** group. Although the computer logon event does not have visual confirmation in a GUI, the event log records it. Also, if you have enabled auditing, the security log of **Event Viewer** records additional events.

What's a TGT and how does Windows Server use it?

Kerberos is the protocol Windows Server uses to provide authentication in AD DS. In the original design of Kerberos, a Kerberos client on the workstation derived a master key for a user from a password that the user entered during sign in. The Kerberos client used this master key to decrypt session keys received from a Kerberos domain controller (KDC).

However, there's a problem with this approach. The Kerberos client needs the user's master key each time it decrypts a session key from the KDC. This means that either the Kerberos client stores the user's key locally on the workstation or prompts the user to provide their password for each subsequent client/server communication. The first option is not secure, and the second option can lead to a frustrating user experience.

To address this issue, the Microsoft implementation of Kerberos generates a temporary key during sign in.

Note: This key is only valid for the current logon session.

When a user signs in, Windows Server implements the following process:

1. The Kerberos client on the workstation requests a ticket from the KDC. The KDC responds by creating:
 - A logon session key
 - A ticket for a special server, the KDC's full ticket-granting service
2. One copy of the logon session key is embedded in the ticket, which is encrypted with the KDC's master key. Another copy of the logon session key is encrypted with the user's master key.
3. Next, the KDC sends both the ticket and the encrypted session key back to the client. When the client receives the reply, it decrypts the logon session key with the user's master key.

Note: The client no longer needs the key derived from the user's password because the client will now use the logon session key to decrypt its copy of any server session key it gets from the KDC.

4. The client stores the logon session key in its ticket cache along with its ticket for the KDC's full ticket-granting service.

Note: The ticket for the full ticket-granting service is called a TGT.

5. When the client requests the KDC for a ticket to a server, it presents credentials in the form of an authenticator message and a ticket, in this case a TGT, just as it would present credentials to any other service.
6. The ticket-granting service opens the TGT with its master key, extracts the logon session key for this client, and uses the logon session key to encrypt the client's copy of a session key for the server.

Overview of AD DS administration tools

Managing the Active Directory Domain Services (AD DS) environment is one of the most common tasks an IT pro performs. You typically manage your domain controllers remotely. There are several tools that you can use to manage AD DS.

Active Directory Administrative Center

The Active Directory Administrative Center provides a GUI that is based on Windows PowerShell. This enhanced interface allows you to perform AD DS object management by using task-oriented navigation, and it replaces the functionality of Active Directory Users and Computers. Tasks that you can perform by using the Active Directory Administrative Center include:

- Creating and managing user, computer, and group accounts.
- Creating and managing organizational units (OUs).
- Connecting to and managing multiple domains within a single instance of the Active Directory Administrative Center.
- Searching and filtering AD DS data by building queries.
- Creating and managing fine-grained password policies.
- Recovering objects from the Active Directory Recycle Bin.
- Managing objects that the **Dynamic Access Control** feature requires.

Windows Admin Center

Windows Admin Center is a web-based console that you can use to manage server computers and computers that are running Windows 10. Typically, you use Windows Admin Center to manage servers instead of using Remote Server Administration Tools (RSAT).

Windows Admin Center works with any browser that is compliant with modern standards, and you can install it on computers that run Windows 10 and Windows Server with Desktop Experience.

Note: You can't install Windows Admin Center on a server computer that is configured as an AD DS domain controller.

With a decreasing number of exceptions, Windows Admin Center supports virtually all current Windows Server and Windows 10 administrative functionality. However, Microsoft intends that Windows Admin Center will eventually support all the administrative functionality that is presently available through RSAT.

To use Windows Admin Center, you must first download and install it. You can download Windows Admin Center from the Microsoft download website. After downloading and installing Windows Admin Center, you must enable the appropriate TCP port on the local firewall. On a Windows 10 computer, this defaults to 6516, but you can change it during setup.

Note: Unless you are using a certificate from a trusted CA, the first time you run Windows Admin Center, it prompts you to select a client certificate. Ensure you select the certificate labeled Windows Admin Center Client.

Remote Server Administration Tools

RSAT is a collection of tools that you can download from Microsoft and install on Windows 10 computers, which enables you to manage Windows Server roles and features remotely.

Note: You don't need to download RSAT. Instead, you enable it from the Settings app. In **Settings**, search for **Manage optional features**, select **Add a feature**, and then select the appropriate RSAT tools from the returned list. Select **Install** to add the feature.

You can install the consoles available within RSAT on computers running Windows 10 or on server computers that are running the Server with Desktop Experience option of a Windows Server installation. Until the introduction of Windows Admin Center, RSAT consoles were the primary graphical tools for administering the Windows Server operating system.

Other AD DS management tools

Other management tools that you will use to perform AD DS administration include:

- Active Directory module for Windows PowerShell. The Active Directory module for Windows PowerShell supports AD DS administration, and it is one of the most important management components. Server Manager and the Active Directory Administrative Center are based on Windows PowerShell and use cmdlets to perform their tasks.
- Active Directory Users and Computers. Active Directory Users and Computers is a Microsoft Management Console (MMC) snap-in that manages most common resources, including users, groups, and computers. Although many administrators are familiar with this snap-in, the Active Directory Administrative Center replaces it and provides more capabilities.
- Active Directory Sites and Services. The Active Directory Sites and Services MMC snap-in manages replication, network topology, and related services.
- Active Directory Domains and Trusts. The Active Directory Domains and Trusts MMC snap-in configures and maintains trust relationships at the domain and forest functional levels.
- Active Directory Schema snap-in. The Active Directory Schema MMC snap-in examines and modifies the definitions of AD DS attributes and object classes. You do not need to review or change it often. Therefore, by default, the Active Directory Schema snap-in is not registered.

Demonstration: Use tools to manage objects and properties in AD DS

In this demonstration, you will learn how to:

- Navigate within the Active Directory Administrative Center.
- Perform an administrative task within the Active Directory Administrative Center.
- Create objects.
- View all object attributes.
- Use the Windows PowerShell History viewer.

Demonstration steps

1. On **SEA-ADM1**, open the **Active Directory Administrative Center**.
2. On the **navigation** pane, select **Contoso (local)**, select **Dynamic Access Control**, and then select **Global Search**.
3. In the navigation pane, switch to the **tree** view, and then expand **Contoso.com**.
4. Go to the **Overview** view.
5. Reset the password for **Contoso\Bruno** to **Pa55w.rd** so that the user does not have to change the password at the next sign-in.
6. Use the **Global Search** section to find any objects that match the **sea** search string.
7. Open the **Properties** page for **SEA-CL1**, navigate to the **Extensions** section, and then select the **Attribute Editor** tab.
8. Review the object's AD DS attributes.
9. Open the **Windows PowerShell History** pane.
10. Review the Windows PowerShell cmdlet that you used to perform the most recent task.
11. On **SEA-ADM1**, close all open windows.

Test Your Knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

What is the Active Directory Domain Services (AD DS) schema?

Question 2

Is the Computers container an organizational unit (OU)?

Deploying Windows Server domain controllers

Lesson overview

Domain controllers authenticate all users and computers in a domain. Therefore, domain controller deployment is critical for the network to function correctly. This lesson examines domain controllers, the sign-in process, and the importance of Domain Name System (DNS) in that process. In addition, this lesson discusses the purpose of the global catalog.

All domain controllers are the same, with two exceptions. Read-only domain controllers (RODCs) contain a read-only copy of the Active Directory Domain Services (AD DS) database, while other domain controllers have a read/write copy. Also, you can perform certain operations only on specific domain controllers called operations masters, which this lesson explains.

Lesson objectives

After completing this lesson, you'll be able to:

- Describe the purpose of domain controllers.
- Describe the purpose of the global catalog.
- Explain the functions of operations masters.
- Describe how to upgrade from earlier versions of AD DS.
- Describe how to clone domain controllers.
- Describe the importance of DNS and service records (SRV records).
- Explore SRV records in DNS.
- Describe operations master role transfer and seizing.
- Explain how to deploy a domain controller in Azure infrastructure as a service (IaaS).

What is a DC?

A *domain controller (DC)* is a server that stores a copy of the Active Directory Domain Services (AD DS) directory database (**Ntds.dit**) and a copy of the **SYSVOL** folder. All domain controllers except read-only domain controllers (RODCs) store a read/write copy of both **Ntds.dit** and the **SYSVOL** folder.

Note: **Ntds.dit** is the database itself, and the **SYSVOL** folder contains all the template settings and files for Group Policy Objects (GPOs).

Domain controllers use a multi-master replication process to copy data from one domain controller to another. This means that for most operations, you can modify data on any domain controller, except for an RODC. The AD DS replication service then synchronizes the changes to the AD DS database with all the other domain controllers in the domain. In Windows Server 2019, you can use only Distributed File System (DFS) replication to replicate the **SYSVOL** folders.

Note: Earlier versions of Windows Server used the file replication service (FRS) to replicate the folders, but FRS is obsolete for several versions of Windows.

Domain controllers host other services related to AD DS. These include the Kerberos authentication service, which user and computer accounts use for sign-in authentication, and the Key Distribution Center (KDC), which issues the ticket-granting ticket (TGT) to an account that signs in to the AD DS domain.

All users in an AD DS domain exist in the AD DS database. If the database is unavailable for any reason, all operations that depend on domain-based authentication will fail. As a best practice, an AD DS domain should have at least two domain controllers. This makes the AD DS database more available and spreads the authentication load during peak sign-in times.

Note: Consider two domain controllers as the absolute minimum for most enterprises to help ensure high availability and performance.

Branch office considerations

When you deploy a domain controller in a branch office that can't guarantee physical security, you can use additional measures to reduce the impact of a security breach.

One option is to deploy an RODC. The RODC contains a read-only copy of the AD DS database, and by default, it does not cache any user passwords. You can configure the RODC to cache the passwords for users in the branch office. If an RODC is compromised, the risk of potential loss of information is much lower than with a full read/write domain controller.

Another option is to use BitLocker Drive Encryption to encrypt the domain controller's hard drive. If someone steals the hard drive, BitLocker will help to ensure that a malicious hacker has difficulty getting any useful information from it.

Note: BitLocker is a drive-encryption feature that is available for Windows Server operating systems and Windows client operating systems. BitLocker encrypts the entire drive to help prevent the computer from starting unless it receives a private key and (optionally) passes an integrity check. The hard drive remains encrypted even if you transfer it to another computer.

What is the global catalog?

The *global catalog* is a partial, read-only, searchable copy of all the objects in the forest. The global catalog can help speed up searches for objects that might be stored on domain controllers in a different domain in the forest.

Within a single domain, the AD DS database on each domain controller contains all the information about every object in that domain. However, only a subset of this information replicates on the global catalog servers in other domains in the forest. Within a domain, a query for an object is directed to one of the domain controllers in that domain, but that query does not include results about objects in other domains in the forest. For a query to include results from other domains in the forest, you must query a domain controller that is a global catalog server.

Note: By default, all domain controllers are also global catalogs.

The global catalog does not contain all the attributes for each object. Instead, the global catalog maintains the subset of attributes that are most likely to be useful in cross-domain searches. These attributes include **givenName**, **displayName**, and **mail**. You can change the set of attributes replicated to the global catalog by modifying the partial attribute set (PAS) in the schema.

In a multidomain forest, searching the global catalog can be useful in many situations. For example, when a server that is running Microsoft Exchange Server receives an incoming email, it must search for the recipient's account so that it can decide how to route the message. By automatically querying the global catalog, the server can find the recipient in a multidomain environment. In addition, when users sign in to their Active Directory accounts, the domain controller that performs the authentication must contact the global catalog to check for universal group memberships before authenticating the users.

In a single domain, you should configure all the domain controllers to hold a copy of the global catalog. However, in a multidomain environment, the infrastructure master should not be a global catalog server unless all the domain controllers in the domain are also global catalog servers.

When you have multiple sites, you should also make at least one domain controller at each site a global catalog server, so that you are not dependent on other sites when you require global catalog queries. Deciding which domain controllers to configure to hold a copy of the global catalog depends on replication traffic and network bandwidth. Many organizations opt to make every domain controller a global catalog server.

What are operations masters?

Certain operations can be performed only by a specific role, on a specific domain controller. A domain controller that holds one of these roles is an **operations master**. An **operations master** role is also known as a **flexible single master operations (FSMO)** role. Five operations master roles exist. You can find all five on a single domain controller or spread them across several domain controllers.

By default, the first domain controller installed in a forest has all five roles. However, you can move these roles after building more domain controllers. By allowing changes only on a single domain controller, the **operations master** roles help to prevent conflicts in Active Directory Domain Services (AD DS) because of replication latency. When making changes to data on an **operations master**, you must connect to the domain controller that holds the role.

The five **operations master** roles have the following distribution:

- Each forest has one **schema master** and one **domain naming master**.
- Each AD DS domain has one **relative ID (RID) master**, one **infrastructure master**, and one **primary domain controller (PDC) emulator**.

Forest operations masters

A forest has the following single master roles:

- **Domain naming master.** This is the domain controller that you must contact when you add or remove a domain or make domain name changes.
If the domain naming master is unavailable, you will not be able to add domains to the forest.
- **Schema master.** This is the domain controller in which you make all schema changes. To make changes, you typically sign in to the **schema master** as a member of both the **Schema Admins** and the **Enterprise Admins** groups. A user who is a member of both groups and who has the right permissions can also edit the schema by using a script.
If the **schema master** is unavailable, you will not be able to make changes to the schema. This prevents the installation of applications that require schema changes, such as Microsoft Exchange Server.

Note: The Windows PowerShell command **Get-ADForest**, from the Active Directory module for Windows PowerShell, displays the forest properties, including the current **domain naming master** and **schema master**.

Domain operations masters

A domain has the following single master roles:

- **RID master.** Whenever you create an object in AD DS, the domain controller where you created the object assigns the object a unique identifying number known as a SID. To ensure that no two domain

controllers assign the same SID to two different objects, the RID master allocates blocks of RIDs to each domain controller within the domain to use when building SIDs. If the RID master is unavailable, you might experience difficulties adding new objects to the domain. As domain controllers use their existing RIDs, they eventually run out of them and are unable to create new objects.

- **Infrastructure master.** This role maintains interdomain object references, such as when a group in one domain has a member from another domain. In this situation, the infrastructure master handles maintaining the integrity of this reference. For example, when you review the **Security** tab of an object, the system references the listed SIDs and translates them into names. In a multidomain forest, the **infrastructure master** references SIDs from other domains.

If the **infrastructure master** is unavailable, domain controllers that are not global catalogs will not be able to check universal group memberships or authenticate users.

The **infrastructure master** role should not reside on a global catalog server unless you have a single-domain forest. The exception is when you follow best practices and make every domain controller a global catalog. In that case, the **infrastructure master** role is not necessary, because every domain controller knows about every object in the forest.

- **PDC emulator master.** The domain controller that holds the **PDC emulator master** is the time source for the domain. The **PDC emulator master** in each domain in a forest synchronize their time with the **PDC emulator master** in the forest root domain. You set the **PDC emulator master** in the forest root domain to synchronize with a reliable external time source.

The **PDC emulator master** is also the domain controller that receives urgent password changes. If a user's password changes, the domain controller holding the **PDC emulator master** role receives this information immediately. This means that if the user tries to sign in, the domain controller in the user's current location will contact the domain controller holding the **PDC emulator master** role to check for recent changes. This will occur even if a domain controller, in a different location that had not yet received the new password information, authenticated the user.

If the **PDC emulator master** is unavailable, users might have trouble signing in until their password changes have replicated to all the domain controllers.

The **PDC emulator master** also plays a role in editing GPOs. When you open a GPO other than a local GPO for editing, the **PDC emulator master** stores the edited copy. This prevents conflicts if two administrators attempt to edit the same GPO at the same time on different domain controllers. However, you can choose to use a specific domain controller to edit the GPOs. This is especially useful when editing GPOs in a remote office with a slow connection to the **PDC emulator master**.

Note: The Windows PowerShell command **Get-ADDomain**, from the Active Directory module for Windows PowerShell, displays the domain properties, including the current **RID master**, **infrastructure master**, and **PDC emulator master**.

Note: The global catalog is not one of the operations master roles.

Install a DC

Install a domain controller from Server Manager

The domain controller installation and promotion process has two steps. First, you install the files that the domain controller role uses by using **Server Manager** to install the Active Directory Domain Services (AD DS) role. At the end of the initial installation process, you have installed the AD DS files but not yet configured AD DS on the server.

The second step is to configure AD DS by using the **Active Directory Domain Services Configuration Wizard**. You start the wizard by selecting the AD DS link in **Server Manager**. The wizard allows you to do one of the following:

- Add a domain controller to an existing domain.
- Add a new domain to an existing forest.
- Add a new forest.

Before installing a new domain controller, you should answer the questions in the following table.

Table 1: Planning to deploy a domain controller

Question	Comments
Are you installing a new forest, a new tree, or an additional domain controller for an existing domain?	Answering this question determines what additional information you might need, such as the parent domain name.
What is the DNS name for the AD DS domain?	When you create the first domain controller for a domain, you must specify the fully qualified domain name (FQDN). When you add a domain controller to an existing domain or forest, the wizard provides the existing domain information.
Which level will you choose for the forest functional level?	The forest functional level determines the available forest features and the supported domain controller operating system. This also sets the minimum domain functional level for the domains in the forest.
Which level will you choose for the domain functional level?	The domain functional level determines the domain features that will be available and the supported domain controller operating system.
Will the domain controller be a DNS server?	Your DNS must be functioning well to support AD DS.
Will the domain controller host the global catalog?	This option is selected by default for the first domain controller in a forest.
Will the domain controller be an RODC?	This option is not available for the first domain controller in a forest.
What will be the Directory Services Restore Mode (DSRM) password?	This is necessary for recovering the AD DS database from a backup.
What is the NetBIOS name for the AD DS domain?	When you create the first domain controller for a domain, you must specify the NetBIOS name for the domain.
Where will the database, log files and SYSVOL folders be created?	By default, the database and log files folder is C:\Windows\NTDS. By default, the SYSVOL folder is C:\Windows\SYSVOL.

Install a domain controller on a Server Core installation of Windows Server

A Windows Server computer that is running a Server Core installation does not have the **Server Manager** GUI, so you must 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 any supported version of Windows Server that has Desktop Experience or any Windows client such as Windows 8.1 or Windows 10.

To install the AD DS files on the server, you can do one of the following:

- Use **Server Manager** to connect remotely to the server running the Server Core installation, and then install the AD DS role as described in the previous topic.
- Use the Windows PowerShell command **Install-WindowsFeature AD-Domain-Services** to install the files.

After you install the AD DS files, you can complete the rest of the configuration process 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 Windows PowerShell cmdlet **Install-ADDSDomainController**, supplying the required information on the command line.

Note: You can use **dcpromo.exe** to promote a domain controller on Server Core.

Note: In Windows Server, running a cmdlet automatically loads the cmdlet's module, if it is available. For example, running the **Install-ADDSDomainController** cmdlet automatically loads the **ADDSDeployment** module into your current Windows PowerShell session. If a module is not loaded or available, you will receive an error message when you run the cmdlet to indicate that it is not a valid cmdlet.

You can still manually import the module that you need. However, in Windows Server, you do so only when necessary, such when you are pointing to a source to install the module.

Install a domain controller by installing from media

If you have a network connection between sites 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 installing it from media rather than by 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 AD DS database and the **SYSVOL** folder will copy to the new domain controller over a potentially unreliable WAN connection. Alternatively, 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 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 occurs locally. In this scenario, only security-related traffic uses the wide-area network (WAN) link. The WAN link also helps ensure that the new domain controller receives any changes made to the central AD DS after you created the **Install from media** backup.

To install a domain controller by installing from media, browse to a domain controller that is not an RODC. Use the **ntdsutil** command-line tool to create a snapshot of the AD DS 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 providing the local path to the **Install from media** directory that you previously created.

The procedure is as follows:

1. On the full domain controller, at an administrative command prompt, enter the following commands, where **C:\IFM** is the destination directory that will contain the snapshot of the AD DS database:

```
Ntdstil
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:
 1. Use **Server Manager** to add the **AD DS** role.
 2. Wait while the AD DS files install.
 3. In **Server Manager**, select the **Notification** icon, and then under **Post-Deployment Configuration**, select **Promote this server to a domain controller**. The **Active Directory Domain Services Configuration Wizard** runs.
 4. On the appropriate page of the wizard, select the **Install from media** option, and then provide the local path to the snapshot directory. AD DS installs from the snapshot.
3. Note that when the domain controller restarts, it contacts the other domain controllers in the domain and updates AD DS with any changes made after the creation of the snapshot.

Note: You can use the option to install from media (IFM) only when you are adding additional domain controllers to an existing domain. You can't use IFM when you are creating a new domain in the forest or when you are creating a new forest.

Upgrade from a previous version of AD DS

The process for upgrading a domain controller is the same for any version of Windows Server starting from Windows Server 2012 R2 through Windows Server 2019. You can upgrade to a Windows Server 2019 domain in one of the following two ways:

- You can upgrade the operating system on existing domain controllers that are running Windows Server 2012 R2 or later.
- You can add servers running Windows Server 2019 as domain controllers in a domain that already has domain controllers running earlier versions of Windows Server.

We recommend the latter method, because when you finish, you will have a clean installation of the Windows Server 2019 operating system and the Active Directory Domain Services (AD DS) database. Whenever you add a new domain controller, Windows Server updates the domain DNS records, and clients will immediately find and use this domain controller.

Upgrade to Windows Server 2019

To upgrade an AD DS domain running at the functional level of an earlier version of Windows Server to an AD DS domain running at the functional level of Windows Server 2019, you must first upgrade all the domain controllers to the Windows Server 2019 operating system. You can perform this upgrade by upgrading all the existing domain controllers to Windows Server 2019. Alternatively, you can introduce new domain controllers that are running Windows Server 2019 and then phase out the existing domain controllers.

Although an in-place operating system upgrade performs automatic schema and domain preparation, you can perform these steps manually. If you want to prepare the forest and domain manually prior to upgrading the domain controllers, you can do so by using the command-line commands **adprep.exe /forestprep** and **adprep.exe /domainprep**.

Note: The installation media in the **\Support\Adprep** folder includes **adprep.exe**.

There are no additional configuration steps after that point, and you can continue to run the Windows Server operating system upgrade.

When you promote a server running Windows Server 2019 to be a domain controller in an existing domain, and you have signed in as a member of the **Schema Admins** and **Enterprise Admins** groups, the AD DS schema automatically updates to Windows Server. In this scenario, you do not need to run the **adprep.exe** command before you start the installation.

Deploying Windows Server 2019 domain controllers

To upgrade the operating system of a domain controller running Windows Server 2012 R2 or later to Windows Server 2019, perform the following steps:

1. Insert the installation media for Windows Server 2019, and then run **Setup**. The **Windows Setup Wizard** opens.
2. Complete the **Windows Setup Wizard** and select the **Upgrade: Install Windows and keep files, settings, and applications** option.

Note: With this type of upgrade, you do not need to preserve users' settings and reinstall applications because the upgrades occur in place. Remember to check for hardware and software compatibility before you perform an upgrade.

To introduce a clean installation of Windows Server 2019 as a domain controller, perform the following steps:

1. Deploy and configure a new installation of Windows Server 2019, and then join it to the domain.
2. Promote the new server to be a domain controller in the domain by using Server Manager or one of the other methods previously described.

DC cloning

The fastest way to deploy multiple computers with identical configurations, especially when those computers run in a virtualized environment such as Microsoft Hyper-V, is to clone those computers. Cloning copies the virtual hard disks (VHDs) of the computers and changes minor configurations such as computer names and IP addresses to be unique, which makes the computers immediately operational. This process, also referred to as provisioning computers, is a central technology of private clouds. In Windows Server 2012 and newer, you can clone domain controllers. The following scenarios benefit from virtual domain controller cloning:

- Rapidly deploying additional domain controllers.
- Quickly restoring business continuity during disaster recovery. You can restore Active Directory Domain Services (AD DS) capacity by using cloning to quickly deploy domain controllers.
- Optimizing private cloud deployments. You can take advantage of the flexible provisioning of domain controllers to accommodate increased scale requirements.
- Rapidly provisioning test environments. This allows for the deployment and testing of new features and capabilities before a production rollout.
- Quickly meeting increased capacity needs in branch offices. You can do this either by cloning existing domain controllers in branch offices or by cloning them in the datacenter, and then transferring them to branches by using Hyper-V.

To clone domain controllers, you will require the following:

- A hypervisor that supports virtual machine generation identifiers, such as Hyper-V in Windows Server 2012 and later.
- Domain controllers as guest operating systems based on Windows Server.
- A domain controller that you want to clone, or a source domain controller, which must run as a virtual machine guest on the supported hypervisor.
- A primary domain controller (PDC) emulator that runs on Windows Server 2012 or later. However, the domain controller that holds the PDC emulator operations master role must support the cloning process. The PDC emulator must be online when the virtual domain controller clones start for the first time.

To help ensure that AD DS administrators authorize cloning virtualized domain controllers, a member of the **Domain Admins** group should prepare a computer for cloning. Hyper-V administrators cannot clone a domain controller without AD DS administrators, and similarly AD DS administrators cannot clone a domain controller without Hyper-V administrators.

Prepare the source virtual domain controller

To prepare to deploy virtual domain controllers, follow these steps:

1. Add the source domain controller to the **Cloneable Domain Controllers** group.
2. Verify that the apps and services on the source domain controller support the cloning process. You can do this by running the following Windows PowerShell cmdlet:

```
Get-ADDCCloneingExcludedApplicationList
```

- You must remove or test any apps or services that do not support cloning. If they work after cloning, put the apps or services in the **CustomDCCloneAllowList.xml** file.
- You can create **CustomDCCloneAllowList.xml** by using the same cmdlet and appending the *GenerateXML* parameter. Optionally, you can append the *-Force* parameter if you want to overwrite an existing **CustomDCCloneAllowList.xml** file, as the following syntax demonstrates:

```
Get-ADDCCloneingExcludedApplicationList -GenerateXML [-Force]
```

3. Create a **DCCloneConfig.xml** file. You must create this file so that the cloning process recognizes it and creates a new domain controller from the clone. By creating this file, you can specify a custom computer name, TCP/IP address settings, and the site name where the new domain controller should reside. If you do not specify one or all of these parameters, Windows Server generates a computer name automatically and sets the IP address settings to dynamic. This requires a Dynamic Host Configuration Protocol (DHCP) server on the network and assumes that the domain controller clones reside in the same site as the source domain controller. You can use Windows PowerShell to create the **DCCloneConfig.xml** file, as the following syntax demonstrates:

```
New-ADDCCloneConfigFile [-CloneComputerName <String>] [-IPv4DNSResolver <String[]>] [-Path <String>] [-SiteName <String>]
```

Note: If you want to create more than one clone and you want to specify settings such as computer names and TCP/IP addressing information, you must modify the **DCCloneConfig.xml** file. Alternatively, you can create a new, individual one for each clone prior to starting it for the first time.

4. Export the source virtual domain controller.

Prepare multiple domain controller clones

If you want to prepare multiple domain controller clones, do not provide any additional parameters, and allow the automatic generation of the computer name. In addition, use DHCP to provide TCP/IP addressing information.

Alternatively, you can customize each clone by creating individual **DCCloneConfig.xml** files. To do this, follow these steps:

1. Create the cloned VHDs by exporting and importing the virtual computer.
2. Mount the newly cloned VHDs by doing one of the following:
 - In File Explorer, double-click the cloned VHD or select the cloned VHD, and then select Enter.
 - Use **Diskpart.exe** with the **assign** command at an elevated command prompt.
 - Use the **Mount-DiskImage** Windows PowerShell cmdlet.
3. Use the *-Offline* and *-Path* parameters with the **New-ADDCCloneConfigFile** cmdlet. Change *E* to the drive letter that you used when mounting the **.vhdx** file in the previous step, as the following cmdlet indicates:

```
New-ADDCCloneConfigFile -CloneComputerName <SEA-DC3> -Offline -Path <E>:\Windows\NTDS
```

4. Unmount the VHD files by using **Diskpart.exe** or the **Dismount-DiskImage** Windows PowerShell cmdlet.

Use dynamically assigned computer names

If you do not configure **DCCloneConfig.xml** with a static computer name—for example, to create multiple clones without individual configurations—the computer name of the new clone is automatically generated based on the following algorithm:

- The prefix consists of the first eight characters of the computer name of the source domain controller. For example, the source computer name **SourceComputer** abbreviates to the prefix **SourceCo**.
- A unique naming suffix of the format *-CLnnnn* appends to the prefix, where *nnnn* is the next available value from 0001 through 9999 that the **PDC emulator** determines is not currently in use.

Finalize the domain controller cloning

When a new domain controller clone starts, the following steps take place automatically:

1. The clone checks whether a virtual machine generation identifier exists. This is required, and if a virtual machine generation identifier does not exist, the computer either starts normally when no **DCCloneConfig.xml** exists or renames **DCCloneConfig.xml** and restarts in Directory Services Restore Mode (DSRM). Starting in DSRM is a safeguard, and a domain administrator should pay close attention and fix the issue to make the domain controller work as intended.

2. The clone checks whether the virtual machine generation identifier changed and performs one of the following actions:
 - If it did not change, it is the original source domain controller. If **DCCloneConfig.xml** exists, then it is renamed. In both cases, a normal startup occurs, and the domain controller is functional again.
 - If it did change, the virtualization safeguards trigger, and the process continues.
3. The clone checks whether **DCCloneConfig.xml** exists. If not, a check for a duplicate IP address determines whether the computer starts normally or in DSRM. If the **DCCloneConfig.xml** file exists, the computer gets the new computer name and IP address settings from the file. The AD DS database is modified, and the initialization steps continue, thereby creating a new domain controller.

Overview of DC SRV records

When users sign in to Active Directory Domain Services (AD DS), their computers examine DNS for service records (SRV records) to locate the nearest domain controller. SRV records specify information about available services. Each domain controller dynamically registers its addresses in DNS at startup by registering an updated SRV record in DNS. Clients can locate a suitable domain controller to service their sign-in requests by using DNS lookups, which use these SRV records.

SRV records for AD DS follow the following pattern:

```
_Service._Protocol.DomainName
```

For example, if a client wants to locate a server that is running the Lightweight Directory Access Protocol (LDAP) service in the `Adatum.com` domain, it queries for **_ldap._tcp.Adatum.com**.

Sites and SRV records

A client uses sites when it needs to contact a domain controller. It starts by examining SRV records in DNS. The response to the DNS query includes:

- A list of the domain controllers in the same site as the client.
- A list of the domain controllers from the next closest site that does not include a read-only domain controller (RODC), if no domain controllers were available in the same site and the **Try Next Closest Site Group Policy** setting is enabled.
- A random list of available domain controllers in the domain if there is no domain controller in the next closest site.

Administrators can define sites in AD DS. When you define sites, you should consider which parts of the network have good connectivity and bandwidth. For example, if a branch office connects to the main datacenter through an unreliable wide area network (WAN) link, you should define the branch office and the datacenter as separate sites.

The Net Logon service that runs on each domain controller registers the SRV records in DNS. If the SRV records are not entered in DNS correctly, you can trigger the domain controller to reregister those records by restarting the Net Logon service on that domain controller. Note that this process reregisters only the SRV records. If you want to reregister the host (A) record information in DNS, you must run **ipconfig /registerdns** from a command prompt, just as you would for any other computer.

Demonstration: Explore DC SRV records in DNS

In this demonstration, you will explore domain controller (DC) service records (SRV records) in Domain Name System (DNS).

Demonstration steps

Review the SRV records by using DNS Manager

1. On **SEA-ADM1**, sign in with the username **Contoso\Administrator** and the password **Pa55w.rd**.
2. Open the **DNS Manager** window, and then explore the DNS domains that begin with an underscore (.).
3. Observe the service records (SRV records) that domain controllers have registered.

Note: These records provide alternate paths so that clients can discover them.

Transfer and seize roles

In an Active Directory Domain Services (AD DS) environment where you distribute operations master roles among domain controllers, you might need to move a role from one domain controller to another. When you plan this role move—for example, to decommission servers or balance workloads—the move is known as transferring the role. If you do not plan the move—for example, in the case of a hardware or system failure—the move is known as seizing the role. When you transfer a role, the latest data from the domain controller in that role replicates to the target server. You should seize a role only as a last resort.

Transfer operations master roles

You can transfer operations master roles through the GUI by using the AD DS snap-ins that the following table lists.

Table 1: Appropriate tools for operations master transfer

Role	Snap-in
Schema master	Active Directory Schema
Domain naming master	Active Directory Domains and Trusts
Infrastructure master	Active Directory Users and Computers
RID master	Active Directory Users and Computers
PDC emulator	Active Directory Users and Computers

Seize operations master roles

You cannot use the snap-ins to seize operations master roles. Instead, you must use the **ntdsutil.exe** command-line tool or Windows PowerShell to seize roles. You can also use these tools to transfer roles. The syntax for transferring a role and seizing a role is similar within Windows PowerShell, as the following syntax line demonstrates:

```
Move-ADDirectoryServerOperationsMasterRole -Identity "< servername >"
-OperationsMasterRole < rolenamelist > -Force
```

For the preceding syntax, the noteworthy definitions are as follows:

- `< servername >`. The name of the target domain controller to which you are transferring one or more roles.
- `< rolenamelist >`. A comma-separated list of AD DS role names to move to the target server.
- `-Force`. An optional parameter that you include to seize a role instead of transferring it.

Deploy a DC in Azure IaaS

Microsoft Azure provides infrastructure as a service (IaaS), which is virtualization in the cloud. All the considerations for virtualizing applications and servers in an on-premises infrastructure apply to deploying the same applications and servers in Azure.

When deploying Active Directory Domain Services (AD DS) on Azure IaaS, you are installing the domain controller on a virtual machine, so all the rules that apply to virtualizing a domain controller apply to deploying AD DS in Azure. You can install AD DS on Azure virtual machines to support a variety of scenarios, which include:

- Disaster recovery. In a scenario in which your on-premises domain controllers are destroyed or are otherwise unavailable, Azure-based virtual machines that are running as replica domain controllers will have a complete copy of your AD DS database. This can help speedy recovery and is a low-cost alternative for organizations that do not have a physical disaster recovery site.
- Geo-distributed domain controllers. If your organization is highly decentralized, Azure-based virtual machines that are running as replica domain controllers can provide lower latency connections for improved authentication performance. You can achieve this by running domain controllers in different Azure regions that correspond to the locations where it is not cost effective for your organization to deploy physical infrastructure.
- User authentication for isolated applications. If you need to deploy an application with an AD DS dependency, but that application does not require connectivity with the organizational AD DS environment, you could deploy a separate forest on Azure virtual machines.

Note: Although on-premises member servers and clients can communicate with Azure-based domain controllers, these domain controllers should never be the only domain controllers in a hybrid environment. Loss of connectivity between an on-premises environment and Azure prevents authentication and other domain functions if you are not also running AD DS services in your on-premises environment.

When you implement AD DS in Azure, consider the following:

- Network topology. To meet AD DS requirements, you must create an Azure Virtual Network and attach your virtual machines to it. If you intend to join an existing on-premises AD DS infrastructure, you can opt to extend network connectivity to your on-premises environment. You can achieve this through a standard virtual private network (VPN) connection or an Azure ExpressRoute circuit, depending on the speed, reliability, and security that your organization requires.

Note: An ExpressRoute circuit is a method of connecting an on-premises infrastructure to Microsoft cloud services through a dedicated connectivity provider that does not use the public internet.

- Site topology. As with a physical site, you should define and configure an AD DS site that corresponds to the IP address space of your Azure Virtual Network. Because the use of an Azure Virtual Network incurs additional gateway costs for all outbound traffic to your on-premises environment, you should carefully plan your AD DS sites and site links to minimize cost. Because AD DS site link transitivity is enabled by default, you should consider disabling the option to bridge all site links if you have more than two sites. If you leave site link bridging enabled, AD DS assumes that all sites in your deployment

have direct connectivity with one another, which might result in your Azure AD DS site having multiple replication partners.

Ensure that you do not enable change notification on site links that contain your Azure AD DS site. If you enable change notification, it will override any replication intervals that are configured on the site link, resulting in frequent and often unnecessary replication. If a writable copy of AD DS is not necessary, you should consider deploying a read-only domain controller (RODC) to further limit the amount of outbound traffic that AD DS replication creates.

- Service healing. Domain controller replication depends on the update sequence number (USN). When an AD DS system rolls back, Windows Server could create duplicate USNs.

To prevent this, Windows Server uses an identifier named VM-Generation ID. VM-Generation ID can detect a rollback and prevent a virtualized domain controller from replicating changes outbound until the virtualized AD DS has converged with the other domain controllers in the domain.

Note: Azure virtual machines that are running the domain controller role should always be shut down through the guest operating system and never through the Azure portal. Initiating a shutdown through the Azure portal deallocates the virtual machine, causing a reset of the VM-Generation ID identifier.

- IP addressing. All Azure virtual machines receive Dynamic Host Configuration Protocol (DHCP) addresses by default, but you can configure static addresses through Azure PowerShell that will persist across restarts, shutdowns, and service healing. Azure virtual machines that are to host a domain controller role, a Domain Name System (DNS) role, or both should have the initial dynamic IP address configured as static by using the **Set-AzureStaticVNetIP** cmdlet so that the IP never deallocates if the virtual machine shuts down. You must first provision the Azure Virtual Network before you provision the Azure-based domain controllers.
- DNS. Azure's built-in DNS does not meet the requirements of AD DS, such as Dynamic DNS and service (SRV) resource records. Before you can extend your on-premises AD DS environment to an Azure virtual machine, you must provision and configure the Azure Virtual Network to an on-premises DNS server.
- Disks. Azure virtual machines use read/write host caching for operating system (OS) virtual hard disks. Although this can improve virtual machine performance, if AD DS components are installed on the OS disk, data loss is possible if there is a disk failure. You can turn off caching in additional Azure hard disks that are attached to a virtual machine. When you install AD DS in Azure, you should put the **NTDS.DIT** and **SYSVOL** folders on an additional data disk on the Azure virtual machine with the Host Cache Preference setting configured to **NONE**. However, keep in mind that Azure data disks have a maximum size of 32 terabytes (TBs).

Test Your Knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

What's a domain controller?

Question 2

What is the primary domain controller (PDC) Emulator?

Overview of Azure AD

Lesson overview

Microsoft Azure Active Directory (Azure AD) is part of the platform as a service (PaaS) offering and operates as a directory service that Microsoft manages in the cloud. You can use it to provide authentication and authorization for cloud-based services and apps offered by Microsoft. In this lesson, you will learn about the features of Azure AD, how it differs from Active Directory Domain Services (AD DS), and how to implement synchronization in hybrid environments.

Lesson objectives

After completing this lesson, you'll be able to:

- Describe Azure AD.
- Differentiate between the different versions of Azure AD.
- Explain how to connect AD DS with Azure AD by using Azure AD Connect.
- Identify the benefits of AD DS and Azure AD hybrid configurations.

What is Azure AD?

Microsoft Azure Active Directory (Azure AD) is not a part of the 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 don't have to dedicate resources to its deployment or maintenance.

With Azure AD, you have access to a set of features that aren't natively available in Active Directory Domain Services (AD DS), such as support for multifactor authentication, identity protection, and self-service password reset.

You can use Azure AD to provide more secure access to cloud-based resources for organizations and individuals by:

- Configuring access to applications.
- Configuring single sign-on (SSO) to cloud-based software-as-a-service (SaaS) applications.
- Managing users and groups.
- Provisioning users.
- Enabling federation between organizations.
- Providing an identity management solution.
- Identifying irregular sign-in activity.
- Configuring multifactor authentication.
- Extending existing on-premises Active Directory implementations to Azure AD.
- Configuring Application Proxy for cloud and local applications.
- Configuring conditional access for users and devices.

Azure AD is a separate Azure service. Its most basic service is the free tier, which any new Azure subscription automatically includes. If you subscribe to any Microsoft Online business services, such as Microsoft Office 365 or Microsoft Intune, you automatically get Azure AD with access to all the free features.

Note: By default, when you create a new Azure subscription by using a Microsoft account, the subscription automatically includes a new Azure AD tenant with a default directory.

The more advanced identity management features require paid versions of Azure AD, which are offered in the form of free and premium tiers. Azure AD includes some of these features as part of Office 365 subscriptions.

Note: You will learn about differences between Azure AD versions later in this lesson.

Implementing Azure AD is not the same as deploying virtual machines in Azure, adding AD DS, and then deploying domain controllers for a new forest and domain. Azure AD is a different service, focused on providing identity management services to web-based apps, unlike AD DS, which focuses on on-premises apps.

Azure AD vs. AD DS

You could consider Azure AD simply as the cloud-based counterpart of AD DS. However, while Azure AD and AD DS share some common characteristics, there are several significant differences between them.

Characteristics of AD DS

AD DS is the traditional deployment of Windows Server-based Active Directory on a physical or virtual server. Although AD DS is commonly considered to be primarily a directory service, it's only one component of the Windows Active Directory suite of technologies, which also includes Active Directory Certificate Services (AD CS), Active Directory Lightweight Directory Services (AD LDS), Active Directory Federation Services (AD FS), and Active Directory Rights Management Services (AD RMS).

When comparing AD DS with Azure AD, it's important to note the following characteristics of AD DS:

- AD DS is a true directory service, with a hierarchical X.500-based structure.
- AD DS uses Domain Name System (DNS) for locating resources such as domain controllers.
- You can query and manage AD DS by using Lightweight Directory Access Protocol (LDAP) calls.
- AD DS primarily uses the Kerberos protocol for authentication.
- AD DS uses organizational units (OUs) and Group Policy Objects (GPOs) for management.
- AD DS includes computer objects, representing computers that join an Active Directory domain.
- AD DS uses trusts between domains for delegated management.

You can deploy AD DS on an Azure virtual machine to enable scalability and availability for an on-premises AD DS. However, deploying AD DS on an Azure virtual machine does not make any use of Azure AD.

Note: Deploying AD DS on an Azure virtual machine requires one or more additional Azure data disks, because you should not use drive **C** for AD DS storage. Windows Server uses these disks to store the AD DS database, logs, and **SYSVOL**. You must set the **Host Cache Preference** setting for these disks to **None**.

Characteristics of Azure AD

Although Azure AD has similarities to AD DS, there are also many differences. It's important to realize that using Azure AD isn't the same as deploying an Active Directory domain controller on an Azure virtual machine and adding it to your on-premises domain.

When comparing Azure AD with AD DS, it's important to note the following characteristics of Azure AD:

- Azure AD is primarily an identity solution, and it's designed for internet-based applications by using HTTP (port 80) and HTTPS (port 443) communications.
- Azure AD is a multi-tenant directory service.
- Azure AD users and groups are created in a flat structure, and there are no OUs or GPOs.
- You cannot query Azure AD by using LDAP. Instead, Azure AD uses the REST API over HTTP and HTTPS.
- Azure AD does not use Kerberos authentication. Instead, it uses HTTP and HTTPS protocols such as Security Assertion Markup Language (SAML), WS-Federation, and OpenID Connect for authentication, and uses OAuth for authorization.
- Azure AD includes federation services, and many third-party services are federated with Azure AD and trust it.

Azure AD authentication protocols

Azure AD authentication protocols differ from AD DS authentication protocols. Often, AD DS administrators are less experienced with web-based authentication protocols. Azure AD supports 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 temporary. OAuth allows decoupling of the authentication credentials, so that credentials don't pass to the destination.
- SAML 2.0. SAML is an open standard XML protocol that consists of security tokens. A security token contains claims, which are typically Active Directory attributes 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. For instance, they are all 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.

Azure AD Join

Joining your organization's devices to your AD DS 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 Windows devices from the cloud by using mobile device management (MDM) instead of using GPOs or with Microsoft Endpoint Configuration Manager.

Usage scenarios

When determining whether to implement Azure AD Join, consider the following scenarios:

- Your organization's apps and resources are mostly cloud-based. If your organization currently uses or is planning to use SaaS apps, such as 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, which connects 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 experience 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.

Azure AD versions

Microsoft Azure Active Directory (Azure AD) has four editions: Free, Office 365 apps, Premium P1, and Premium P2. Microsoft includes the Free edition with a subscription of a commercial online service such as Microsoft Azure, Microsoft Dynamics 365, Microsoft Intune, or Microsoft Power Platform. Office 365 subscriptions include the Free edition, but Office 365 E1, E3, E5, and F1 subscriptions also include the features listed in the Office 365 apps column. The Azure Active Directory Premium P1 and P2 editions provide additional features for enterprise users.

The following table identifies the key differences between the editions of Azure AD.

Table 1: Azure AD editions

Feature	Free	Office 365 Apps	Premium P1	Premium P2
Number of objects	500,000	Unlimited	Unlimited	Unlimited
Single sign-on	Up to 10 apps	Up to 10 apps	Unlimited	Unlimited
Azure AD Connect	Yes	Yes	Yes	Yes
Device registration	Yes	Yes	Yes	Yes
User and group management	Yes	Yes	Yes	Yes
Self-service password change	Yes	Yes	Yes	Yes
Multifactor authentication	Yes	Yes	Yes	Yes
Device write-back (device objects two-way synchronization between on-premises directories and Azure)	No	Yes	Yes	Yes
Password protection	No	No	Yes	Yes
Group access management	No	No	Yes	Yes

Feature	Free	Office 365 Apps	Premium P1	Premium P2
Azure AD Join + Mobile Device Management (MDM) auto-enrollment	No	No	Yes	Yes
Dynamic groups	No	No	Yes	Yes
Azure Information Protection integration	No	No	Yes	Yes
MFA with conditional access	No	No	Yes	Yes
Microsoft Cloud App Security integration	No	No	Yes	Yes
Privileged Identity Management (PIM)	No	No	No	Yes
Vulnerabilities and risky accounts detection	No	No	No	Yes
Risk-based Conditional Access policies	No	No	No	Yes

Overview of Azure AD Premium P1 and P2

The Azure AD Premium tier editions provide additional functionality compared to the Free and Office 365 editions. However, these editions require additional cost per user provisioning. As previously described, Azure AD Premium has two versions: P1 and P2. You can procure it as an additional license or as a part of the Microsoft Enterprise Mobility + Security, which also includes the license for Azure Information Protection and Intune.

Microsoft provides a free trial period of 90 days that covers 100 user licenses that you can use to experience the full functionality of the Azure AD Premium P2 edition. The following features are available with the Azure AD Premium P1 edition:

- Self-service group management. It simplifies the administration of groups where users are given the rights to create and manage the groups. End users can create requests to join other groups, and group owners can approve requests and maintain their groups' memberships.
- Advanced security reports and alerts. You can monitor and protect access to your cloud applications by observing detailed logs that display advanced anomalies and inconsistent access pattern reports. Advanced reports are machine learning-based and can help you gain new insights to improve access security and respond to potential threats.
- Multifactor authentication (MFA). Full MFA works with on-premises applications that use virtual private network (VPN), Remote Authentication Dial-In User Service (RADIUS), and other methods to access resources. It also works with Azure, Office 365, Dynamics 365, and third-party Azure AD gallery applications. It does not work with non-browser off-the-shelf apps, such as Microsoft Outlook.
- Microsoft Identity Manager (MIM) licensing. MIM integrates with Azure AD Premium to provide hybrid identity solutions. MIM can bridge multiple on-premises authentication stores such as Active

Directory Domain Services (AD DS), Lightweight Directory Access Protocol (LDAP), and other applications with Azure AD. This provides consistent experiences to on-premises line-of-business (LOB) applications and SaaS solutions.

- Enterprise SLA of 99.9%. Enterprise SLA guarantees at least 99.9% availability of the Azure AD Premium service.
- Password reset with writeback. Self-service password reset follows the Active Directory on-premises password policy.
- Cloud App Discovery feature of Azure AD. This feature discovers the most frequently used cloud-based applications.
- Conditional Access based on device, group, or location. This lets you configure conditional access for critical resources, based on several criteria.
- Azure AD Connect Health. You can use this tool to gain operational insight into Azure AD. It works with alerts, performance counters, usage patterns, and configuration settings, and presents the collected information in the Azure AD Connect Health portal.

In addition to these features, the Azure AD Premium P2 license provides two additional functionalities:

- Azure AD Identity Protection. This feature provides enhanced functionalities for monitoring and protecting user accounts. You can define user risk policies and sign-in policies. In addition, you can review users' behavior and flag users for risk.
- Azure AD Privileged Identity Management. This functionality lets you configure additional security levels for privileged users such as administrators. With Privileged Identity Management, you define permanent and temporary administrators. You also define a policy workflow that activates whenever someone wants to use administrative privileges to perform some task.

Connect AD DS with Azure AD by using Azure AD Connect

Although most deployment scenarios for Microsoft Azure Active Directory (Azure AD) do not involve an on-premises Active Directory Domain Services (AD DS) environment, some do. For those organizations that have some services on their on-premises networks and some services in the cloud, synchronization and integration between on-premises AD DS and Azure AD is the way to deliver the best user experience.

Directory synchronization enables user, group, and contact synchronization between on-premises AD DS and Azure AD. In its simplest form, you install a directory synchronization component on a server in your on-premises AD DS domain. All your user accounts, groups, and contacts from AD DS then replicate to Azure AD. Users with access to those accounts can sign in and access Azure services.

With Azure AD Free or Azure AD Office 365 apps, the synchronization flow is one-way from on-premises AD DS to Azure AD. However, with Azure AD Premium, you can replicate some attributes from Azure AD to AD DS. For example, you can configure Azure AD to write passwords back to your on-premises AD DS.

Azure AD Connect

Microsoft provides Azure AD Connect to perform directory synchronization between Azure AD and AD DS. By default, Azure AD Connect synchronizes all users and groups. If you don't want to synchronize your entire on-premises AD DS, directory synchronization for Azure AD supports a degree of filtering and customization of attribute flow based on the following values:

- Group

- Organizational unit (OU)
- Domain
- User attributes
- Applications

When you enable directory synchronization, you have the following authentication options:

- Separate cloud password. When you synchronize a user identity and not the password, the cloud-based user account will have a separate unique password, which can be confusing for users.
- Synchronized password. If you enable password hash synchronization, the AD DS user password hash syncs with the identity in Azure AD. This allows users to authenticate by using the same credentials, but it doesn't provide seamless single sign-on (SSO), because users still receive prompts to authenticate cloud services.
- Pass-through authentication. When you enable pass-through authentication, Azure AD uses the cloud identity to verify that the user is valid and then passes the authentication request to Azure AD Connect. This option provides true SSO because users don't receive multiple prompts to authenticate with cloud services.
- Federated identities. If you configure federated identities, the authentication process is similar to pass-through authentication, but Active Directory Federation Services (AD FS) performs authentication on-premises instead of Azure AD Connect. This authentication method provides claims-based authentication that multiple cloud-based apps can use.

When you install Azure AD Connect, you must sign in as a local administrator of the computer on which you are performing the installation.

Additionally, you will receive prompts for credentials to the local AD DS and Azure AD. The account you use to connect to the local AD DS account must be a member of the **Enterprise Admins** group. The Azure AD account you specify must be a global administrator. If you're using AD FS or a separate SQL Server instance, you will also receive prompts for credentials with management permissions for those resources.

The computer that is running Azure AD Connect must be able to communicate with Azure AD. If the computer needs to use a proxy server for internet access, then additional configuration is necessary.

Note: You don't need inbound connectivity from the internet because Azure AD Connect initiates all communication.

Azure AD Connect must be on a domain member. When you install Azure AD Connect, you can use express settings or custom settings. Most organizations that synchronize a single AD DS forest with an Azure AD tenant use the express settings option.

Note: Installing Azure AD Connect on a domain controller is supported, but this typically occurs only in smaller organizations with limited licensing.

When you choose express settings, the following options are selected:

- SQL Server Express is installed and configured.
- All identities in the forest are synchronized.
- All attributes are synchronized.
- Password synchronization is enabled.
- An initial synchronization is performed immediately after install.
- Automatic upgrade is enabled.

You can enable additional options during installation when you select custom settings, such as:

- Pass-through authentication.
- Federation with AD FS.
- Select an attribute for matching existing cloud-based users.
- Filtering based on OUs or attributes.
- Exchange hybrid.
- Password, group, or device writeback.

After deploying Azure AD Connect, the following occurs:

- New user, group, and contact objects in on-premises Active Directory are added to Azure AD. However, no licenses for cloud services, such as Office 365, are automatically assigned to these objects.
- Attributes of existing user, group, or contact objects that are modified in on-premises Active Directory are modified in Azure AD. However, not all on-premises Active Directory attributes synchronize with Azure AD.
- Existing user, group, and contact objects that are deleted from on-premises Active Directory are deleted from Azure AD.

Existing user objects that are disabled on-premises are disabled in Azure. However, licenses aren't automatically unassigned.

Federation support

The primary feature that AD FS and Web Application Proxy facilitate is federation support. A federation resembles a traditional trust relationship, but it relies on claims (contained within tokens) to represent authenticated users or devices. It relies on certificates to establish trusts and to facilitate secure communication with an identity provider. In addition, it relies on web-friendly protocols such as HTTPS, Web Services Trust Language (WS-Trust), Web Services Federation (WS-Federation), or OAuth to handle transport and processing of authentication and authorization data. Effectively, AD DS, in combination with AD FS and Web Application Proxy, can function as a claims provider that is capable of authenticating requests from web-based services and applications that are not able to, or not permitted to, access AD DS domain controllers directly.

Benefits of integrating Azure AD with AD DS

In a purely on-premises environment, management of authentication, authorization, and other security-related settings, together with device and app management, is provided by tools such as Group Policy and Endpoint Configuration Manager. For organizations operating solely within the cloud, authentication and authorization is provided by Azure Active Directory (Azure AD) so that users can gain access to cloud apps and resources. For these organizations, Microsoft Intune and the Microsoft Store for Business provide device and app management.

Some organizations might choose to implement Azure AD in addition to their on-premises AD DS infrastructure. For some, this might be a temporary measure during a staged migration from AD DS to the cloud. For others, it might be a permanent configuration. In *hybrid* environments, administrators can use on-premises and cloud-based provisioning along with management tools to manage their users' devices. Consequently, there are numerous benefits to operating in this hybrid environment. The following sections discuss some of these benefits.

Azure Information Protection

Azure Information Protection is a set of cloud-based technologies that provide classification, labeling, and data protection. You can use Azure Information Protection to classify, label, and protect data such as emails and documents created in Microsoft Office apps or other supported apps. Instead of focusing only on data encryption, Azure Information Protection has a wider scope. It provides mechanisms to recognize sensitive data, alert users when they deal with sensitive data, and track critical data usage. However, the key component of Azure Information Protection is data protection based on rights management technologies. In hybrid environments, you can extend the reach of Azure Information Protection to your on-premises apps such as Microsoft Exchange Server and Microsoft SharePoint Server.

Note: Microsoft intends to replace Azure Information Protection with Microsoft Identity Protection around 2021, after which Azure Information Protection will no longer be available in the Azure portal.

Classification, labeling, and protection

To use Azure Information Protection, you should configure rules and policies for classification, labeling, and protection. For example, you can configure some data types, keywords, or phrases to be conditions for automatic or recommended classification. The Azure Information Protection client component monitors documents or emails in real time. If it detects a keyword or a phrase, it recommends a proper classification for a document.

Note: You can also configure Azure Information Protection to apply classification automatically. For example, you can configure an automatic classification rule that classifies a document as restricted if it contains a credit card number.

The result of classification is a label. A label is metadata for a document that appears in files and email headers in clear text. The label has clear text so that other services, such as data loss prevention (DLP) solutions or protection solutions, can identify the classification and take appropriate action. For example, a label can be confidential, restricted, or public. The label also contains protection configuration if a specific label requires the protection.

For example, this protection could provide read-only access for certain users within the company. After Azure Information Protection applies protection to a document or an email, the protection remains until an author or a super user removes it.

Note: Azure Information Protection is available in Azure AD Premium P1 and P2.

Self-service password reset

You can enable self-service password reset (SSPR) on Azure AD, either for all or selected users. Using this feature, users can change their passwords or unlock their accounts after their password expires. This feature only affects the Azure AD account. In a hybrid environment, where Azure AD connects to an on-premises AD DS, this process can result in a password mismatch. However, you can implement password writeback to synchronize password changes in Azure AD back to your on-premises AD DS environment.

By bringing your devices to Azure AD, you maximize your users' productivity through single sign-on (SSO) across your cloud and on-premises resources. At the same time, you can secure access to your cloud and on-premises resources with conditional access, which is a capability of Azure AD. With conditional access, you can implement automated access control decisions for accessing applications such as SharePoint Online or Exchange Online, which are based on conditions. For example, you could create a conditional access policy that requires admin users to sign in with multifactor authentication. Or you could require a device to be compliant in terms of security settings before it could access Exchange.

Endpoint co-management

If you have an on-premises Active Directory environment and you want to co-manage your domain-joined devices, you can do this by configuring hybrid Azure AD-joined devices. Co-management is managing Windows 10 devices with on-premises technologies, such as Group Policy and Endpoint Configuration Management, and by using Intune (Endpoint Manager) policies. You can manage some aspects by using Endpoint Configuration Manager and other aspects by using either Intune or Endpoint Manager.

Intune, which provides mobile device management (MDM), enables you to configure settings that achieve your administrative intent without exposing every setting. In contrast, Group Policy exposes fine-grained settings that you control individually. With MDM, you can apply broader privacy, security, and application management settings through lighter and more efficient tools. MDM also allows you to target internet-connected devices to manage policies without using Group Policy that requires on-premises domain-joined devices. This makes MDM the best choice for devices that are constantly on the go.

Note: Intune is a cloud-based service that you can use to manage computers, laptops, tablets, other mobile devices, applications running on these devices, and data that you store on these devices. Intune is a component of the Microsoft 365 platform.

After you join your on-premises AD DS devices to Azure AD, you can immediately use the following Intune features.

Remote actions:

- Factory reset
- Selective wipe
- Delete devices
- Restart device
- Fresh start

Orchestration with Intune for the following workloads:

- Compliance policies
- Resource access policies
- Windows Update policies
- Endpoint Protection
- Device configuration
- Office Click-to-Run apps

Manage apps

Intune provides mobile application management (MAM) capabilities, in addition to MDM. You can use Intune to deploy, configure, and manage apps within your organization for devices that are Azure AD-joined, Azure AD-registered, and Azure AD hybrid-joined.

Test Your Knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

What would you use to synchronize user details to Microsoft Azure Active Directory (Azure AD) from Active Directory Domain Services (AD DS)?

Question 2

Which version of Azure AD supports Azure AD Join + mobile device management (MDM) autoenrollment?

Implementing Group Policy

Lesson overview

Since early versions of Windows Server, the Group Policy feature of Windows operating systems has provided an infrastructure with which administrators can define settings centrally and then deploy them to computers across their organizations. In an environment managed by a well-implemented Group Policy infrastructure, an administrator rarely configures settings by directly touching a user's computer. You can define, enforce, and update the entire configuration by using Group Policy Object (GPO) settings. By using GPO settings, you can affect an entire site or a domain within an organization, or you can narrow your focus to a single organizational unit (OU). Filtering based on security group membership and physical computer attributes allows you to define the target for your GPO settings even further. This lesson explains what Group Policy is, how it works, and how best to implement it in your organization.

Lesson objectives

After completing this lesson, you'll be able to:

- Describe GPOs.
- Describe GPO scope and inheritance.
- Explain domain-based GPOs.
- Identify the default domain-based GPOs.
- Create and configure GPOs.
- Explain GPO storage.
- Describe Start GPOs.
- Explain administrative templates.
- Describe how to use the Central Store.

What are GPOs?

Overview

Consider a scenario in which you have only one computer in your home environment and you wish to modify the desktop background. You can do it in several different ways. Often people open **Personalization** from the **Settings** app in Windows 10 and then make the change by using the Windows operating system interface. Although that works well for one computer, it might be tedious if you want to make the same change across multiple computers. With multiple computers, it is more difficult to implement changes and maintain a consistent environment.

Group Policy is a framework in Windows operating systems with components that reside in Active Directory Domain Services (AD DS), on domain controllers, and on each Windows server and client. By using these components, you can manage configuration in an AD DS domain. You define Group Policy settings within a Group Policy Object (GPO). A GPO is an object that contains one or more policy settings that apply to one or more configuration settings for a user or a computer.

Group Policy is a powerful administrative tool. You can use GPOs to push various settings to a large number of users and computers. Because you can apply them to different levels, from the local computer to domain, you also can focus these settings precisely. Primarily, you use Group Policy to configure

settings that you do not want users to configure. Additionally, you can use Group Policy to standardize desktop environments on all computers in an organizational unit (OU) or in an entire organization. You also can use Group Policy to provide additional security, to configure some advanced system settings, and for other purposes that the following sections detail.

Apply security settings

In the Windows Server operating system, GPOs include a large number of security-related settings that you can apply to both users and computers. For example, you can enforce settings for the domain password policy, for Windows Defender Firewall, and you can configure auditing and other security settings. You also can configure full sets of user-rights assignments.

Manage desktop and application settings

You can use Group Policy to provide a consistent desktop and application environment for all users in your organization. By using GPOs, you can configure each setting that affects the representation of the user environment. You also can configure settings for some applications that support GPOs.

Deploying software

With Group Policy, you can deploy software to users and computers. You can use Group Policy to deploy all software that is available in the .msi format. Additionally, you can enforce automatic software installation, or you can let your users decide whether they want the software to deploy to their computers.

Note: Deploying large software packages with GPOs might not be the most efficient way to distribute an application to your organization's computers. In some circumstances, it might be more effective to distribute applications as part of the desktop computer image.

Manage Folder Redirection

With the **Folder Redirection** option, it is easier to back up users' data files. By redirecting folders, you also ensure that users have access to their data regardless of the computer to which they sign in. Additionally, you can centralize all users' data to one place on a network server, while still providing a user experience that is similar to storing these folders on their computers. For example, you can configure **Folder Redirection** to redirect users' **Documents** folders to a shared folder on a network server.

Configuring network settings

By using Group Policy, you can configure various network settings on client computers. For example, you can enforce settings for wireless networks to allow users to connect only to specific Wi-Fi network service set identifiers (SSIDs) and with predefined authentication and encryption settings. You also can deploy policies that apply to wired network settings, and some Windows Server roles use Group Policy to configure the client side of services, such as DirectAccess.

Group Policy Objects

The most granular component of Group Policy is an individual policy setting. An individual policy setting defines a specific configuration, such as a policy setting that prevents a user from accessing registry-editing tools. If you define that policy setting and then apply it to a user, that user will be unable to run tools such as **Regedit.exe**.

Note that some settings affect a user, known as user configuration settings or user policies, and some affect the computer, known as computer configuration settings or computer policies. However, settings do not affect groups, security principals other than user objects, computer objects, or other directory objects.

Group Policy manages various policy settings, and the Group Policy framework is extensible. You can manage almost any configurable setting with Group Policy.

In the Group Policy Management Editor, you can define a policy setting by selecting it and then selecting Enter. The policy setting **Properties** dialog box appears. Most policy settings can have three states: **Not Configured**, **Enabled**, and **Disabled**.

GPOs store Group Policy settings. In a new GPO, every policy setting defaults to **Not Configured**. When you enable or disable a policy setting, Windows Server makes a change to the configuration of users and computers to which the GPO is applied. When you return a setting to its **Not Configured** value, you return it to its default value.

To create a new GPO in a domain, right-click or access the context menu for the **Group Policy Objects** container, and then select **New**. To modify the configuration settings in a GPO, right-click or access the context menu for the GPO, and then select **Edit**. This opens the Group Policy Management Editor snap-in.

The Group Policy Management Editor displays all the policy settings that are available in a GPO in an organized hierarchy that begins with the division between computer settings and user settings: the **Computer Configuration** node and the **User Configuration** node.

GPOs display in a container named Group Policy Objects. The next two levels of the hierarchy are nodes named **Policies** and **Preferences**. Progressing through the hierarchy, the Group Policy Management Editor displays folders, called nodes or policy setting groups. The policy settings are within the folders.

Overview of GPO scope and inheritance

Policy settings in Group Policy Objects (GPOs) define configuration. However, you must specify the computers or users to which the GPO applies before the configuration changes in a GPO will affect computers or users in your organization. This is called scoping a GPO. The scope of a GPO is the collection of users and computers that will apply the settings in the GPO. You can use several methods to manage the scope of domain-based GPOs.

Scope a GPO

The first is the GPO link. You can link GPOs to sites, domains, and organizational units (OUs) in Active Directory Domain Services (AD DS). The site, domain, or OU then becomes the maximum scope of the GPO. The configurations that the policy settings in the GPO specify will affect all computers and users within the site, domain, or OU, including those in child OUs.

Note: You can link a GPO to more than one domain, OU, or site. Linking GPOs to multiple sites can introduce performance issues when applying the policy, and you should avoid linking a GPO to multiple sites. This is because, in a multiple-site network, the GPOs are stored on the domain controllers in the domain where the GPOs were created. The consequence of this is that computers in other domains might need to traverse a slow wide area network (WAN) link to obtain the GPOs.

You can further narrow the scope of the GPO with one of two types of filters:

- Security filters. These specify security groups or individual user or computer objects that relate to a GPO's scope, but to which the GPO explicitly should or should not apply.
- Windows Management Instrumentation (WMI) filters. These specify a scope by using characteristics of a system, such as an operating system version or free disk space.

Use security filters and WMI filters to narrow or specify the scope within the initial scope that the GPO link created.

GPO processing order

The GPOs that apply to a user, computer, or both do not apply all at once. GPOs apply in a particular order. Conflicting settings that process later might overwrite settings that process first.

Group Policy follows the following hierarchical processing order:

1. Local GPOs. Each computer has at least one local Group Policy. The local policies apply first when you configure such policies.
2. Site-linked GPOs. Policies linked to sites process second. If there are multiple site policies, they process synchronously in the listed preference order.
3. Domain-linked GPOs. Policies linked to domains process third. If there are multiple domain policies, they process synchronously in the listed preference order.
4. OU-linked GPOs. Policies linked to top-level OUs process fourth. If there are multiple top-level OU policies, they process synchronously in the listed preference order.
5. Child OU-linked GPOs. Policies linked to child OUs process fifth. If there are multiple child OU policies, they process synchronously in the listed preference order. When there are multiple levels of child OUs, policies for higher-level OUs apply first and policies for the lower-level OUs apply next.

In Group Policy application, the general rule is that the last policy applied prevails. For example, a policy that restricts access to the Control Panel applied at the domain level could be reversed by a policy applied at the OU level for the objects contained in that particular OU.

If you link several GPOs to an OU, their processing occurs in the order that the administrator specifies on the OU's **Linked Group Policy Objects** tab in the **Group Policy Management Console (GPMC)**. By default, processing is enabled for all GPO links. You can disable a container's GPO link to block the application of a GPO completely for a given site, domain, or OU. For example, if you made a recent change to a GPO and it is causing production issues, you can disable the link or links until the issue resolves. Note that if the GPO is linked to other containers, they will continue to process the GPO if their links are enabled.

You also can disable the user or computer configuration of a particular GPO independently from either the user or computer. If one section of a policy is known to be empty, disabling the other side speeds up policy processing slightly. For example, if you have a policy that only delivers user desktop configuration, you could disable the computer side of the policy.

GPO inheritance

You can configure a policy setting in more than one GPO, which might result in GPOs conflicting with each other. For example, you might enable a policy setting in one GPO, disable it in another GPO, and then not configure it in a third GPO. In this case, the precedence of the GPOs determines which policy setting the client applies. A GPO with higher precedence prevails over a GPO with lower precedence.

The **GPMC** has precedence as a number. The smaller the number—that is, the closer the number is to 1—the higher the precedence. Therefore, a GPO that has a precedence of 1 will prevail over all other GPOs. Select the relevant AD DS container, and then select the **Group Policy Inheritance** tab to review the precedence of each GPO.

When you enable or disable a policy setting in a GPO with higher precedence, the configured setting takes effect. However, remember that by default, Windows Server sets policy settings to **Not Configured**.

If you do not configure a policy setting in a GPO with higher precedence, the policy setting, either enabled or disabled, in a GPO with lower precedence will take effect. If multiple GPOs link to an AD DS container object, the objects' link order determines their precedence.

The default behavior of Group Policy is that GPOs linked to a higher-level container are inherited by lower-level containers. When a computer starts up or a user signs in, the Group Policy Client Extensions examines the location of the computer or user object in AD DS and evaluates the GPOs with scopes that include the computer or user. Then, the client-side extensions apply policy settings from these GPOs. Policies apply sequentially, beginning with the policies that link to the site, followed by those that link to the domain, followed by those that link to OUs—from the top-level OU down to the OU in which the user or computer object exists. It is a layered application of settings, so a GPO that applies later in the process overrides settings that applied earlier in the process because it has higher precedence.

The sequential application of GPOs creates an effect called policy inheritance. Policies are inherited, which means that the Resultant Set of Policies (RSOPs) for a user or computer will be the cumulative effect of site, domain, and OU policies.

By default, inherited GPOs have lower precedence than GPOs that link directly to a container. For example, you might configure a policy setting to disable the use of registry-editing tools for all users in the domain by configuring the policy setting in a GPO that links to the domain. All users within the domain inherit that GPO and its policy setting. However, because you probably want administrators to be able to use registry-editing tools, you will link a GPO to the OU that contains administrators' accounts and then configure the policy setting to allow the use of registry-editing tools. Because the GPO that links to the administrators' OU takes higher precedence than the inherited GPO, administrators will be able to use registry-editing tools.

Precedence of multiple linked GPOs

You can link more than one GPO to an AD DS container object. The link order of GPOs determines the precedence of GPOs in such a scenario. GPOs with a higher link order take precedence over GPOs with a lower link order. When you select an OU in the **GPMC**, the **Linked Group Policy Objects** tab displays the link order of GPOs that link to that OU.

To change the precedence of a GPO link, follow this procedure:

1. Select the AD DS container object in the **GPMC** console tree.
2. Select the **Linked Group Policy Objects** tab in the details pane.
3. Select the GPO.
4. Use the **Up**, **Down**, **Move To Top**, and **Move To Bottom** arrows to change the link order of the selected GPO.

Block Inheritance

You can configure a domain or OU to prevent the inheritance of policy settings. This is known as blocking inheritance. To block inheritance, right-click or access the context menu for the domain or OU in the **GPMC** console tree, and then select **Block Inheritance**.

The **Block Inheritance** option is a property of a domain or OU, so it blocks all Group Policy settings from GPOs that link to parents in the Group Policy hierarchy. For example, when you block inheritance on an OU, GPO application begins with any GPOs that link directly to that OU. Therefore, GPOs that are linked to higher-level OUs, the domain, or the site will not apply.

You should use the **Block Inheritance** option sparingly because blocking inheritance makes it more difficult to evaluate Group Policy precedence and inheritance. With security group filtering, you can

carefully scope a GPO so that it applies to only the correct users and computers in the first place, making it unnecessary to use the **Block Inheritance** option.

Enforce a GPO link

Additionally, you can set a GPO link to be enforced. To enforce a GPO link, right-click or access the context menu for the GPO link in the console tree, and then select **Enforced** from the shortcut menu. When you set a GPO link to **Enforced**, the GPO takes the highest level of precedence. Policy settings in that GPO will prevail over any conflicting policy settings in other GPOs. Furthermore, an enforced link will apply to child containers even when those containers are set to **Block Inheritance**. The **Enforced** option causes the policy to apply to all objects within its scope. The **Enforced** option will cause policies to override any conflicting policies and will apply, regardless of whether a **Block Inheritance** option is set. Enforcement is useful when you must configure a GPO that defines a configuration that is mandated by your corporate IT security and usage policies. Therefore, you want to ensure that other GPOs that are linked to the same or lower levels do not override those settings. You can do this by enforcing the GPO's link.

Evaluating precedence

To facilitate evaluation of GPO precedence, you can simply select an OU or domain, and then select the **Group Policy Inheritance** tab. This tab will display the resulting precedence of GPOs, accounting for GPO link, link order, inheritance blocking, and link enforcement. This tab does not account for policies that are linked to a site, for GPO security, or WMI filtering.

What are domain-based GPOs?

You can create domain-based Group Policy Objects (GPOs) in Active Directory Domain Services (AD DS) and store them on domain controllers. You can use them to manage configuration centrally for the domain's users and computers. The other type of a GPO is a local GPO, which is linked to a specific computer. The remainder of this lesson refers to domain-based GPOs rather than local GPOs, unless otherwise specified.

When you install AD DS, Windows Server creates two default GPOs:

- Default Domain Policy
- Default Domain Controllers Policy

Note: Windows computers also have local GPOs, which are primarily used when computers are not connected to domain environments. All Windows operating systems support the existence of multiple local GPOs. As with domain-based GPOs, it is a good practice to create new GPOs for customizations. In the **Computer Configuration** node, you can configure all computer-related settings. In the **User Configuration** node, you can configure settings that you want to apply to all users on a computer. The user settings in the local computer GPO can be modified by the user settings in two new local GPOs: Administrators and Non-Administrators. These two GPOs apply user settings to signed-in users according to the group to which they belong. If they are members of the local **Administrators** group, the users would use the Administrators GPO and if they are not members of the **Administrators** group, they would use the Non-Administrators GPO. You can further refine the user settings with a local GPO that applies to a specific user account. User-specific local GPOs are associated with local user accounts and not with the domain.

Note: Domain-based GPO settings combine with those applied by using local GPOs, but because domain-based GPOs apply after local GPOs and there are conflicting settings, the settings from the do-

main-based GPOs take precedence over the settings from local GPOs. Also, note that you can disable local GPOs by using a domain-based GPO.

Default domain GPOs

As mentioned earlier, a domain has two default Group Policy Objects (GPOs):

- Default Domain Policy
- Default Domain Controllers Policy

These are linked to the domain object and to the built-in Domain Controllers organizational unit (OU), respectively.

Default Domain Policy

The Default Domain Policy GPO is linked to the domain, and it applies to Authenticated Users. This GPO does not have any Windows Management Instrumentation (WMI) filters. Therefore, it affects all users and computers in the domain. This GPO contains policy settings that specify password, account lockout, and Kerberos version 5 authentication protocol policies.

These settings are of critical importance to the Active Directory Domain Services (AD DS) environment, and thus, make the Default Domain Policy a critical component of Group Policy. You should not add unrelated policy settings to this GPO. If you need to configure other settings to apply broadly in your domain, create additional GPOs that link to the domain.

Default Domain Controllers Policy

The Default Domain Controllers Policy GPO links to the organizational unit (OU) of the domain controllers. Because computer accounts for domain controllers are kept exclusively in the Domain Controllers OU, and other computer accounts should be kept in other OUs, this GPO affects only domain controllers or other computer objects that are in the Domain Controllers OU.

You should modify GPOs linked to the Domain Controllers OU to implement your auditing policies and to assign user rights that are required on domain controllers.

Demonstration: Create and configure a domain-based GPO

In this demonstration, you will learn how to:

- Manage objects in AD DS.
- Create and edit a GPO.
- Link the GPO.
- View the effects of the GPOs settings.
- Create and link the required GPOs.
- Verify the order of precedence.
- Configure the scope of a GPO with security filtering.
- Verify the application of settings.

Demonstration steps

Manage objects in AD DS

1. Switch to **SEA-ADM1** and then switch to Windows PowerShell.
2. Create an organizational unit (OU) called **Seattle** in the domain.
3. Create a user account for **Ty Carlson** in the **Seattle** OU.
4. Test the account by switching to **SEA-CL1**, and then signing in as **Ty**.
5. Create a group called **SeattleBranchUsers** and add **Ty** to the group.

Create and edit a GPO

1. On **SEA-ADM1**, open **Group Policy Management Console**.
2. Create a GPO named **Contoso Standards** in the **Group Policy Objects** container.
3. Edit the **Contoso Standards** policy, and configure the **Screen saver** timeout policy to **600** seconds.
4. Enable the **Password protect the screen saver** policy setting.

Link the GPO

- Link the **Contoso Standards** GPO to the `Contoso.com` domain.

Review the effects of the GPO's settings

1. Sign in to **SEA-CL1** as **Contoso\Administrator**.
2. Allow **Remote Event Log Management** and **Windows Management Instrumentation (WMI)** traffic through Windows Defender Firewall.
3. Sign out and then sign in as **Contoso\Ty** with the password **Pa55w.rd**.
4. Attempt to change the screen saver wait time and resume settings. Group Policy prevents you from doing this.
5. Attempt to run Registry Editor. Group Policy prevents you from doing this.

Create and link the required GPOs

1. Create a new GPO named **Seattle Application Override** that is linked to the **Seattle** OU.
2. Configure the **Screen saver timeout** policy setting to be disabled.

Verify the order of precedence

- In the **Group Policy Management Console** tree, select the **Seattle** OU and select the **Group Policy Inheritance** tab.

Notice that the Seattle Application Override GPO has precedence over the Contoso Standards GPO. The screen saver time-out policy setting that you just configured in the Seattle Application Override GPO will be applied after the setting in the Contoso Standards GPO. Therefore, the new setting will overwrite the

standards setting and will prevail. Screen saver time-out will be unavailable for users within the scope of the Seattle Application Override GPO.

Configure the scope of a GPO with security filtering

1. Select the **Seattle Application Override** GPO. Notice that in the **Security Filtering** section, the GPO applies to all authenticated users by default.
2. In the **Security Filtering** section, remove **Authenticated Users**, and then add the **SeattleBranchUsers** group and **SEA-CL1**.

Verify the application of settings

1. Launch the **Group Policy Results Wizard**.
2. Select the **SEA-CL1** computer and the **CONTOSO\Ty** user account.
3. After the report is created, in the details pane, select the **Details** tab, and then select **show all**.
4. In the report, scroll down until you locate the **User Details** section, and then locate the **Control Panel/Personalization** section. You should notice that the **Screen save timeout** settings are obtained from the Seattle Application Override GPO.

Overview of GPO storage

Group Policy settings present as Group Policy Settings (GPOs) in Active Directory Domain Services (AD DS) user interface tools, but a GPO actually includes two components:

- The **Group Policy** container. The **Group Policy Objects** container stores the **Group Policy** container, an AD DS object, within the domain-naming context of the directory.
- The **Group Policy** template. This template is a collection of files stored in the **SYSVOL** of each domain controller in the **%SystemRoot%\SYSVOL\Domain\Policies\GPOGUID** path, where **GPOGUID** is the globally unique identifier (GUID) of the **Group Policy** container.

Similar to all AD DS objects, each **Group Policy** container includes a GUID attribute that uniquely identifies the object within AD DS. The **Group Policy** container defines basic attributes of the GPO, but it does not contain any of the settings. The **Group Policy** template contains the **Group Policy** settings.

When you change the settings of a GPO, the changes save to the **Group Policy** template of the server from which you opened the GPO. By default, when Group Policy refresh occurs, the client-side extensions apply settings in a GPO only if the GPO has been updated.

The Group Policy client can identify an updated GPO by its version number. Each GPO has a version number that increments each time you make a change. The version number is stored as a **Group Policy** container attribute and in a text file, **Group Policy GPO.ini**, in the **Group Policy** template folder. The Group Policy client knows the version number of each GPO it has previously applied. If, during the Group Policy refresh, the Group Policy client discovers that the version number of the **Group Policy** container has changed, Windows Server will inform the client-side extensions that the GPO is updated.

GPO replication

The **Group Policy** container and the **Group Policy** template both replicate between all domain controllers in AD DS. However, these two items use different replication mechanisms.

- The **Group Policy** container in AD DS replicates by using the Directory Replication Agent (DRA). The DRA uses a topology that the Knowledge Consistency Checker generates, which you can define or refine manually. The result is that the **Group Policy** container replicates within seconds to all domain controllers in a site and replicates between sites based on your intersite replication configuration.
- The **Group Policy** template in the **SYSVOL** replicates by using the Distributed File System (DFS) Replication.

Because the **Group Policy** container and **Group Policy** template replicate separately, it is possible for them to become out-of-sync for a brief time. Typically, when this happens, the **Group Policy** container will replicate to a domain controller first.

Systems that obtained their ordered list of GPOs from that domain controller will identify the new **Group Policy** container. Those systems will then attempt to download the **Group Policy** template, and they will notice that the version numbers are not the same. A policy processing error will record in the event logs.

If the reverse happens, and the GPO replicates to a domain controller before the **Group Policy** container, clients that obtain their ordered list of GPOs from that domain controller will not be notified of the new GPO until the **Group Policy** container has replicated.

What are Starter GPOs?

You can use a Starter GPO as a template from which to create other Group Policy Objects (GPOs) within the **Group Policy Management Console (GPMC)**. Starter GPOs only have **Administrative Template** settings. You might use a Starter GPO to provide a starting point to create new GPOs in your domain. The Starter GPO might already have specific settings that are best practices for your environment. You can export starter GPOs to, and import them from, cabinet (.cab) files to make distribution to other environments simple and efficient.

The **GPMC** stores Starter GPOs in a folder, **StarterGPOs**, which is in **SYSVOL**. Microsoft includes preconfigured Starter GPOs for Windows client operating systems. These Starter GPOs have **Administrative Template** settings that reflect best practices that Microsoft recommends for the configuration of the client environment.

What are administrative templates?

Administrative template files provide most of the available Group Policy Objects (GPO) settings, which modify specific registry keys. The use of administrative templates is known as a registry-based policy, because all the settings you configure in administrative templates result in changes to the registry. For many apps, using a registry-based policy is the simplest and best way to support the centralized management of policy settings.

You can use administrative templates to control the environment of an operating system and the user experience. There are two sets of administrative templates:

- User-related settings
- Computer-related settings

When you configure settings in the **Administrative Templates** node of the GPO, you make modifications to the registry. Administrative templates have the following characteristics:

- They have subnodes that correspond to specific areas of the environment, such as network, system, and Windows components.
- The settings in the computer section of the **Administrative Templates** node edit the **HKEY_LOCAL_MACHINE** hive in the registry, and the settings in the user section of the **Administrative Templates** node edit the **HKEY_CURRENT_USER** hive in the registry.
- Some settings exist for both user and computer. In the case of conflicting settings, the computer setting will prevail.
- Some settings are available only to certain versions of Windows operating systems. For example, you can apply several new settings only to Windows 10. You can double-click or select the setting, and then select Enter to display the supported versions for that setting.

The following table details the organization of the **Administrative Templates** node.

Table 1: Administrative template nodes

Administrative template section	Settings
Computer configuration	Control Panel, Network, Printers, Server, Start Menu and Taskbar, System, Windows Components, All Settings
User configuration	Control Panel, Desktop, Network, Shared Folders, Start Menu and Taskbar, System, Windows Components, All Settings

Most of the nodes contain multiple subfolders that enable you to organize settings into logical groupings. Even with this organization, finding the setting that you need might be a daunting task. The **All Settings** node contains an alphabetically sorted list of all settings contained in all the other nodes. Later in this lesson, you will learn how to filter settings in the **Administrative Templates** node to help you locate settings.

What are .admx files?

All the settings in the **Administrative Templates** node of a GPO are stored in files. All currently supported operating systems store the settings in .admx files.

These settings use a standards-based XML file format known as .admx files. By default, Windows Server stores .admx files in the **Windows\PolicyDefinitions** folder, but you can store them in a central location, which you will learn about in the next topic.

The .admx files are language neutral. The plain language descriptions of the settings are not part of the .admx files. Instead, they are stored in language-specific .adml files. This means that administrators can review the same GPO and observe the policy descriptions in their own language because they can each use their own language-specific .adml files.

The **PolicyDefinitions** folder stores .adml files subfolders. Each language has its own folder. For example, the **en-US** folder stores the English files, and the **es-ES** folder stores the Spanish files. By default, only the .adml language files for the language of the installed operating system are present.

Overview of the Central Store

In domain-based enterprises, you can create a Central Store location for .admx files, which anyone with permissions to create or edit Group Policy Objects (GPOs) can access. The Group Policy Management Editor automatically reads and displays administrative templates policy settings from .admx files in the Central Store, and then ignores the .admx files stored locally. If the domain controller or Central Store is not available, the Group Policy Management Editor uses the local store.

The advantages of creating a Central Store are:

- You ensure that whenever someone edits a GPO, the settings in the **Administrative Templates** node are always the same.
- When Microsoft releases .admx files for new operating systems, you only need to update the .admx files in one location.

You must create the Central Store manually, and then update it manually on a domain controller. The use of .admx files is dependent on the operating system of the computer where you create or edit the GPO. The domain controllers use Active Directory Domain Services (AD DS) replication and Distributed File System (DFS) Replication to replicate the data.

To create a Central Store for .admx and .adml files, create a folder and name it **PolicyDefinitions** in the **\\FQDN\SYSVOL\FQDN\Policies** location, where **FQDN** is the domain name for your AD DS domain.

For example, to create a Central Store for the `Test.Microsoft.com` domain, create a **PolicyDefinitions** folder in the following location:

```
\\Test.Microsoft.Com\SYSVOL\Test.Microsoft.Com\policies
```

A user must copy all files and subfolders of the **PolicyDefinitions** folder, which on a Windows computer resides in the **Windows** folder. The **PolicyDefinitions** folder stores all .admx files, and subfolders store .adml files for all languages enabled on the client computer. For example, on a Windows Server computer that has English enabled, **C:\Windows\PolicyDefinitions** will contain the .admx files and in the subfolder **en-US**, the .adml files will contain English-based descriptions for the settings defined in the .admx files.

Note: You must update the **PolicyDefinitions** folder for each feature update and for other software, such as Windows 10 Version 2004 and Microsoft Office 2019 .admx files.

Test Your Knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

If you linked a Group Policy Object (GPO) to the domain object in your Active Directory Domain Services (AD DS), what are the different ways to prevent this policy from applying to all users in the domain?

Question 2

What are the default domain GPOs called?

Overview of AD CS

Lesson overview

The public key infrastructure (PKI) consists of several components, such as certification authority (CA), that help you secure organizational communications and transactions. You can use CAs to manage, distribute, and validate the digital certificates that you use to secure information.

You use digital certificates as a form of authentication. This might be for a computer to identify itself with a wireless access point, or for a user to identify a server that they want to communicate with. You also can use certificates in file and volume encryption, document signing and sealing, on the wire network authentication and encryption, and for many other security-related purposes.

You can install Active Directory Certificate Services (AD CS) as a root CA or a subordinate CA in your organization. In this lesson, you will learn about deploying and managing CAs.

Lesson objectives

After completing this lesson, you will be able to:

- Describe AD CS.
- Explain options for implementing CA hierarchies.
- Compare standalone and enterprise class CAs.
- Describe certificate templates.
- Describe how to manage CAs.
- Explain the purpose of certificate revocation lists (CRLs) and CRL distribution lists.
- Configure trust for certificates.
- Describe how to enroll in a certificate.

What is AD CS?

To use certificates in your Active Directory Domain Services (AD DS) infrastructure, you need to use externally provided certificates or deploy and configure at least one certification authority (CA). The first CA that you deploy is a root CA. After you install the root CA, you can install a subordinate CA to apply policy restrictions and issue certificates.

Active Directory Certificate Services (AD CS) is an identity technology in Windows Server that allows you to implement PKI so that you can easily issue and manage certificates to meet your organization's requirements.

Overview of PKI

Public key infrastructure (PKI) is the combination of software, encryption technologies, processes, and services that enables an organization to secure its communications and business transactions. PKI relies on the exchange of digital certificates between authenticated users and trusted resources. You use certificates to secure data and to manage identification credentials from users and computers both within and outside of your organization.

You can design a PKI solution by using Active Directory Certificate Services (AD CS) to meet the following security and technical requirements of your organization:

- **Confidentiality.** PKI gives you the ability to encrypt stored and transmitted data. For example, you can use a PKI-enabled Encrypting File System (EFS) to encrypt and secure data. You can also maintain the confidentiality of transmitted data on public networks by using PKI-enabled Internet Protocol security (IPsec).
- **Integrity.** You can use certificates to sign data digitally. A digital signature will identify whether any data was modified while communicating information. For example, a digitally signed email message will help ensure that the message's content was not modified while in transit. Additionally, in a PKI, the issuing CA digitally signs certificates that are issued to users and computers, which proves the integrity of the issued certificates.
- **Authenticity.** A PKI provides several authenticity mechanisms. Authentication data passes through hash algorithms such as Secure Hash Algorithm 2 (SHA-2) to produce a message digest. The message digest then is digitally signed by using the sender's private key from the certificate to prove that the sender produced the message digest.
- **Nonrepudiation.** When data is digitally signed with an author's certificate, the digital signature provides both proof of the integrity of the signed data and proof of the data's origin.
- **Availability.** You can install multiple CAs in your CA hierarchy to issue certificates. If one CA is not available in a CA hierarchy, other CAs can continue to issue certificates.

AD CS in Windows Server

Windows Server deploys all PKI-related components as role services of the AD CS server role. Each role service is responsible for a specific portion of the certificate infrastructure while working together to form a complete solution.

The role services of the AD CS role in Windows Server are as follows:

- **Certification Authority.** The main purposes of CAs are to issue certificates, to revoke certificates, and to publish authority information access (AIA) and revocation information. When you install the first CA, it establishes the PKI in your organization. You can have one or more CAs in one network, but only one CA can be at the highest point in the CA hierarchy. The root CA is the CA at the highest point in the hierarchy. However, you can have more than one CA hierarchy, which allows you to have more than one root CA. After a root CA issues a certificate for itself, subordinate CAs that are lower in the hierarchy receive certificates from the root CA.
- **Certification Authority Web Enrollment.** This component provides a method to issue and renew certificates for users, computers, and devices that are not joined to the domain, are not connected directly to the network, or are for users of operating systems other than Windows.
- **Online Responder.** You can use this component to configure and manage Online Certificate Status Protocol (OCSP) validation and revocation checking. An Online Responder decodes revocation status requests for specific certificates, evaluates the status of those certificates, and returns a signed response that has the requested certificate status information.
- **Network Device Enrollment Service (NDES).** With this component, routers, switches, and other network devices can obtain certificates from AD CS.
- **Certificate Enrollment Web Service (CES).** This component works as a proxy client between a computer running Windows and the CA. CES enables users, computers, or applications to connect to a CA by

using web services to:

- Request, renew, and install issued certificates.
- Retrieve certificate revocation lists (CRLs).
- Download a root certificate.
- Enroll over the internet or across forests.
- Renew certificates automatically for computers that are part of untrusted AD DS domains or are not joined to a domain.
- Certificate Enrollment Policy Web Service. This component enables users to obtain certificate enrollment policy information. Combined with CES, it enables policy-based certificate enrollment when the client computer is not a member of a domain or when a domain member is not connected to the domain.

The AD CS server role, in addition to all related role services, can run on Windows Server with a full desktop experience or a Server Core installation. You can deploy the AD CS role services in Windows Server by using Windows Admin Center, **Server Manager**, or Windows PowerShell command-line interface cmdlets. Additionally, you can deploy the role services while working locally at the computer or remotely over the network.

AD CS in Windows Server also supports Trusted Platform Module (TPM) key attestation. AD CS allows you to use the Microsoft Smart Card key storage provider (KSP) for TPM key attestation so that devices that are not domain members can enroll for certificates attesting to a TPM-protected private key by using NDES enrollment.

Options for implementing CA hierarchies

When you decide to implement Active Directory Certificate Services (AD CS) in your organization, one of the first decisions you must make is how to design your certification authority (CA) hierarchy. CA hierarchy determines the core design of your internal public key infrastructure (PKI) and determines the purpose of each CA in the hierarchy. Each CA hierarchy usually includes two or more CAs. Usually, the second CA and all others after that deploy with a specific purpose. Only the root CA is mandatory.

Note: Having a multilevel CA hierarchy that is deployed to use PKI and certificates is not mandatory. For smaller and simpler environments, you can have a CA hierarchy with just one deployed CA. This CA usually deploys as an enterprise root CA. Additionally, you might choose not to deploy an internal CA at all and use externally provided certificates.

If you decide to implement a CA hierarchy and have deployed a root CA already, you must decide which roles to assign CAs on the second and third tiers. In general, we do not recommend building a CA hierarchy deeper than three levels unless it is in a complex and distributed environment.

Most commonly, CA hierarchies have two levels, with the root CA at the top level and the subordinate issuing CA on the second level. Usually, the root CA is taken offline while the subordinate CA issues and manages certificates for all clients. However, in some more complex scenarios, you can also deploy other types of CA hierarchies.

In general, CA hierarchies are in one of following categories:

- CA hierarchies with a policy CA. Policy CAs are types of subordinate CAs that are directly under the root CA in a CA hierarchy. You use policy CAs to issue CA certificates to subordinate CAs that are directly under the policy CA in the hierarchy. The role of a policy CA is to describe the policies and procedures that an organization implements to secure its PKI, the processes that validate the identity

of certificate holders, and the processes that enforce the procedures that manage certificates. A policy CA issues certificates only to other CAs. The CAs that receive these certificates must uphold and enforce the policies that the policy CA defined. Using policy CAs is not mandatory unless different divisions, sectors, or locations of your organization require different issuance policies and procedures. However, if your organization requires different issuance policies and procedures, you must add policy CAs to the hierarchy to define each unique policy. For example, an organization can implement one policy CA for all certificates that it issues internally to employees and another policy CA for all certificates that it issues to users who are not employees.

- CA hierarchies with cross-certification trust. In this scenario, two independent CA hierarchies interoperate when a CA in one hierarchy issues a cross-certified CA certificate to a CA in another hierarchy. When you do this, you establish mutual trust between different CA hierarchies.
- CAs with a two-tier hierarchy. In a two-tier hierarchy, there is a root CA and at least one subordinate CA. In this scenario, the subordinate CA is responsible for policies and for issuing certificates to requestors.

Standalone vs. enterprise CAs

In Active Directory Certificate Services (AD CS), you can deploy two types of certification authorities (CAs): standalone and enterprise CAs. These types of CAs are not about hierarchy, but instead, they are about functionality and configuration storage. The most important difference between these two CA types is Active Directory Domain Services (AD DS) integration and dependency. A standalone CA can work without AD DS and does not depend on it in any way. An enterprise CA requires AD DS, but it also provides several benefits, including autoenrollment. The **Autoenrollment** feature allows users and domain member devices to enroll automatically for certificates if you have enabled automatic certificate enrollment through Group Policy.

The following table details the most significant differences between standalone and enterprise CAs.

Table 1: Comparing standalone and enterprise CAs

Characteristic	Standalone CA	Enterprise CA
Typical usage	You typically use a standalone CA for offline CAs, but you can also use it for a CA that consistently is available on the network.	You typically use an enterprise CA to issue certificates to users, computers, and services, and you cannot use it as an offline CA.
AD DS dependencies	A standalone CA does not depend on AD DS, and you can deploy it in environments other than AD DS.	An enterprise CA requires AD DS, which you use as a configuration and registration database. An enterprise CA also provides a publication point for certificates that issue to users and computers.
Certificate request methods	Users can request certificates only from a standalone CA by using a manual procedure or web enrollment.	Users can request certificates from an enterprise CA by using the following methods: Manual enrollment, Web enrollment, Autoenrollment, Enrollment on behalf, Web services
Certificate issuance methods	A certificate administrator must approve all requests manually.	Requests can be issued or denied automatically based on issuance-requirements settings.

Considerations for deploying a root CA

Before you deploy a root CA, you should decide several aspects. First, you should decide whether you need to deploy an offline root CA. Based on that decision, you also need to decide if you are going to deploy a standalone root CA or an enterprise root CA.

Usually, if you deploy a single-layer CA hierarchy, which means that you deploy only a single CA, it is most common to choose an enterprise root CA. However, if you deploy a two-layer hierarchy with a subordinate CA, the most common scenario is to deploy a standalone root CA. This makes the root CA more secure and allows it to be taken offline except for when it needs to issue certificates for new subordinate CAs.

The next factor to consider is the operating system installation type. Both the Desktop Experience and the Server Core installation scenarios support AD CS. Server Core installation provides a smaller attack surface and less administrative overhead, and therefore, you should consider it for AD CS in an enterprise environment. In Windows Server, you also can use Windows PowerShell to deploy and manage the AD CS role.

You should be aware that you cannot change computer names, domain name, or computer domain memberships after you deploy a CA of any type on that computer. Therefore, it is important to determine these attributes before installing a CA.

If you decide to deploy an offline, standalone root CA, you should consider the following:

- Before you issue a subordinate certificate from the root CA, make sure that you provide at least one certificate revocation list distribution point (CDP) and authority information access (AIA) location that will be available to all clients. This is because, by default, a standalone root CA has the CDP and AIA located on itself. Therefore, when you take the root CA off the network, a revocation check will fail because the CDP and AIA locations will be inaccessible. When you define these locations, you should manually copy certificate revocation list (CRL) and AIA information to that location.
- Set a validity period for CRLs that the root CA publishes to a long period of time, for example, one year. This means that you will have to turn on the root CA once per year to publish a new CRL, and then you will have to copy it to a location that is available to clients. If you fail to do so, after the CRL on the root CA expires, revocation checks for all certificates will also fail.
- Use Group Policy to publish the root CA certificate to a trusted root CA store on all server and client computers. You must do this manually because a standalone CA cannot do it automatically, unlike an enterprise CA. You can also publish the root CA certificate to AD DS by using the **certutil** command-line tool.

Demonstration: Manage CAs

In this demonstration, you will learn how to:

- Create a new template based on the Web Server template.
- Configure templates so that they can be issued.

Demonstration steps

Create a new template based on the Web Server template

1. On **SEA-ADM1**, in **Server Manager**, select **Tools**, and then select **Certification Authority**.
2. Retarget the console to point to **SEA-DC1**.

3. In the **Certification Authority** console, open the **Certificate Templates Console**.
4. Duplicate the **Web Server template**.
5. Create a new certification authority (CA) template, and then name it **Production Web Server**.
6. Configure validity for **3** years.
7. Configure the private key as exportable.
8. Publish the certificate revocation list on **SEA-DC1**.

Configure templates so that they can be issued

- Issue the certificates based on the **Production Web Server** template.

What are certificate templates?

Certificate templates define how you request and use a certificate for file encryption or email signing, for example. You configure templates on the certification authority (CA), and the Active Directory Domain Services (AD DS) database stores them. There are several different versions of templates that correlate to the operating system on the CA. Windows Server supports version 4 templates and earlier template versions.

The two types of certificate categories are certificate templates for users and certificate templates for computers. You can use both the user and computer templates for multiple purposes. You can assign permissions to certificate templates to define who can manage them and who can perform enrollment or autoenrollment. You also can update certificate templates by modifying the original certificate template, copying a template, or superseding existing certificate templates.

A certificate is a small file that contains several pieces of information about its owner. This data can include the owner's email address, the owner's name, the certificate usage type, the validity period, and the URLs for authority information access (AIA) and certificate revocation list distribution point (CDP) locations.

A certificate also contains the key pair, which is the private key and its related public key. You can use these keys in processes of validating identities, digital signatures, and encryption. The key pair that each certificate generates works under the following conditions:

- When content is encrypted with the public key, it can be decrypted only with the private key.
- When content is encrypted with the private key, it can be decrypted only with the public key.
- No other key is involved in the relationship between the keys from a single key pair.
- The private key cannot be derived in a reasonable amount of time from a public key, and vice versa.

During the enrollment process, the client generates the public/private key pair. The client then sends the public key to CA, which confirms client information, signs it with its own private key, and then sends the certificate, which includes client public key, back to the client.

You can think of a certificate as being similar to a driver's license. Many businesses accept a driver's license as a form of identification because the community accepts the license issuer (a government institution) as trustworthy. Because businesses understand the process by which someone can obtain a driver's license, they trust that the issuer verified the identity of the individual before issuing the license. Therefore, the driver's license is acceptable as a valid form of identification. A certificate trust is established in a similar way.

Certificate templates

Certificate templates allow administrators to customize the distribution method of certificates, define certificate purposes, and mandate the type of usage that a certificate allows. Administrators can create templates and then can deploy them quickly to an enterprise by using built-in GUI or command-line management tools.

Associated with each certificate template is its discretionary access control list (DACL). The DACL defines which security principals have permissions to read and configure the template and what security principals can enroll or use autoenrollment for certificates based on the template. Certificate templates and their permissions are defined in AD DS and are valid within the forest. If more than one enterprise CA is running in the AD DS forest, permission changes will affect all CAs.

When you define a certificate template, the definition of the certificate template must be available to all CAs in the forest. You do this when you store the certificate template information in the configuration-naming context of AD DS. The replication of this information depends on the AD DS replication schedule, and the certificate template might not be available to all CAs until replication completes. Storage and replication occur automatically.

Template versions

The CA in Windows Server AD CS supports four versions of certificate templates. Aside from corresponding with Windows Server operating system versions, certificate template versions also have the following functional differences:

- Version 1 templates. The only modification allowed to version 1 templates is the ability to change the permissions to read, write, allow, or disallow enrollment of the certificate template. When you install a CA, version 1 certificate templates are created by default.
- Version 2 templates. You can customize several settings in version 2 templates. The default installation of AD CS provides several preconfigured version 2 templates. You also can create version 2 templates based on the requirements of your organization. Alternatively, you can duplicate a version 1 certificate template to create a new version 2 template. You then can modify and secure the newly created version 2 certificate template. Templates must be a minimum of version 2 to support autoenrollment.
- Version 3 templates. Version 3 certificate templates support Cryptography Next Generation (CNG). CNG provides support for Suite B cryptographic algorithms such as elliptic curve cryptography. You can duplicate default version 1 and version 2 templates to upgrade them to version 3. When you use the version 3 certificate templates, you can use CNG encryption and hash algorithms for certificate requests, issued certificates, and protection of private keys for key exchange and key archival scenarios.
- Version 4 templates. Version 4 certificate templates are available only to Windows Server 2012, Windows 8, and later operating systems. To help administrators determine which operating system versions support which features, Microsoft added the **Compatibility** tab to the certificate template **Properties** tab. It marks options as unavailable in the certificate template properties, depending on the selected operating system versions of a certificate client and CA. Version 4 certificate templates also support both cryptographic service providers (CSPs) and key storage providers. You can also configure them to require renewal with the same key.

What are CRLs and CRL distribution lists?

Revocation is the process in which you disable the validity of one or more certificates. By initiating the revocation process, you publish a certificate thumbprint in the corresponding certificate revocation list (CRL). This indicates that a specific certificate is no longer valid.

An overview of the certificate revocation lifecycle is as follows:

1. You revoke a certificate from the certification authority (CA) Microsoft Management Console (MMC) snap-in. Specify a reason code and a date and time during revocation. This is optional but recommended.
2. The CRL publishes by using the CA console, or the scheduled revocation list publishes automatically based on the configured value. CRLs can publish in Active Directory Domain Services (AD DS), in a shared folder location, or on a website.
3. When client computers running Windows are presented with a certificate, they use a process to verify the revocation status by querying the issuing CA and certificate revocation list distribution point (CDP) location. This process determines whether the certificate is revoked and then presents the information to the application that requested the verification. The client computer running Windows uses one of the CRL locations specified in the certificate to check its validity.

Windows operating systems include CryptoAPI, which is responsible for the certificate revocation and status-checking processes. CryptoAPI uses the following phases in the certificate-checking process:

- Certificate discovery. Certificate discovery collects CA certificates, authority information access (AIA) information in issued certificates, and details of the certificate enrollment process.
- Path validation. Path validation is the process of verifying the certificate through the CA chain, or path, until the root CA certificate is reached.
- Revocation checking. Each certificate in the certificate chain is verified to ensure that none of the certificates are revoked.
- Network retrieval and caching. Network retrieval is performed by using Online Certificate Status Protocol (OCSP). CryptoAPI is responsible for checking the local cache first for revocation information, and if there is no match, making a call by using OCSP, which is based on the URL that the issued certificate provides.

What is an Online Responder service?

You also can use an Online Responder service, which is a more effective way to check certificate revocation status. By using the OCSP, an Online Responder service provides clients with an efficient way to determine the revocation status of a certificate. OCSP submits certificate status requests by using HTTP.

Clients access CRLs to determine the revocation status of a certificate. CRLs might be large, and clients might use a significant amount of time to search through these CRLs. An Online Responder service can search these CRLs dynamically for the clients and respond to the client about the status of the requested certificate. You can use a single Online Responder to determine revocation status information for certificates that are issued by a single CA or by multiple CAs. You also can use more than one Online Responder to distribute CA revocation information.

You should install an Online Responder and a CA on different computers. You must configure the CAs to include the URL of the Online Responder in the AIA extension of issued certificates. The OCSP client uses this URL to validate the certificate status. You also must issue the OCSP Response Signing certificate template, so the Online Responder also can enroll that certificate.

Configure trust for certificates

Certificates play a crucial role in helping to manage authentication and security-related issues. Therefore, it's important that you understand about certificate trust. In general, for a certificate to have meaning as a form of identity, an authority that the recipient of the certificate trusts must issue it.

Note: A passport is only useful as a form of identity if a legitimate and recognized authority issues it. It is not enough for the passport to represent a good likeness of the individual presenting it.

When using certificates for different purposes, it is important that you consider who (or what) might be expected to assess the digital certificate as a form of proof of identity. There are three types of certificates that you can use:

- Internal certificates from a private certification authority (CA) such as a server installed with the Active Directory Certificate Services (AD CS) role.
- External certificates from a public CA such as an organization on the internet that provides cybersecurity software or identity services.
- A self-signed certificate.

Understand certificate trust

If you deploy an internal public key infrastructure (PKI), and distribute certificates to your users' devices, those certificates are issued to those devices from a trusted authority. However, if some of your users use devices that are not part of your Active Directory Domain Services (AD DS) environment, those devices will not trust certificates issued by your internal CAs. To mitigate this issue, you can:

- Obtain public certificates from an external CA for those devices. This comes with a cost attached.
- Configure your users' devices to trust the internal CA. This requires additional configuration.

Manage certificates in Windows

You can manage certificates that are stored in the local machine by using Windows PowerShell, Windows Admin Center, or by using the management console with the Certificates snap-in. The easiest way to access this is to search for *certificates* in **Settings**. You can then choose to manage certificates assigned to the user or to the local computer. In either case, you will be able to access many certificates folders, including the following nodes:

- **Personal.** Contains certificates issued to the local device or local user, depending on whether you are observing the computer or the user certificate store.
- **Trusted Root Certification Authorities** Contains certificates for the CAs you trust. Sometimes called the **Root**.
- **Enterprise Trust.** Certificates here define which CAs your device trusts for user authentication.
- **Intermediate Certification Authorities.** Certificates here are used to verify the path, or chain, of trusts.

To enable a computer to trust your internal certificates, you must export your enterprise CA's root certificate, and then distribute it for import into the **Trusted Root Certification Authorities (Root)** node on all appropriate devices.

Note: You can also use the **certutil.exe** command line tool to import certificates.

Create a self-signed certificate for testing purposes

Self-signed certificates can be useful for testing purposes. You can use the Windows PowerShell **New-SelfSignedCertificate** cmdlet to create a self-signed certificate for testing purposes. By using the *CloneCert* parameter, you can create a test certificate based on an existing certificate, which has all

settings copied from the original certificate except for the public key. The cmdlet creates a new key of the same algorithm and length.

The following example creates a self-signed SSL server certificate in the computer MY store with the subject alternative name set to `www.fabrikam.com`, `www.contoso.com` and Subject and Issuer name set to `www.fabrikam.com`.

```
New-SelfSignedCertificate -DnsName "www.fabrikam.com", "www.contoso.com"  
-CertStoreLocation "cert:\LocalMachine\My"
```

Demonstration: Enroll for a certificate

In this demonstration, you will learn how to enroll a Web Server certificate on **sea-adm1**.

Demonstration steps

Enroll the Web Server certificate on sea-adm1

1. Install the web server feature on the server.
2. Open **Internet Information Services (IIS) Manager**.
3. Enroll for a domain certificate by using the following settings:
 - Common name: `sea-adm1.contoso.com`
 - Organization: **Contoso**
 - Organizational unit: **IT**
 - City/locality: **Seattle**
 - State/province: **WA**
 - Country/region: **US**
 - Friendly name: **sea-adm1**
4. Create an HTTPS binding for the **Default Web Site**, and then associate it with the **sea-adm1** certificate.

Test Your Knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

In general, what are the three categories of certification authority (CA) hierarchies?

Question 2

What is certificate revocation?

Module 02 lab and review

Lab 02: Implementing identity services and Group Policy

Scenario

You are working as an administrator at Contoso Ltd. The company is expanding its business with several new locations. The Active Directory Domain Services (AD DS) Administration team is currently evaluating methods available in Windows Server for rapid and remote domain controller deployment. The team is also searching for a way to automate certain AD DS administrative tasks. Additionally, the team wants to establish configuration management based on Group Policy Objects (GPO) and enterprise certification authority (CA) hierarchy.

Objectives

After completing this lab, you'll be able to:

- Deploy a new domain controller on Server Core.
- Configure Group Policy.
- Deploy, manage, and use digital certificates.

Estimated Time: 60 minutes

Module review

Use the following questions to check what you've learned in this module.

Question 1

What are the two reasons to create organizational units (OUs) in a domain?

Question 2

If the domain controller that holds the primary domain controller (PDC) Emulator operations master role is going to be offline for an extended period, what should you do?

Question 3

True or false? Azure Active Directory (Azure AD) is hierarchical.

Question 4

If you have a new version of Microsoft Office to deploy in your on-premises environment, and you want to configure settings with GPOs, what would you do?

Question 5

What is a certificate template?

Answers

Question 1

What is the Active Directory Domain Services (AD DS) schema?

The AD DS schema is the component that defines all the object classes and attributes that AD DS uses to store data.

Question 2

Is the Computers container an organizational unit (OU)?

No, it is an object of the Container class.

Question 1

What's a domain controller?

A domain controller is a server that stores a copy of the Active Directory Domain Services (AD DS) directory database (Ntds.dit) and a copy of the SYSVOL folder. All domain controllers except read-only domain controllers (RODCs) store a read/write copy of both Ntds.dit and the SYSVOL folder.

Question 2

What is the primary domain controller (PDC) Emulator?

The PDC Emulator is an operations master role. The domain controller that holds the PDC emulator master is the time source for the domain. The PDC emulator masters in each domain in a forest synchronize their time with the PDC emulator master in the forest root domain. You set the PDC emulator master in the forest root domain to synchronize with a reliable external time source. The PDC emulator master is also the domain controller that receives urgent password changes. If a user's password changes, the domain controller holding the PDC emulator master role receives this information immediately. This means that if the user tries to sign in, the domain controller in the user's current location will contact the domain controller holding the PDC emulator master role to check for recent changes. This will occur even if the user has been authenticated by a domain controller in a different location that had not yet received the new password information.

Question 1

What would you use to synchronize user details to Microsoft Azure Active Directory (Azure AD) from Active Directory Domain Services (AD DS)?

You would use Azure AD Connect to synchronize user details to Azure AD from AD DS.

Question 2

Which version of Azure AD supports Azure AD Join + mobile device management (MDM) autoenrollment?

Both Azure AD Premium P1 and P2 editions support this feature.

Question 1

If you linked a Group Policy Object (GPO) to the domain object in your Active Directory Domain Services (AD DS), what are the different ways to prevent this policy from applying to all users in the domain?

There are several possible approaches, which include:

You could unlink the GPO from the domain. You could use Security Group filtering to target the GPO settings to a specific group. You could also use Block Inheritance on child OUs.

Question 2

What are the default domain GPOs called?

The two default GPOs are called Default Domain Policy and Default Domain Controllers Policy.

Question 1

In general, what are the three categories of certification authority (CA) hierarchies?

The three categories of CA hierarchies include CA hierarchies with a policy CA, CA hierarchies with cross-certification trust, and CAs with a two-tier hierarchy.

Question 2

What is certificate revocation?

Revocation is the process in which you disable the validity of one or more certificates. By initiating the revocation process, you publish a certificate thumbprint in the corresponding certificate revocation list (CRL). This indicates that a specific certificate is no longer valid.

Question 1

What are the two reasons to create organizational units (OUs) in a domain?

The first reason is because you want to group users and computers, for example, by geography or department. The second reason is that you might want to delegate administration on the OU or configure the objects in an OU by using Group Policy Objects (GPOs).

Question 2

If the domain controller that holds the primary domain controller (PDC) Emulator operations master role is going to be offline for an extended period, what should you do?

You should transfer the operations master role to another server in the same domain ahead of the planned outage.

Question 3

True or false? Azure Active Directory (Azure AD) is hierarchical.

False. Azure AD has a flat structure.

Question 4

If you have a new version of Microsoft Office to deploy in your on-premises environment, and you want to configure settings with GPOs, what would you do?

You could download and install the latest .admx files for Office. If you install these into the Central Store, you could configure the new Office settings in one location.

Question 5

What is a certificate template?

Certificate templates define how you can request or use a certificate, such as for file encryption or email signing.



Module 3 Network infrastructure services in Windows Server

Deploying and managing DHCP

Lesson overview

Dynamic Host Configuration Protocol (DHCP) is used to dynamically configure devices with an IP address and other IP configuration information, such as a default gateway and Domain Name System (DNS) server. You can configure Windows Server to be a DHCP server. The scopes you configure on the DHCP server contain the information that is provided to DHCP clients. Before a Windows-based DHCP server can begin servicing clients, you need to authorize it. To ensure that DHCP service is not interrupted for clients, you can configure DHCP Failover.

Lesson objectives

After completing this lesson, you'll be able to:

- Describe the DHCP Server role.
- Describe how to install and configure the DHCP Server role.
- Configure DHCP options.
- Configure the DHCP Server role.
- Describe how to configure DHCP scopes.
- Create and configure a DHCP scope.
- Identify how to authorize a DHCP server in Active Directory Domain Services (AD DS).
- Describe high availability options for DHCP.
- Describe how to use DHCP Failover.

Overview of the DHCP role




Figure 1: DHCP server and DHCP client communication process

The Dynamic Host Configuration Protocol (DHCP) is used to automatically configure network devices with IP address configuration information. If you don't use DHCP, each time you add a client to a network, you need to configure its network interface with information about the network you connect it to. The information that you must configure includes the IP address, the network's subnet mask, and the default gateway for access to other networks. The four step communication process is demonstrated in Figure 1.

Benefits of DHCP

The main benefit of using DHCP is reducing the maintenance required to configure IP address information on network devices. Many organizations manage thousands of computer devices, including printers, scanners, smartphones, desktop computers, and laptops. Because of this, performing manual management of the network IP configurations for organizations of this size isn't practical.

Because DHCP is an automated process, it is more accurate than manually configuring IP address information. This is particularly important for users that wouldn't know or understand the configuration process.

DHCP makes it easier to update IP address configuration information. As an administrator, when you make a network service change, such as providing a new Domain Name System (DNS) server, you only make a single update on the DHCP servers, and that change is received by all of the DHCP clients. For example, a mobile user with a laptop using DHCP automatically gets new IP address configuration information when they connect to a new network.

Note: By default, all Windows operating systems are configured to automatically get an IP address after the initial installation of the operating system (OS).

How DHCP works

The DHCP Client service runs on all Windows computers that have their TCP/IP properties set to automatically obtain an IP Address. The DHCP client communicates with a DHCP Server to obtain IP configuration information. Clients can use the assigned DHCP address for a certain period, known as a *lease*. The DHCP

server is configured with an address pool and configuration options. This information determines what IP address configuration information is handed out to clients.

Communication for DHCP lease generation uses IP broadcasts. Because IP broadcasts are not routed, you need to configure a DHCP server on each subnet or configure a DHCP relay. Many routers include DHCP relay functionality.

The four steps in lease generation are:

1. The DHCP client broadcasts a DHCPDISCOVER packet. The only computers that respond are computers that have the DHCP Server role, or computers or routers that are running a DHCP relay agent. In the latter case, the DHCP relay agent forwards the message to the DHCP server that you have configured to relay requests.
2. A DHCP Server responds with a DHCPOFFER packet, which contains a potential address for the client. If multiple DHCP servers receive the DHCPDISCOVER packet, then multiple DHCP servers can respond.
3. The client receives the DHCPOFFER packet. If multiple DHCPOFFER packets are received, the first response is selected. The client then sends a DHCPREQUEST packet that contains a server identifier. This informs the DHCP servers that receive the broadcast which server's DHCPOFFER the client has chosen to accept.
4. The DHCP servers receive the DHCPREQUEST. Servers that the client has not accepted use this message as the notification that the client has declined that server's offer. The chosen server stores the IP address-client information in the DHCP database and responds with a DHCPACK message. If the DHCP server can't provide the address that was offered in the initial DHCPOFFER, the DHCP server sends a DHCPNAK message.

DHCP lease renewal

When the DHCP lease reaches 50 percent of the lease time, the client automatically attempts to renew the lease. This process occurs in the background. It is possible for a computer to have the same DHCP-assigned IP address for a long time. This is because the computer renews the lease multiple times.

To attempt to renew the IP address lease, the client sends a unicast DHCPREQUEST message. The server that originally leased the IP address sends a DHCPACK message back to the client. This message contains any new parameters that have changed since the original lease was created. Note that these packets do not broadcast, because at this point the client has an IP address that it can use for unicast communications.

Note: When you update DHCP configuration options, clients might not get the updated options until 50 percent of the lease time is complete. For example, if you configure a six-day lease time, clients might not get updated options for three days.

If the DHCP client cannot contact the DHCP server, then the client waits until 87.5 percent of the lease time expires. At this point, the client sends a DHCPREQUEST broadcast (rather than a unicast) to obtain a renewal, and the request goes to all DHCP servers, not just the server that provided the original lease. However, this broadcast request is for a renewal, not a new lease.

Because client computers might be moved while they are turned off (for example, a laptop computer that is plugged into a new subnet), client computers also attempt renewal during the startup process, or when the computer detects a network change. If renewal is successful, the lease period will reset.

DHCP version 6

DHCP version 6 (DHCPv6) stateful and stateless configurations are supported for configuring clients in an IPv6 environment. Stateful configuration occurs when the DHCPv6 server assigns the IPv6 address to the client, along with additional DHCP data. Stateless configuration occurs when the router assigns the IPv6 address automatically, and the DHCPv6 server only assigns other IPv6 configuration settings.

Install and configure the DHCP role

You can install the Dynamic Host Configuration Protocol (DHCP) Server role only on Windows Server operating systems. You can install the DHCP server on a domain controller, and any server that is running Windows Server can host the DHCP server. For example, a branch office file and print server also might function as the local DHCP server. In addition, you must have local administrative rights to perform the installation, and the server must have a static IP address.

Install the DHCP Server role

You can install the DHCP Server role by using **Roles & Features** in Windows Admin Center, the **Add Roles and Features Wizard** in the **Server Manager** console, or by using the following Windows PowerShell command:

```
Add-WindowsFeature DHCP -IncludeManagementTools
```

Note: The *-IncludeManagementTools* parameter is optional.

DHCP management tools

When you install the DHCP Server role by using **Server Manager**, the DHCP management tools are also installed that server by default. The Windows PowerShell cmdlets are installed, and if the server includes the Desktop Experience, then the DHCP Management console is installed too. You can install the **DHCP management** console or Windows PowerShell cmdlets from the Remote Server Administration Tools (RSAT) on a Windows server or Windows client for remote administration.

To manage a DHCP server by using Windows Admin Center, the DHCP management cmdlets need to be installed on the DHCP server. If they aren't installed, Windows Admin Center displays the message "DHCP PowerShell tools (RSAT) are not installed" and provides an install button to perform the installation remotely.

DHCP management groups

To delegate management of DHCP, the **DHCP Administrators** local group and the **DHCP Users** local group are created on each DHCP server. The **DHCP Administrators** group can manage the local DHCP server, and the **DHCP Users** group can examine configuration and status information on the local DHCP server.

When you use **Server Manager** to install the DHCP Server role, the DHCP Post-Install Configuration Wizard creates both groups. If you use Windows Admin Center or Windows PowerShell to install the DHCP role, then you need to manually trigger the creation of the groups.

To create the DHCP management groups by using Windows PowerShell, run the following command:

```
Add-DhcpServerSecurityGroup -Computer DhcpServerName
```

Additional reading: For additional information about managing servers in Windows Admin Center, refer to **Manage Servers with Windows Admin Center**¹.

Configure DHCP options

Dynamic Host Configuration Protocol (DHCP) servers can configure more than just an IP address. They also provide information about network resources, such as Domain Name System (DNS) servers and the default gateway. DHCP options are values for common configuration data that apply to the server, scopes, reservations, and class options. You can apply DHCP options at the server, scope, class, and reservation levels. An option code identifies the DHCP options, and most option codes come from the Request for Comments (RFC) documentation found on the Internet Engineering Task Force (IETF) website.

The following table lists common option codes that Windows DHCP clients use.

Table 1: Option codes

Option code	Name
1	Subnet mask
3	Router
6	DNS servers
15	DNS domain name
31	Perform router discovery
33	Static route
43	Vendor-specific information
47	NetBIOS scope ID
51	Lease time
58	Renewal (T1) time value
59	Rebinding (T2) time value
60	Pre-Boot Execution (PXE) client
66	Boot server host name
67	Bootfile name
249	Classless static routes

How DHCP options are applied

The DHCP client service applies the options in an order of precedence at four different levels. Going from least specific to most specific, they are:

1. Server level. Assigns a server-level option to all DHCP clients of the DHCP server.
2. Scope level. Assigns a scope-level option to all clients of a scope. Scope options override server options.
3. Class level. Assigns a class-level option to all clients that identify themselves as members of a class. Class options override both scope and server options.
4. Reserved client level. Assigns a reservation-level option to one DHCP client. Reserved client options apply to devices that have a DHCP reservation.

If you apply DHCP option settings at each level and they conflict, the option that you applied last overrides the previously applied setting. For example, if you configure the default gateway at the scope level,

¹ <https://aka.ms/manage-servers-windows-admin-center>

and then apply a different default gateway for a reserved client, the reserved client setting is the effective setting.

Note: Currently, you can't manage server-level DHCP options by using Windows Admin Center, and there are only a few scope-level options that you can manage.

Where to configure DHCP options

Most DHCP options are configured at the scope level because the scope level represents a subnet where all clients have the same configuration needs. For example, the router and subnet mask options are unique for a subnet and you should configure them at the scope level.

It's common to configure the DNS server option at the server level because clients on multiple subnets in a location typically use the same DNS servers. For example, all clients at the London site use the same DNS servers even though the London site has clients on several subnets. You can also configure the DNS domain name at the server level if all clients are joined to the same Active Directory Domain Services (AD DS) domain.

The class level is used to support specialized devices such as IP phones. The devices report their vendor class as part of the DHCP leasing process and the DHCP server provides the DHCP options specific to that vendor class. This is useful when the specialized devices are on a different virtual LAN (VLAN) but still part of the same broadcast domain. For example, when an office has a single network connection, and an IP phone and computer use the same network connection.

Note: The class level includes user classes and vendor classes. You can manually configure the user class on computers running Windows operating systems. Typically, the manufacturer of the device configures the vendor class.

Demonstration: Configure the DHCP role

In this demonstration, you will learn how to install and configure the DHCP Server role.

Demonstration steps

Install the DHCP Server role

1. On **SEA-ADM1**, open Microsoft Edge and sign in to Windows Admin Center as **Contoso\Administrator** with the password **Pa55w.rd**.
2. In Windows Admin Center, connect to **SEA-SVR1**.
3. Use **Roles & features** to install the DHCP role.

Install the DHCP PowerShell tools

1. Reconnect to **SEA-SVR1** and notice DHCP on the **Tools** pane.
2. Use DHCP to install the **DHCP PowerShell tools** on **SEA-SVR1**.
3. Use PowerShell to run `Add-DhcpServerSecurityGroup` on **SEA-SVR1**.

Configure a DHCP server option

1. Open **Server Manager** and start the **DHCP management** console.

2. In the **DHCP management** console, add **SEA-SVR1**.
3. For IPv4, enable the **006 DNS Servers** option with a value of **172.16.10.10**.

Configure DHCP scopes

A Dynamic Host Configuration Protocol (DHCP) scope is a range of IP addresses that are available for lease and that a DHCP server manages. A DHCP scope typically is confined to the IP addresses in a given subnet. For example, a DHCP scope for the network 192.168.1.0/24 can support a range from 192.168.1.1 through 192.168.1.254. When a computer or device on the 192.168.1.0/24 subnet requests an IP address, the scope that defined the range in this example allocates an address between 192.168.1.1 and 192.168.1.254.

Although it's not typical, a DHCP server could host scopes for multiple different subnets, in which case DHCP relay agents would distribute those addresses to clients on other subnets.

You do not need to assign all IP addresses in a given subnet to the scope. Usually, some IP addresses are excluded from the scope so that they are available for assignment as static addresses. For example, the first 20 addresses of the scope might be excluded and then statically assigned to routers, printers, and servers on the subnet.

DHCP scope properties

To create and configure a scope, you must define the following properties:

- Name and description. This property identifies the scope. The name is mandatory.
- IP address range. This property lists the range of addresses that can be offered for lease. This property is mandatory.
- Subnet mask. Client computers use this property to determine their location in the organization's network infrastructure. This property is mandatory.
- Exclusions. This property lists single addresses or blocks of addresses that are within the IP address range, but that will not be offered for lease. This property is optional.
- Delay. This property indicates the amount of time to delay before sending DHCP OFFER. The default setting is 0 milliseconds.
- Lease duration. This property lists the lease duration. Use shorter durations for scopes that have limited IP addresses and use longer durations for more-static networks.
- Options. You can configure many optional properties on a scope, but typically you configure the following properties:
 - Option 003. Router (the default gateway for the subnet)
 - Option 006. DNS servers
 - Option 015. DNS suffix
- Activation. You must activate the scope before it can lease IP addresses.

Creating DHCP scopes by using Windows PowerShell

You can use Windows PowerShell to configure DHCP scopes and retrieve information. This is useful when you want to create scripts that automate scope management tasks. The following table lists some of the Windows PowerShell cmdlets that you can use for scope management.

Table 1: PowerShell cmdlets for scope management

Cmdlet name	Description
Add-DhcpServerv4Scope	Adds an IPv4 scope on the DHCP server.
Get-DhcpServerv4Scope	Returns the IPv4 scope configuration of the specified scopes.
Get-DhcpServerv4ScopeStatistics	Gets the IPv4 scope statistics corresponding to the IPv4 scope identifiers specified for a DHCP server service.
Remove-DhcpServerv4Scope	Deletes the specified IPv4 scopes from the DHCP server service.
Set-DhcpServerv4Scope	Sets the properties of an existing IPv4 scope on the DHCP server.

Additional reading: For a list of all available DHCP management cmdlets, refer to [DhcpServer²](#).

DHCP reservations

If you want a computer or device to obtain a specific address from the scope range, you can permanently reserve that address for assignment to that device in DHCP. Reservations are useful for tracking IP addresses assigned to devices such as printers. To create a reservation, select the scope in the **DHCP** console, and from the **Action** menu, select **New Reservation**. You need to provide the following information to create the reservation in the **New Reservation** dialog box:

- Reservation name. A friendly name to reference the reservation.
- IP address. The IP address from the scope that you want to assign to the device.
- MAC address. The MAC address of the interface to which you want to assign the address.
- Description. An optional field in which you can provide a comment about the reservation.

Note: If a client has already obtained an IP address from a DHCP server, you can convert the existing lease to a reservation in the **DHCP** console.

Demonstration: Create and configure a DHCP scope

In this demonstration you will learn how to create and configure and DHCP scope.

Demonstration steps

Create a DHCP scope

1. On **SEA-ADM1**, open **Microsoft Edge** and sign in to **Windows Admin Center** as **Contoso\Administrator** with the password **Pa55w.rd**.
2. In **Windows Admin Center**, connect to **SEA-SVR1**.

² <https://aka.ms/dhcpserver>

3. Use **Dynamic Host Configuration Protocol (DHCP)** to create a new scope with the following information:

- Protocol: **IPv4**
- Name: **ContosoClients**
- Starting IP address: **10.10.100.10**
- Ending IP address: **10.10.100.200**
- DHCP client subnet mask: **255.255.255.0**
- Router: **10.10.100.1**
- Lease duration: **4 days**

Create a DHCP reservation

• Create a new reservation in the **ContosoClientsScope** with the following information:

- Reservation name: **Printer**
- IP address: **10.10.100.199**
- MAC address: **00-14-6D-01-73-6B**

DHCP AD DS authorization

Dynamic Host Configuration Protocol (DHCP) communication typically occurs before any user or computer authentication. Because the DHCP protocol is based on IP broadcasts, an unknown DHCP server can provide invalid information to clients. You can avoid this by authorizing the server. As the domain administrator, you can use a process called DHCP authorization to register the DHCP server in the Active Directory domain before it can support DHCP clients. Authorizing the DHCP server is one of the post-installation tasks that you must perform after you install the DHCP server.

Active Directory requirements

You must authorize the Windows Server DHCP Server role in Active Directory Domain Services (AD DS) before it can begin leasing IP addresses. It's possible to have a single DHCP server providing IP addresses for subnets that contain multiple AD DS domains. Because of this, you must use an Enterprise Administrator account to authorize the DHCP server. In a single-domain environment, membership in **Domain Admins** is sufficient to authorize a DHCP server.

You can authorize a DHCP server by using the DHCP management console or Windows PowerShell. To authorize a DHCP server by using Windows PowerShell, run the following command:

```
Add-DHCPServerinDC <hostname or IP address of DHCP server>
```

Standalone DHCP server considerations

A standalone DHCP server is a computer that is running Windows Server, is not a member of an AD DS, and has the DHCP Server role installed and configured. If a standalone DHCP server detects an authorized DHCP server in the domain, it does not lease IP addresses and automatically shuts down.

Unauthorized DHCP servers

Many network devices have built-in DHCP server software and can act as a DHCP server. The DHCP servers on these devices do not typically recognize authorization in AD DS. Therefore, these DHCP servers will lease IP addresses when they are connected to the network and the DHCP server software is enabled. To find these unauthorized DHCP servers, you must perform an investigation. When you detect unauthorized DHCP servers, you should disable the DHCP service on them. You can find the IP address of the unauthorized DHCP server by running the `ipconfig /all` command on the DHCP client computer that obtained the incorrect IP address information.

High availability options for DHCP

Dynamic Host Configuration Protocol (DHCP) is a critical component in most modern networks and needs to be available when clients request IP addresses. Options for making DHCP available include using the DHCP Server role in failover clustering, split scopes, and DHCP Failover.

DHCP Failover is the preferred option for highly available DHCP service. It provides the widest variety of options for configuration. DHCP Failover is also the only high availability method in which the DHCP servers actively coordinate to provide service. DHCP clustering and split scopes are high availability methods that might be used when highly available DHCP is implemented on versions of Windows Server earlier than Windows Server 2012.

Note: The topic *DHCP Failover* describes DHCP Failover.

DHCP clustering




Figure 1: A two-member DHCP server cluster. The DHCP information is stored on Shared Storage.

You can configure the DHCP Server role to run in a failover cluster. After installing the DHCP Server role on all cluster nodes and creating the failover cluster, you add the DHCP Server role to the failover cluster. As part of the configuration process, you need to provide an IP address for the DHCP server and shared storage. In this scenario, the DHCP configuration information is stored on shared storage, as illustrated in

Figure 1. If one cluster member fails, another cluster member detects the failure and starts the DHCP service to continue providing service.

Split scopes




Figure 2: DHCP servers with a split scope, where each server controls a portion of the IP address range.

A split scope scenario also involves two DHCP servers. In this case, each DHCP server controls a part of the entire range of IP addresses, and both servers are active on the same network. For example, as Figure 2 illustrates, if your subnet is 192.168.0.0/24, you might assign an IP address range of 192.168.0.1 through 192.168.0.150 to the DHCP server A, the primary server, and assign 192.168.0.151 through 192.168.0.254 to DHCP server B, which acts as a DHCP secondary server. You can control which server is the primary server assigning addresses by setting the **Delay configuration** attribute in the properties of scope on the secondary server. This ensures that the primary server will be the first server to respond to client requests. If the primary server fails and stops responding to requests, then the secondary server's response will be the one the client accepts.

DHCP Failover

The Dynamic Host Configuration Protocol (DHCP) Failover feature allows two DHCP servers to work together to provide IP address information to clients. The two DHCP servers replicate lease information between them. If one of the DHCP servers fails, the remaining DHCP server continues to use the scope information to provide IP addresses to clients.

Note: You can configure only two DHCP servers in a failover relationship, and you can configure these only for IPv4 scopes.

Configuring DHCP Failover

To configure DHCP Failover, you establish a failover relationship between the two DHCP servers and give the relationship a unique name. The failover partners exchange this name during configuration. This enables a single DHCP server to have multiple failover relationships with other DHCP servers, so long as

they all have unique names. To configure failover in the **DHCP Management** console, use the **Configuration Failover Wizard**, which you launch by right-clicking or accessing the context menu for the **IP** node or the **scope** node.

Note: DHCP Failover is time sensitive. If the time difference between the partners is greater than one minute, the failover process halts with a critical error.

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

Table 1: Configuration modes

Mode	Characteristics
Load balance	This is the default mode. In this mode, both servers supply IP configuration to clients simultaneously. Which server responds to IP configuration requests depends on how the administrator configured the load distribution ratio. The default ratio is 50:50.
Hot standby	In this mode, one server is the primary server and the other is the secondary server. The primary server actively assigns IP configurations for the scope or subnet. The secondary DHCP server assumes this role only if the primary server becomes unavailable. A DHCP server can act simultaneously as the primary server for one scope or subnet, and the secondary server for another.

As an administrator, you must configure a percentage of the scope addresses to be assigned to the standby server. These addresses are supplied during the Maximum Client Lead Time (MCLT) interval if the primary server is down. By default, five percent of the scope addresses are reserved for the standby server. The secondary server takes control of the entire IP range after the MCLT interval has passed.

Hot standby mode is best for deployments in which a disaster-recovery site is at a different location. This way the DHCP server will not service clients unless there is a main server outage.

MCLT

Configure the MCLT parameter to specify the amount of time that a DHCP server should wait when a partner is unavailable before it assumes control of the address range. The default value is one hour, and it can't be zero. If required, you can adjust the MCLT by using Windows PowerShell.

Auto state switchover interval

A communication-interrupted state occurs when a server loses contact with its partner. Because the server has no way of knowing what is causing the communication loss, it remains in this state until the administrator manually changes it to a **partner down** state. You can enable automatic transition to **partner down** state by configuring the auto state switchover interval. The default value for this interval is 60 minutes.

Message authentication

Windows Server enables you to authenticate the failover message traffic between the replication partners. The administrator can establish a shared secret—much like a password—in the **Configuration**

Failover Wizard for DHCP Failover. This validates that the failover message comes from the failover partner.

Firewall considerations

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

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

Test your knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

If you configure a DHCP scope with a lease length of four days, when will computers attempt to renew the lease for the first time?

- 1 day
- 2 days
- 3 days
- 3.5 days

Question 2

Which permissions are required to authorize a DHCP server in a multiple domain AD DS forest?

- Member of "Enterprise Admins" group
- Member of "Domain Admins" group
- Member of local "Administrators" group on the DHCP server
- Member of "DHCP Administrators"

Deploying and managing DNS services

Lesson overview

Active Directory Domain Services (AD DS) and general network communication require Domain Name System (DNS) as a critical network service. Windows Server is often used as a DNS server in companies that use AD DS. The DNS clients make requests to DNS servers that host DNS zones, which contain resource records. You can create these resource records manually or the records can be created by dynamic DNS. If the DNS server doesn't have the required information, it can use root hints or forwarding to find the required information. Additionally, DNS policies allow you to provide different DNS information to groups of users or computers.

After configuring DNS, to enhance security, you should consider implementing Domain Name System Security Extensions (DNSSEC) which digitally signs DNS records to verify authenticity.

Lesson objectives

After completing this lesson, you'll be able to:

- List the components in DNS name resolution.
- Describe DNS zones.
- Describe DNS records.
- Install and configure the DNS server role.
- Manage DNS services.
- Create records in DNS.
- Configure DNS zones.
- Describe DNS forwarding.
- Understand DNS integration in AD DS.
- Describe DNS policies.
- Describe how to use DNSSEC.

DNS components

The most common use for Domain Name System (DNS) is resolving host names, such as `dc1.contoso.com`, to an IP address. Users require this functionality to access network resources and websites. Administrators use name resolution in DNS when configuring apps and when managing them. It's much easier to remember names than IP addresses. Domain-joined Windows clients and servers also use DNS to locate domain controllers in an Active Directory Domain Services (AD DS) domain.

DNS domain names

The naming structure used in DNS is called the *DNS namespace*. It is hierarchical, which means that it starts with a root domain. That root domain can itself have any number of subdomains underneath it. Each subdomain can, in turn, have any number of subdomains underneath it.

Domain names can be either public (internet-facing), or private. If they are private, you decide on your own how to define your namespace. If they are public, you must work with the Internet Corporation for

Assigned Names and Numbers (ICANN) or other internet naming registration authorities that can delegate, or sell, unique names to you. From these names, you can create subnames.

Note: To aid in obtaining trusted certificates for apps and authentication, it is typical to use a public domain name that is registered on the internet.

DNS servers

A DNS server responds to requests for DNS records that are made by DNS resolvers. For example, a Windows 10 client can send a DNS request to resolve `dc1.contoso.com` to a DNS server, and the DNS server response includes the IP address of `dc1.contoso.com`. A DNS server can retrieve this information from a local database that contains resource records. Alternatively, if the DNS server doesn't have the requested information, it can forward DNS requests to another DNS server. A DNS server can also cache previously requested information from other DNS servers.

When AD DS is used, it's common to have domain controllers that are also configured as DNS servers. However, it's possible to use member servers or other devices as DNS servers.

Note: Windows Server is configured to be a DNS server when you install the DNS server role.

DNS zones and resource records

When a DNS server is responsible for resolving requests for a specific namespace, you create a zone on the DNS server that corresponds to the namespace. For example, if a DNS server is responsible for `contoso.com`, you would create a `contoso.com` zone. Inside the zone, you create resource records that contain the information that's used to respond to queries.

DNS resolvers

A DNS resolver is a client—such as a Windows client—that needs to resolve DNS records. In Windows, the DNS Client service sends DNS requests to the DNS server configured in the properties of IP. After receiving a response to a DNS request, the response is cached for future use. This is called the *DNS resolver cache*.

Note: You can access the contents of the DNS resolver cache by using the **Get-DnsClientCache** cmdlet. You can clear the contents of the DNS resolver cache by using the **Clear-DnsClientCache** cmdlet.

You can manually configure name resolution in Windows by editing the Hosts file located in **C:\Windows\System32\Drivers\etc**. You can add a simple name to IP address mapping in a Hosts file, but not more complex resource records. When you enter information into the Hosts file, it overrides information found in the local DNS resolver cache, which includes information that was resolved by using DNS.

What are DNS zones?

A *DNS zone* is the specific portion of a Domain Name System (DNS) namespace (such as `contoso.com`) that is hosted on a DNS server. A DNS zone contains resource records, and the DNS server responds to queries for records in that namespace. For example, the DNS server that is authoritative for resolving `www.contoso.com` to an IP address would contain the `contoso.com` zone.

You can store DNS zone content in a file or in the Active Directory Domain Services (AD DS) database. When the DNS server stores the zone in a file, that file is located in a local folder on the server. When the zone is not stored in AD DS, only one copy of the zone is a writable copy, and all the other copies are read only.

Forward lookup zones

Forward lookup zones can hold a wide variety of different resource records, but the most common record type is a *host (A) record*. A host record is used to resolve a name to an IP address. Figure 1 has a client querying a host record in a forward lookup zone.

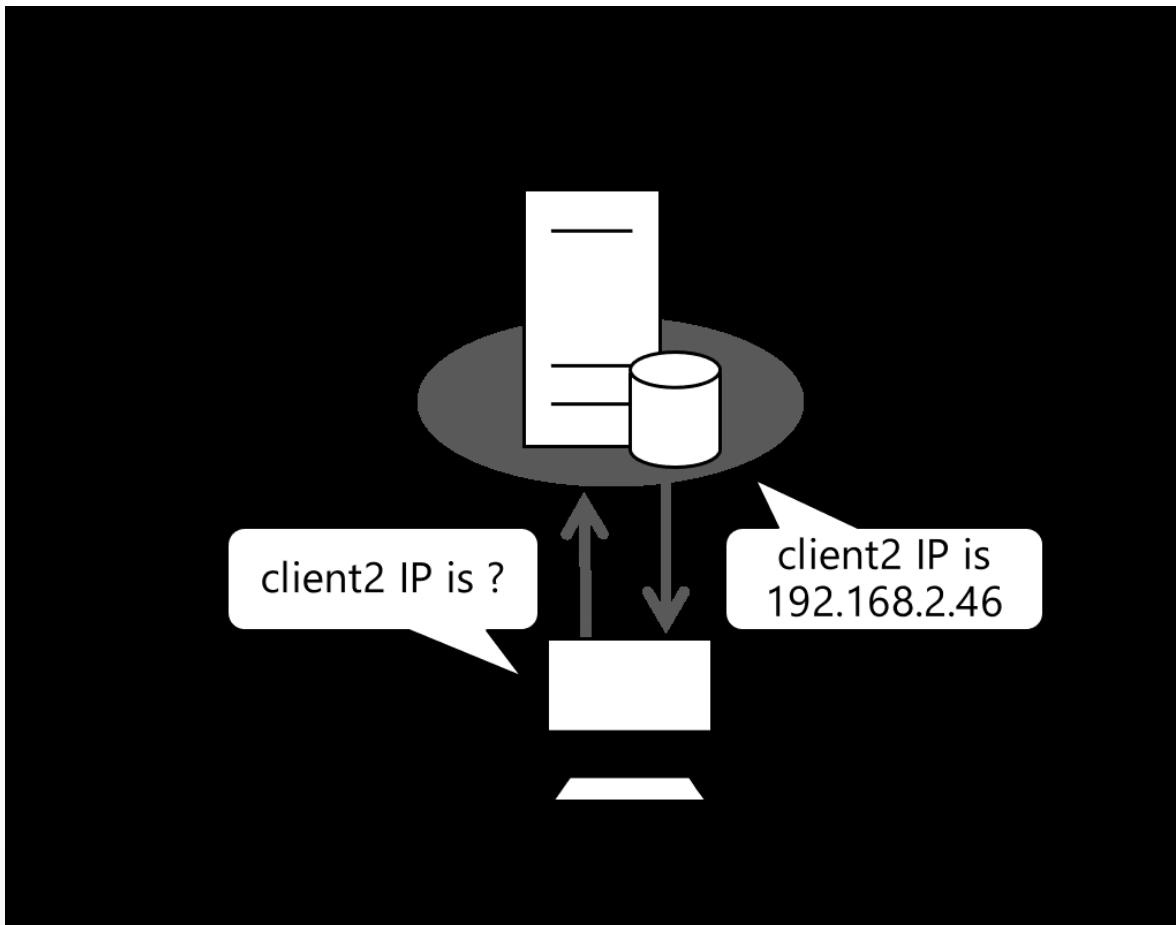


Figure 1: DNS client petitioning its configured DNS server

When using AD DS, the internal DNS servers for your organization have a zone that corresponds to the AD DS domain. For example, if the domain name for AD DS is `contoso.com`, then there will also be a `contoso.com` DNS name. AD DS is used to store resource records that Windows servers and clients use to locate network services. There is no requirement to make the DNS records containing AD DS resource information available on the internet.

If you are providing name resolution for a zone to internet clients, you can host the zone on a Windows server that is accessible on the internet. You also have the option to host the zone on a third-party DNS service that specializes in providing internet name resolution.

Reverse lookup zones

Reverse lookup zones are used only for resolving an IP address to a name. A variety of apps, and sometimes administrators, use this functionality. For example, an administrator might notice a specific IP address in a log file and use a reverse lookup to identify the name of the computer that corresponds with the IP address.

You create reverse lookup zones only for IP address ranges for which you are responsible. As a best practice, you should create reverse lookup zones for all the IP address ranges on your internal network and host them on your internal DNS servers. The zone name for reverse lookup zones ends with `in-addr.arpa` and is based on the IP address range. For example, the zone name for the **172.16.35.0/24** reverse lookup zone will be **35.16.172.in-addr.arpa**. Reverse lookup zones are always based on a full octet of the IP address.

The internet service provider that provides internet routable IP addresses for your organization often maintains the reverse lookup zones for those IP addresses. If you have been allocated a large block of internet routable IP addresses, you might have the option to maintain your own reverse lookup zone for those IP addresses.

Primary and secondary zones

When you create a zone on a DNS server, you need to identify whether it's a primary zone or a secondary zone. You can only create, edit, or delete resource records in the primary zone. You can't manage the records in a secondary zone.

You can store a standard primary zone in a local file, or you can store the zone data in AD DS. When you store the zone data in AD DS, the zone is called *Active Directory-integrated* and allows additional features, such as secure dynamic updates. Active Directory-integrated zones are available only on domain controllers with the DNS Server role installed. Most Windows-based DNS servers use Active Directory-integrated zones.

A secondary zone is a read-only copy of a primary zone. In most cases, a secondary zone periodically copies resource records directly from the primary zone. But in some complex configurations, a secondary zone can copy resource records from another secondary zone.

Stub zones

The purpose of a stub zone is to provide a list of name servers that can be used to resolve information for a domain without synchronizing all the records locally. To enable this, the following are synchronized: name server records, their corresponding host records, and the start of authority record. You would typically use stub zones when integrating with autonomous systems such as partner organizations.

What are DNS records?

DNS records are the resource records stored in DNS zones. The DNS records contain the information that DNS servers send in response to DNS requests. All forward lookup and reverse lookup DNS zones contain the following records:

- Start of authority (SOA). A *start of authority* record for a zone contains configuration information for the zone, including the name of the primary DNS server and how often secondary servers should be synchronized. There is one start of authority record per zone.
- Name server (NS). A *name server* record identifies a DNS server for the domain. There is one name server record for each DNS server that has a copy of the zone.

Resource records in forward lookup zones

The following table describes some of the resource records available in forward lookup zones.

Table 1: Resource records available in forward lookup zones

DNS record type	Description
Host (A)	Host (A) records are used to resolve a name to an IPv4 address. You can create multiple host records with a single name to allow a name to resolve to multiple IPv4 addresses.
Host (AAAA)	Host (AAAA) records are used to resolve a name to an IPv6 address.
Alias (CNAME)	Alias records are used to resolve a name to another name. For example, an alias can resolve <code>app.contoso.com</code> to <code>sea-svr1.contoso.com</code> . In some cases, this makes it easier to reconfigure names because clients are not pointed at a specific server.
Service location (SRV)	Service location records are used by applications to identify the location of servers hosting that application. For example, Active Directory Domain Services (AD DS) uses service location records to identify the location of domain controllers and global catalog servers.
Mail exchanger (MX)	Mail exchanger records are used to identify email servers for a domain. There can be multiple mail exchanger records for a domain for redundancy.
Text (TXT)	Text records are used to store arbitrary strings of information in DNS. These are often used by services to validate control of a namespace. For example, when you add a domain name to Microsoft 365, you can create a text record with a specified value to prove that you are the owner of the domain.

Resource records in reverse lookup zones

The most common record type created in reverse lookup zones is a pointer (PTR) record. A *pointer* record is used to resolve an IP address to a name. For example, a pointer record could be used to resolve the IP address 172.16.35.100 to `filesrv.contoso.com`. The pointer record would be named **100** and located in the **35.16.172.inaddr.arpa** zone.

Time to live

All resource records are configured with a time to live (TTL). The TTL for a resource record defines how long DNS clients and DNS servers can cache a DNS response for the record. For example, if a record has a TTL of 60 minutes, then the client makes a DNS query for the record, and the response is cached for 60 minutes. When the query result is cached, updates to the record in DNS are not recognized.

Note: When you are troubleshooting cached DNS records, you might need to clear the cache on the DNS client and on the DNS server used by that client.

Demonstration: Install and configure the DNS role

In this demonstration, you will learn how to install the DNS role and create zones.

Demonstration steps

Install the DNS Server role

1. On **SEA-ADM1**, open **Microsoft Edge** and sign in to Windows Admin Center as **Contoso\Administrator** with the password **Pa55w.rd**.
2. In **Windows Admin Center**, connect to **SEA-SVR1**.
3. Use the **Roles & features** tool to install the **DNS role** on **SEA-SVR1**.

Install the DNS PowerShell tools and create a DNS zone

1. Use the **DNS tool** to install the **DNS PowerShell** tools on **SEA-SVR1**.
2. Create a zone on **SEA-SVR1** using the following information:

- Zone type: **Primary**
- Zone name: `adatum.com`
- Zone file: **Create a new file**
- Zone file name: `adatum.com.dns`
- Dynamic update: **Do not allow dynamic update**

Create and verify a host record

1. Create a new DNS record in the `adatum.com` zone using the following information:
 - DNS record type: **Host (A)**
 - Record name: **LON-SVR5**
 - IP address: **172.30.99.234**
 - Time to live: **600**
2. Use the **Resolve-DnsName** cmdlet to verify that **SEA-SVR1** resolves the name `lon-svr5.adatum.com`.

Manage DNS services

When you implement DNS servers for your organization, your first concerns are typically which zones and resource records to create. However, there are ongoing management tasks when you implement DNS servers. You might want to delegate administration of DNS for Windows servers. For troubleshooting DNS server issues, you can enable debug logging. To clean up stale records, you can configure aging and scavenging. Finally, you can back up the DNS database.

Delegate administration of DNS

By default, the **Domain Admins** group has full permissions to manage all aspects of DNS servers in its home domain, and the **Enterprise Admins** group has full permissions to manage all aspects of all DNS servers in any domain in the forest. If you need to delegate the administration of a DNS server to a different user or group, you can add that user or global group to the **DNS Admins** group for a given domain in the forest. Members of the **DNS Admins** group can examine and modify all DNS data, settings, and configurations of DNS servers in their home domain. The **DNS Admins** group is a Domain Local security group, and by default it has no members.

Note: If you implement IP Address Management (IPAM), you can also delegate management of DNS within IPAM.

Configure DNS logging

By default, Domain Name System (DNS) maintains a *DNS server log*, which you can examine in **Event Viewer** or Windows Admin Center. This event log is located in the **Applications and Services Logs** folder. It records common events, such as:

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

For more verbose logging, you can enable *debug logging*. Debug logging captures information about individual DNS requests and responses and writes them to a text file. This provides detailed information that you can use for troubleshooting. When you enable debug logging, you specify a file name and path for the log file, and a maximum size. Debug logging is disabled by default but you can enable it in the properties of the DNS server. You can filter the captured data based on packet direction, transport protocol, packet contents, packet type, and IP address.

As an alternative to storing debug logging in a text file, you can enable analytical logging for DNS that captures information about DNS packets and saves the information to an event log. Analytical logging for DNS does not provide any mechanism to filter which packets are captured, but you can filter or search events in **Event Viewer** or Windows Admin Center. The analytical log for DNS is located at **Applications and Services Logs\Microsoft\Windows\DNS-Server**. Like the debug log, this log is not enabled by default.

Both debug logging and analytical logging can be resource intensive. They can affect overall server performance and consume disk space. Therefore, you should enable them only temporarily when you require detailed information about DNS performance.

Aging and scavenging

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

Aging is determined by using these two parameters:

- The *no-refresh interval* is a period during which the client does not update the DNS record if there are no changes. If the client retains the same IP address, the record is not updated. This prevents a large number of time stamp updates on DNS records from triggering DNS replication. By default, the no-refresh interval is seven days.
- The *refresh interval* is the time span after the no-refresh interval when the client can refresh the record. If the DNS record isn't refreshed during this time span, it becomes eligible for scavenging. If the client refreshes the DNS record, then the no-refresh interval begins again for that record. The default length of the refresh interval is seven days. A client attempts to refresh its DNS record at startup, and every 24 hours while the system is running.

To perform aging and scavenging, you need to enable aging on the zone containing the resource records and enable scavenging on a DNS server. Only one DNS server hosting the zone needs to have scavenging enabled. The DNS server with scavenging enabled is the DNS server that will scan the zone for stale resource records and remove them, if necessary.

Note: Records that are added dynamically to the database are time stamped. Static records that you enter manually have a time-stamp value of 0. Therefore, aging will not affect the records, and they won't be scavenged out of the database.

Back up the DNS database

Backing up the contents of a DNS zone is useful to restore resource records, or even an entire zone, that is accidentally deleted. Your standard backup processes for Active Directory Domain Services (AD DS) and DNS servers might already contain this information, but you can manually back up zone contents. How you back up the DNS database depends on how DNS was implemented into your organization.

You can back up a primary zone that is not stored in AD DS by copying or backing up the individual zone file **zonename.dns**, which is in the **%windir%\System32\DNS** directory. For example, if your DNS primary zone is named **Adatum.com**, the DNS zone file will be named **Adatum.com.dns** by default.

To back up an Active Directory-integrated zone, you can use **dnscmd.exe** or the **Export-DnsServerZone** cmdlet. Both options extract the zone information from AD DS and store it in a file.

To back up a DNS zone by using Windows PowerShell, run the following command:

```
Export-DnsServerZone -Name <ZoneName> -Filename <ZoneBackupFilename>
```

To back up a DNS zone by using **dnscmd.exe**, run the following command:

```
dnscmd.exe /ZoneExport <ZoneName> <ZoneBackupFilename>
```

Create records in DNS

You must create DNS resource records before they can be resolved within the DNS infrastructure. When you create a DNS resource record, it exists within a DNS zone. A DNS zone constitutes several related records. You can manually create DNS records in a zone, but most host and pointer resource records for Windows servers and clients are created dynamically.

Manual creation

When you create resource records to support a specific service or app, you can manually create the resource records. For example, you can create host or CNAME records, such as `app.contoso.com`, for a specific app running on a server. The record name might be easier for users to remember, and users don't need to reference the server name.

You can create resource records by using DNS manager, Windows Admin Center, or Windows PowerShell. The following table lists some Windows PowerShell cmdlets that you can use to create DNS resource records.

Table 1: Windows PowerShell cmdlets to create DNS resource records

Cmdlet	Description
Add-DnsServerResourceRecord	Creates any resource record, specified by type
Add-DnsServerResourceRecordA	Creates a host (A) resource record
Add-DnsServerResourceRecordAAAA	Creates a host (AAAA) resource record
Add-DnsServerResourceRecordCNAME	Creates a CNAME alias resource record
Add-DnsServerResourceRecordMX	Creates an MX resource record
Add-DnsServerResourceRecordPtr	Creates a PTR resource record

Note: By default, the static resource records that you create manually don't have a time-stamp value configured, so aging and scavenging don't remove them.

Dynamic creation

When you allow dynamic updates for a DNS zone, clients that use DNS register with the DNS server. The dynamic update creates host and pointer records for the client. Dynamic DNS makes it easier for you to manage DNS, because when dynamic DNS is enabled, the current IP address for a computer registers automatically after an IP address change.

Note: The Dynamic Host Configuration Protocol (DHCP) client service performs the registration, regardless of whether the client's IP address is obtained from a DHCP server or is static.

Dynamic DNS registration is triggered by the following events:

- When the client starts, and the DHCP client service starts
- Every 24 hours while the DHCP client service is running
- When an IP address is configured, added, or changed on any network connection
- When an administrator executes the **Register-DNSClient** cmdlet
- When an administrator runs the **ipconfig /registerdns** command

Dynamic DNS updates can only be performed when the client communicates with a DNS server holding the primary zone. The client queries DNS to obtain the SOA record for the domain that lists the primary server. If the zone is Active Directory-integrated, the DNS server includes itself in the SOA as the primary server. Also, if you configure the zone for secure dynamic updates, the client authenticates to send the update.

By default, Windows DNS clients perform dynamic DNS registration. However, some non-Windows DNS clients do not support using dynamic DNS. For these clients, you can configure a Windows-based DHCP server to perform dynamic updates on behalf of clients.

Note: Secure dynamic updates create DNS records based on the primary DNS suffix configured on clients. This should be the same as the Active Directory Domain Services (AD DS) domain name to which the client is joined, but sometimes it can become misconfigured.

Configure DNS zones

You can create and configure Domain Name System (DNS) zones by using Windows Admin Center, DNS Manager, and Windows PowerShell. At time of writing, Windows Admin Center provided access to the most configured options such as dynamic updates and zone storage. Other features such as aging and scavenging are available only in DNS Manager and Windows PowerShell.

Zone storage and replication scope

When you create a primary zone, you have the option to create a zone file or store the zone in Active Directory Domain Services (AD DS). If you create a zone file, then the zone is a standard primary zone. If you store the zone in AD DS, then the zone is Active Directory-integrated. A secondary zone is always stored in a zone file.




Figure 1: Replication for Active Directory-integrated zones and traditional DNS zones

If you configure a zone to be Active Directory-integrated, then the zone data is stored in AD DS and replicated to domain controllers, as depicted in Figure 1. You can choose from the following options:

- **To all DNS servers running on domain controllers in this forest.** This option is useful in a multi-domain forest when you want the zone available to DNS servers in all domains. When you select this option, the DNS zone is stored in the **ForestDnsZones** partition.
- **To all DNS servers running on domain controllers in this domain.** This option is selected by default and works well for single-domain environments. When you select this option, the DNS zone is stored in the **DomainDnsZones** partition.
- **To all domain controllers in this domain (for Windows 2000 compatibility).** This option is seldom selected because replicating to all domain controllers in the domain is less efficient than replicating only to DNS servers running on domain controllers in the domain. When you select this option, the DNS zone is stored in the domain partition.
- **To all domain controllers in the scope of this directory partition.** You can use this option to select an application partition that you have created to store the zone.

Zone transfers

Zone records are synchronized from a primary zone to a secondary zone by performing a *zone transfer*. For each zone, you can control which servers hosting secondary zones can request a zone transfer. If you choose to allow zone transfers, you can control them with the following options:

- **To any server.** This option allows any server to request a zone transfer. You should avoid it because of security concerns.
- **Only to servers listed on the Name Servers tab.** This option is useful if you are already adding the DNS servers hosting secondary zones as name servers for the zone.
- **Only to the following servers.** This option allows you to specify a list of servers that are allowed to request zone transfers.

You can also configure notifications for zone transfers. When notifications are enabled, the primary server notifies the secondary server when changes are available to synchronize. This allows for faster synchronization.

Note: After initial replication of a secondary zone is complete, incremental zone transfers are performed.

Security for dynamic updates

You can configure each DNS zone with security settings for dynamic updates. For Active Directory-integrated zones, you can restrict the zone to allow only secure dynamic updates. Secure dynamic updates ensure that only the client that owns a DNS record can update it. If a second device with the same name attempts a secure dynamic update, the record won't update. When a zone is not Active Directory-integrated, you can still allow dynamic updates, but you can't enforce security.

Only domain joined Windows clients can perform secure dynamic updates. To allow non-domain joined Windows clients or non-Windows devices to perform dynamic updates, you need to allow nonsecure dynamic updates.

Windows PowerShell cmdlets

There are many Windows PowerShell cmdlets that you can use to manage DNS zones. The following table lists cmdlets that you can use.

Table 1: Windows PowerShell cmdlets to manage DNZ zones

Cmdlet	Description
Add-DnsServerPrimaryZone	Create a primary DNS zone
Add-DnsServerSecondaryZone	Create a secondary DNS zone
Get-DnsServerZone	View configuration information for a DNS zone
Get-DnsServerZoneAging	View aging configuration for a DNS zone
Remove-DnsServerZone	Removes a DNS zone
Restore-DnsServerPrimaryZone	Reloads the zone content from AD DS or a zone file
Set-DnsServerPrimaryZone	Modifies the settings of a primary DNS zone
Start-DnsServerZoneTransfer	Triggers a zone transfer to a secondary DNS zone

DNS forwarding

When a Domain Name System (DNS) server does not host a primary or secondary zone containing resource records in a DNS request, it needs a mechanism to find the required information. By default, each DNS server is configured with root hints that can be used to resolve DNS requests on the internet by finding the authoritative DNS servers. This works if the DNS server has access to the internet and the resource record being requested is available on the internet. Sometimes, both of these conditions aren't met. For example, it's common that internal DNS is hosted on domain controllers that can't access the internet and therefore can't use root hints. To optimize the name resolution process, you can use forwarding and stub zones.

Forwarders

You can configure each DNS server with one or more forwarders. If a DNS server receives a request for a zone for which it is not authoritative, and is not already cached by the server, the DNS server forwards that request to a forwarder. A DNS server uses a forwarder for all unknown zones.

Forwarders commonly are used for internet name resolution. The internal DNS servers forward requests to resolve internet names to a DNS server that is outside the corporate network. Your organization might configure the external DNS servers in a perimeter network, or use a DNS server provided by your internet service provider. This configuration limits external connectivity and increases security.

Conditional forwarding

You can configure conditional forwarding for individual DNS domains. This is similar to configuring a forwarder, except that it applies only to a single DNS domain. Trusted Active Directory Domain Services (AD DS) forests and partner organizations often use this feature.

When you create a conditional forwarder, you can choose whether to store it locally on a single DNS server or in AD DS. If you store it in AD DS, it can be replicated to all DNS servers running on domain controllers in the domain or forest, depending on the option you select. Storing conditional forwarders in AD DS makes it easier to manage them across multiple DNS servers.

Comparing stub zones and conditional forwarders

You can use stub zones or conditional forwarders to resolve DNS requests for zones for which the local DNS server is not authoritative. The difference is how the remote servers are selected for querying. You configure a conditional forwarder with specific remote DNS servers that are authoritative for the domain, whereas a stub zone replicates and uses all of the name server records configured in the zone.

If firewalls limit network connectivity between your DNS servers and the DNS servers that are authoritative for the zone, you might prefer to use a conditional forwarder. That way, you can configure firewall rules to allow communication only with the servers listed in the conditional forwarder.

If the authoritative DNS servers are likely to change over time, you might want to use a stub zone. A stub zone will automatically update the name server records and use only valid name servers for the zone. However, if firewalls control communication, the updated name servers might be not reachable.

DNS integration in AD DS

Active Directory Domain Services (AD DS) is highly dependent on Domain Name System (DNS). DNS is required to store the SRV records that domain joined clients use to locate domain controllers. In addition,

DNS servers can store zone data in AD DS when the DNS role is installed on domain controllers. It's common for domain controllers to be configured as DNS servers.

SRV records

When you add a domain controller to a domain, the domain controller advertises its services by creating SRV records (also known as *locator records*) in DNS. Unlike host (A) resource records, which map host names to IP addresses, SRV records map services to host names. For example, to publish its ability to provide authentication and directory access, a domain controller registers Kerberos v5 protocol and Lightweight Directory Access Protocol (LDAP) SRV records. These SRV records are added to several folders within the forest's DNS zones.

A typical SRV record contains the following information:

- The service name and port. This portion of the SRV record indicates a service with a fixed port. It does not have to be a well-known port. SRV records in Windows Server include LDAP (port **389**), Kerberos (port **88**), Kerberos password protocol (KPASSWD, port **464**), and global catalog services (port **3268**).
- Protocol. The TCP or User Datagram Protocol (UDP) is indicated as a transport protocol for the service. The same service can use both protocols in separate SRV records. Kerberos records, for example, are registered for both TCP and UDP. Microsoft clients use only TCP, but UNIX clients can use both UDP and TCP.
- Host name. The host name corresponds to the host (A) record for the server hosting the service. When a client queries for a service, the DNS server returns the SRV record and associated host (A) records, so the client does not need to submit a separate query to resolve the IP address of a service.

The service name in an SRV record follows the standard DNS hierarchy with components separated by dots. For example, a domain controller's Kerberos service is registered as **kerberos._tcp.sitename._sites.domainName**, where:

- **kerberos** is a Kerberos Key Distribution Center (KDC) that provides authentication.
- **_tcp** is any TCP-based services in the site.
- **sitename** is the site of the domain controller registering the service.
- **_sites** is all sites registered with DNS.
- **domainName** is the domain or zone, for example, `contoso.com`.

SRV records are dynamically registered by the NetLogon service running on a domain controller. If you need to force a domain controller to recreate its SRV records, you can restart the NetLogon service or restart the domain controller.

Note: Domain controllers perform dynamic DNS updates for A records and PTR records using the same mechanism as member servers or Windows clients.

Active Directory-integrated zones

A single DNS server hosting a primary zone is a single point of failure. If it goes down, because the secondary zone servers are read-only, they can resolve names, but cannot store additional records or accept changes to records.

You can make a DNS zone fault tolerant by storing it in AD DS. When the DNS server stores zone data in this way, the records in the zone are stored as AD DS objects, and the various properties of these objects are AD DS attributes. When a domain controller with an Active Directory-integrated DNS zone fails, the

DNS functionality for that zone, and the domain, continue to operate correctly, as long as there are other domain controllers configured as DNS servers with the Active Directory–integrated zone.

The benefits of an Active Directory–integrated zone are significant, and include:

- **Multi-master updates.** Unlike standard primary zones, which can only be modified by a single primary server, any writable domain controller to which the zone is replicated can write to Active Directory–integrated zones. This builds redundancy into the DNS infrastructure. Multi-master updates are particularly important in organizations that use dynamic update zones and have locations that are distributed geographically, because clients can update their DNS records without having to connect to a potentially geographically distant primary server.
- **Replication of DNS zone data by using AD DS replication.** One characteristic of Active Directory replication is attribute-level replication, in which only changed attributes are replicated. An Active Directory–integrated zone can thus avoid replicating the entire zone file, which is the case in traditional DNS zone transfer models.
- **Secure dynamic updates.** An Active Directory–integrated zone can enforce secure dynamic updates.
- **Detailed security.** As with other Active Directory objects, an Active Directory–integrated zone enables you to delegate administration of zones, domains, and resource records by modifying the access control list (ACL) on the zone.

Overview of DNS policies

You can use DNS policies to manipulate how a DNS server manages queries based on different factors. As an example, you might create a DNS policy to respond to queries asking for the IP address of a web server to respond with a different IP address based on the closest datacenter to the client. This differs from netmask reordering because the client will not have the same local subnet address as the web server, but the particular web server is closer than others, from the perspective of the client.

Scenarios for using DNS policies

You can create several DNS policies depending on your needs. There are various factors that might benefit from creating a DNS policy, based on the following scenarios:

- **Application high availability.** Clients are redirected to the healthiest endpoint for an application, where “healthiest” is determined by high availability factors in a failover cluster.
- **Traffic management.** Clients are redirected to the closest datacenter or server location.
- **Split-brain DNS.** Clients receive a response based on whether they are internal or external, and the DNS records are split into different zone scopes.
- **Filtering.** DNS queries are blocked if they are from a list of malicious IP addresses or fully qualified domain names (FQDNs).
- **Forensics.** Malicious DNS clients are redirected to a sinkhole instead of the computer they are trying to reach.
- **Time-of-day based redirection.** Clients are redirected to datacenters based on the time of the day.

DNS policy objects

To use the previously mentioned scenarios to create policies, you must identify groups of records in a zone, groups of clients on a network, or other elements. You can identify the elements by the following new DNS policy objects:

- **Client subnet.** This represents the IPv4 or IPv6 subnet from which queries are sent to a DNS server. You create subnets to later define policies that you apply based on the subnet that generates the requests. For example, you might have a split-brain DNS scenario where the name resolution request for `www.contoso.com` can be answered with an internal IP address to internal clients, and a different IP address to external clients.
- **Recursion scope.** This represents unique instances of a group of settings that control DNS server recursion. A recursion scope holds a list of forwarders and specifies whether recursion is used. A DNS server can have multiple recursion scopes. You can use DNS server recursion policies to choose a recursion scope for a given set of queries. If the DNS server is not authoritative for certain queries, DNS server recursion policies let you control how to resolve those queries. In this case, you can specify which forwarders to use and whether to use recursion.
- **Zone scopes.** DNS zones can have multiple zone scopes, and each zone scope can contain its own set of DNS resource records. The same resource record can be present across multiple scopes, with different IP addresses depending on the scope. Additionally, zone transfers can occur at the zone-scope level. This will allow resource records from a zone scope in a primary zone to be transferred to the same zone scope in a secondary zone.

Create and manage DNS policies

You create DNS policies based on level and type. You can use query-resolution policies to define how to manage client name resolution queries, and use zone-transfer policies to define zone transfers. You can apply both policy types at the server or zone level.

You can create multiple query resolution policies at the same level if they have a different value for the processing order. Recursion policies are a special type of server-level policy. They control how a DNS server performs query recursion, if at all. Recursion policies only apply when query processing reaches the recursion path. You can choose a value of DENY or IGNORE for recursion for a given set of queries. Otherwise, you can choose a set of forwarders for a set of queries.

The high-level steps to resolve a host record differently for users from a specific IP address range are:

1. Create a DNS server client subnet for the IP address range.
2. Create a DNS server zone scope for the zone containing the host record.
3. Add a host record to the zone that is specific to the zone scope.
4. Add a DNS server query resolution policy that allows the DNS server client subnet to query the zone scope for the zone.

The following is an example of the steps used to configure the DNS policy by using Windows PowerShell:

```
Add-DnsServerClientSubnet -Name "AsiaSubnet" -IPv4Subnet "172.21.33.0/24"
Add-DnsServerZoneScope -ZoneName "Contoso.com" -Name "AsiaZoneScope"
Add-DnsServerResourceRecord -ZoneName "Contoso.com" -A -Name "www" -IPv4Address "172.21.21.21" -ZoneScope "AsiaZoneScope"
Add-DnsServerQueryResolutionPolicy -Name "AsiaPolicy" -Action ALLOW -ClientSubnet "eq,AsiaSubnet" -ZoneScope "AsiaZoneScope,1" -ZoneName "Contoso.com"
```

com"

Additional reading: For detailed information about implementing DNS policies, refer to **DNS Policy Scenario Guide**³.

Overview of DNSSEC

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

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

The high-level steps for deploying DNSSEC are:

1. Sign the DNS zone.
2. Configure the trust anchor distribution.
3. Configure the name resolution policy table (NRPT) on client computers.

Resource records

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

The following table describes the additional resource records used with DNSSEC.

Table 1: Additional resource records used with DNSSEC

Resource record	Purpose
RRSIG (Resource Record Signature)	This record holds a signature for a set of DNS records. DNS clients can use it to check the authority of a response. When a resource record is resolved, an RRSIG record is sent for verification.

³ <https://aka.ms/dns-policy-scenario-guide>

Resource record	Purpose
DNSKEY	This record publishes the public keys for the zone. It allows clients to validate signatures created by the private key held by the DNS server. These keys require periodic replacement through key rollovers. Windows Server supports automated key rollovers. Every zone has multiple DNSKEY records that are broken down to the zone signing key (ZSK) and key-signing key (KSK) level. The ZSK is used to create the RRSIG records for a set of resource records. The KSK is used to create an RRSIG record for the ZSK DNSKEY record.
NSEC (Next Secure)	When the DNS response has no data to provide to the client, this record authenticates that the host does not exist.
NSEC3	This record is a hashed version of the NSEC record, which prevents attacks by enumerating the zone.
DS (Delegation Signer)	This record is a delegation record that contains the hash of the public key of a child zone. This record is signed by the parent zone's private key. If a child zone of a signed parent is also signed, you must manually add the DS records from the child to the parent to create a <i>chain of trust</i> .

Signing the DNS zone

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

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

Trust anchors distribution

A *trust anchor* is an authoritative entity that is represented by a public key. The **TrustAnchors** zone stores preconfigured public keys that are associated with a specific zone. In DNS, the trust anchor is the DNSKEY or DS resource record. Client computers use these records to build trust chains. You must configure a trust anchor from the zone on every domain DNS server to validate responses from that signed zone. All DNS servers hosting a zone need to have the same DNSKEY key records to provide the information required to validate the RRSIG records.

By default, trust anchors are local on the DNS server on which you configured DNS signing. However, you can enable the distribution of trust anchors for a zone. If the DNS server is a domain controller, the trust anchors are distributed to all DNS servers running on domain controllers in the forest. If the DNS server is not a domain controller, then trust anchors are stored in a local trust anchor store at `%windir%\system32\dns\TrustAnchors.dns`.

Name Resolution Policy Table

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

Test your knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

Which type of resource record is used only for IPv6 addresses?

- PTR record
- TXT record
- AAAA record
- CNAME record

Question 2

Which DNS functionality should you use to direct DNS queries for a single domain to a partner organization through firewalls?

- Forwarder
- Stub zone
- Root hints
- Conditional forwarder

Deploying and managing IPAM

Lesson overview

IP Address Management (IPAM) is a service that you can use to simplify the management of IP ranges, Dynamic Host Configuration Protocol (DHCP), and DNS. When you implement IPAM, you have access to all this information in a single console instead of connecting to each server and consolidating information manually.

To implement IPAM, you need to deploy an IPAM server and one or more IPAM clients. After you deploy the IPAM server, you can use Group Policy Objects (GPOs) to configure the managed servers to avoid manually configuring security groups and firewall rules on each managed server.

Lesson objectives

After completing this lesson, you'll be able to:

- Describe IPAM.
- List the requirements for deploying IPAM.
- Describe implementing IPAM.
- Deploy the IPAM role.
- Describe how to administer IPAM.
- Configure IPAM options.
- Manage DNS zones with IPAM.
- Configure DHCP servers with IPAM.
- Use IPAM to manage IP addressing.

What is IPAM?

Managing the allocation of IP addresses can be a complex task in large networks. IP Address Management (IPAM) provides a framework for discovering, auditing, and managing the IP address space of your network. It enables you to monitor and administer both Dynamic Host Configuration Protocol (DHCP) and Domain Name System (DNS) services, and it provides a comprehensive display of where specific IP addresses are allocated.

You can configure IPAM to collect statistics from domain controllers and Network Policy Servers (NPSs). The Windows Internal Database (WID), or optionally, a Microsoft SQL Server database, stores the collected data.

The benefits of using IPAM include:

- IPv4 and IPv6 address space planning and allocation.
- IP address space utilization statistics and trend monitoring.
- Static IP inventory management, lifetime management, and DHCP and DNS record creation and deletion.
- Service and zone monitoring of DNS servers.
- IP address lease and sign-in event tracking.

IPAM consists of the following four modules:

- IPAM discovery. You can configure IPAM to use Active Directory Domain Services (AD DS) for discovering servers that are running Windows Server 2008 and newer, and servers that are domain controllers or that have DHCP or DNS installed. Also, you can add servers manually.
- IP address space management. You can use this module to examine, monitor, and manage the IP address space. You can dynamically issue or statically assign addresses. You can also track address utilization and detect overlapping DHCP scopes.
- Multiserver management and monitoring. You can use this module to manage and monitor multiple DHCP servers. Use multiserver management when you need tasks to run across multiple servers. For example, you can configure and edit DHCP properties and scopes, and you can track the status of DHCP and scope utilization. You can also monitor multiple DNS servers and monitor the health and status of DNS zones across authoritative DNS servers.
- Operational auditing and IP address tracking. You can use the auditing tools to track potential configuration problems. You can collect, manage, and examine details of configuration changes from managed DHCP servers. You can also collect address lease tracking from DHCP lease logs, and sign-in event information from NPSs and domain controllers.

IPAM deployment requirements

Before you deploy IP Address Management (IPAM), you should understand the IPAM architecture so that you can plan a deployment that meets your needs. You also need to understand the deployment requirements and limitations.

An IPAM deployment includes two components:

- IPAM server. The IPAM server performs data collection from the managed servers. Additionally, the IPAM server manages the Windows Internal Database (WID) or a SQL Server database, and it provides role-based access control (RBAC).
- IPAM client. The IPAM client provides the client computer interface and interacts with the IPAM server, invoking Windows PowerShell cmdlets to perform remote management, Dynamic Host Configuration Protocol (DHCP) configuration, and DNS monitoring. The IPAM client can be a Windows client operating system or Windows Server operating system.

IPAM topology

IPAM servers don't coordinate with each other or roll up information from one IPAM server to another. Consequently, the two main topology options are centralized or distributed.

For a centralized topology, you deploy a single IPAM server for your entire forest. A single IPAM server provides centralized control and visibility for IP addressing tasks. You can examine your entire IP addressing infrastructure from a single console when you are using the centralized topology. You can use a single IPAM server for multiple Active Directory Domain Services (AD DS) forests with a two-way trust in place.

For a distributed topology, you deploy an IPAM server to each site in your forest. It's common to use the distributed topology when your organization has multiple sites with significant IP addressing infrastructure in place. Servers in each location can help to distribute a workload that might be too large for a single server to manage. You can also use the distributed topology to enable separate locations or business units to administer their own IP addressing management.

You can also implement a hybrid topology with a centralized IPAM server and an IPAM server at each site. Both the local IPAM server and the centralized IPAM server monitor managed servers. You can also control the management scope by monitoring some services centrally and monitoring others at each site.

IPAM server requirements

The IPAM server must be a member server in the domain because installing the IPAM server on a domain controller is not supported. As a best practice, the IPAM server should be a single-purpose server. You shouldn't install other network roles such as DHCP or DNS on the same server. If you install the IPAM server on a DHCP server, IPAM won't be able to detect other DHCP servers on the network.

IPAM collects data and stores it in a database. You can use WID on the IPAM server or a Microsoft SQL Server database. If you use a SQL Server database for IPAM, you have the option to use a database on a separate server. However, if you use SQL Server to host your IPAM database, that must be the only SQL Server instance running on that server.

IPAM collects a large amount of data from various servers. Ensure that the disk hosting the SQL database is large enough to store the data collected. For example, IP address utilization data for 10,000 clients requires approximately 1GB of disk space per month.

IPAM deployment considerations

Consider the following when implementing IPAM:

- To manage the IPv6 address space, IPv6 must not be disabled on the IPAM server.
- Sign in to the IPAM server with a domain account and not a local account.
- For IPAM's IP address tracking and auditing feature to work, you must enable logging of account sign-in events on domain controllers and Network Policy Servers (NPSs).
- You can define the scope of discovery to a subset of domains in the forest.
- A single IPAM server can support up to 150 DHCP servers, 6,000 DHCP scopes, 500 DNS servers, and 150 DNS zones.
- IP address utilization trends are provided only for IPv4.
- IP address reclamation support is provided only for IPv4.
- IPAM does not check for IP address consistency with routers and switches.

Process for deploying IPAM

Before you deploy IP Address Management (IPAM), you need to go through a planning process where you determine how you will use IPAM and the deployment topology that supports your organizational needs. You should only start deploying IPAM after you complete the planning process. To deploy IPAM, you need to deploy IPAM servers and IPAM clients.

Deploy IPAM servers

Deploying IPAM servers begins with the installation of the IPAM server feature. After determining which IPAM topology to use, you can deploy IPAM servers by performing the following steps:

1. Install the IPAM Server feature. You can install it by using Windows Admin Center, **Server Manager**, or by using the following Windows PowerShell command:

```
Install-WindowsFeature IPAM -IncludeManagementTools
```
2. Provision IPAM servers. After installing the IPAM server feature, you must provision each IPAM server to create the permissions, file shares, and settings on the managed servers. You can perform this

manually or by deploying Group Policy Objects (GPOs). Using GPOs is recommended because it automates the configuration process for managed servers.

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 will run discovery. You can also manually add a server in the **IPAM management console** by specifying the fully qualified domain name (FQDN) of the server that you want to manage.
4. Choose and manage the discovered servers. After discovery completes, and after you manually add any servers that were not discovered, choose the servers that you want to manage by editing the server properties in the IPAM console and changing the Manageability Status to **Managed**. After setting the management permission for a server, note the status indicator displaying **IPAM Access Unblocked** in the IPAM server inventory.

Deploy IPAM clients

You use the IPAM client to configure and manage IPAM servers. If you install the IPAM role on a Windows server with the Desktop Experience, then the IPAM client automatically is installed on the IPAM server. If you install the IPAM role on Server Core, then you need to manually install the IPAM client on another Windows Server used for management or a Windows client to manage IPAM remotely. IPAM installation varies based on the operating system:

- Windows Server. You can install the IPAM client by installing the Windows feature under **Remote Server Administration Tools\Feature Administration Tools\IP Address Management (IPAM) Client**.
- Windows 8.1 and Windows 10. The IPAM client installs automatically when you install Remote Server Administration Tools (RSAT).

Demonstration: Install the IPAM role

This demonstration describes how to install the IP Address Management (IPAM) role.

Demonstration steps

Install the IPAM Server feature

1. On **SEA-ADM1**, open **Microsoft Edge** and sign in to **Windows Admin Center** as **Contoso\Administrator** with the password **Pa55w.rd**.
2. In **Windows Admin Center**, connect to **SEA-SVR2**.
3. Use **Roles & features** to install the **IP Address Management (IPAM) Server** feature on **SEA-SVR2**.

Install the IPAM Client feature

1. In **Windows Admin Center**, connect to **SEA-ADM1**.
2. Use **Roles & features** to install the **IP Address Management (IPAM) Client** feature on **SEA-ADM1**.

Provision the IPAM server

1. Open **Server Manager** and examine the IPAM workspace.

2. Connect to the IPAM server `SEA-SVR2.Contoso.com`.
3. Provision the IPAM server by using the following information:
 - Database: **Windows Internal Database (WID)**
 - Provisioning method: **Group Policy based**
 - GPO name prefix: **IPAM**

Create the IPAM GPOs

1. To create the required GPOs, run the following command at a Windows PowerShell prompt, and then select Enter:

```
Invoke-IpamGpoProvisioning -Domain contoso.com -GpoPrefixName IPAM -IpamServerFqdn sea-svr2.contoso.com
```
2. Verify that three new GPOs are created, and that they include IPAM as a naming prefix.

Add the server to IPAM and review IPAM data

1. In **Server Manager**, to configure server discovery, get the forests and add the domain.
2. Start server discovery.
3. Add **SEA-DC1** as a managed server for **DC**, **DNS**, and **DHCP**.
4. In Windows Admin Center, connect to **SEA-DC1**.
5. Use the PowerShell tool to sign in with the password **Pa55.wrd**, and then run **gpupdate**.
6. In **Server Manager**, refresh the **server access status** for **SEA-DC1**.
7. On the **IPAM Overview** page, retrieve the data from managed servers.
8. Review the collected data.

Administer IPAM

Configuring administration for IP Address Management (IPAM) can be a complex task depending on how your IPAM infrastructure is deployed and who is managing the infrastructure. You can allow an administrator to manage all aspects within IPAM or limit management ability. If you assign specific administrative tasks to administrators, you can limit tasks based on IPAM functional areas or specific servers.

To define and establish fine-grained control for users and groups, you can use role-based access control (RBAC) to customize roles, access scopes, and access policies. This enables users and groups to perform a specific set of administrative operations on specific objects that IPAM manages. You implement role-based management in IPAM by using:

- **Roles.** A *role* is a collection of IPAM operations. You can associate a role with a user or group in Windows by using an access policy. Eight built-in administrator roles are available for convenience, but you can also create customized roles to meet your business requirements. You can create and edit roles from the **Access Control** node in the **IPAM management** console.
- **Access scopes.** An *access scope* determines the objects to which a user has access. You can use access scopes to define administrative domains in IPAM. For example, you might create access scopes based on a user's geographical location. By default, IPAM includes an access scope named **Global**. All other

access scopes are subsets of the **Global** access scope. Users or groups that you assign to the **Global** access scope have access to all objects in IPAM that their assigned role permits. You can create and edit access scopes from the **Access Control** node in the **IPAM management** console.

- Access policies. An *access policy* combines a role with an access scope to assign permissions to a user or group. For example, you might define an access policy for a user with a role named **IP Block Admin**, and an access scope named **Global\Asia**. This user would have permission to edit and delete IP address blocks that are associated with the **Asia** access scope, but would not have permission to edit or delete any other IP address blocks in IPAM. You can create and edit access policies from the **Access Control** node in the **IPAM management** console.

IPAM security groups

IPAM has several built-in role-based security groups that you can use for managing your IPAM infrastructure, as listed in the following table.

Table 1: IPAM security groups

Group name	Description
IPAM Administrators	Members of this group have privileges to access all IPAM data and to perform all IPAM tasks.
IPAM MSM Administrators	Members of this group can manage DHCP servers, scopes, policies, and DNS servers and associated zones and records.
IPAM DNS Administrators	Members of this group can manage DNS servers and their associated DNS zones and resource records.
DNS Record Administrators	Members of this group can manage DNS resource records.
IPAM ASM Administrators	Members of this group can perform IP address space tasks, in addition to common IPAM management tasks.
IP Address Record Administrators	Members of this group can manage IP addresses, including unallocated addresses, and members can create and delete IP address instances.
IPAM DHCP Administrators	Members of this group can manage DHCP servers and their scopes.
IPAM DHCP Scope Administrators	Members of this group can manage DHCP scopes.
IPAM DHCP Reservations Administrators	Members of this group can manage DHCP reservations.

Configure IPAM options

You can configure IP Address Management (IPAM) to suit your environment and provide the level of manageability that you require. You have the option to configure all the managed servers manually by creating and configuring the necessary firewall rules, security groups, and file shares. However, to simplify management, you should use Group Policy Objects (GPOs) to perform provisioning.

When you select Group Policy provisioning, you are prompted to provide a prefix for the GPOs so that they can be easily identified. The names of the GPOs display after you provide the prefix. You need to

create these GPOs after completing the provisioning wizard. The wizard configures IPAM to use the GPOs but does not create them.

You need to create the following GPOs:

- **<Prefix>_DHCP.** This GPO applies settings that allow IPAM to monitor, manage, and collect information from managed DHCP servers on the network. It sets up IPAM provisioning scheduled tasks and adds Windows Defender Firewall inbound rules for Remote Event Log Management (RPC-EMAP and RPC), Remote Service Management (RPC-EMAP and RPC), and DHCP Server (RPCSS-In and RPC-In).
- **<Prefix>_DNS.** This GPO applies settings that allow IPAM to monitor and collect information from managed DNS servers on the network. It sets up IPAM provisioning scheduled tasks and adds Windows Defender Firewall inbound rules for RPC (TCP, Incoming), RPC Endpoint Mapper (TCP, Incoming), Remote Event Log Management (RPC-EMAP and RPC), and Remote Service Management (RPC-EMAP and RPC).
- **<Prefix>_DC_NPS.** This GPO applies settings that allow IPAM to collect information from managed domain controllers and NPSs on the network for IP address tracking purposes. It sets up IPAM provisioning scheduled tasks and adds Windows Defender Firewall inbound rules for Remote Event Log Management (RPC-EMAP and RPC) and Remote Service Management (RPC-EMAP and RPC).

To create the GPOs required by IPAM, you need to use the **Invoke-IpamGpoProvisioning** cmdlet and include the domain in which to create the GPOs, and the prefix for the GPO names. If you don't run **Invoke-IpamGpoProvisioning** from the IPAM server, you need to include the IPAM server name. When you run the cmdlet without a server name, the computer account from the local computer is added as a member of the **IPAMUG** group in Active Directory Domain Services (AD DS). You will need to add the computer account of the IPAM server to this group.

The following example creates the IPAM GPOs with a prefix of **IPAM** in the `contoso.com` domain, and adds the IPAM server **SEA-SVR2** to the **IPAMUG** group. This group is granted permissions on each managed server.

```
Invoke-IpamGpoProvisioning -Domain contoso.com -GpoPrefixName IPAM -Ipam-ServerFqdn SEA-SVR2.contoso.com
```

The three GPOs are automatically linked to the root of the domain, but security filtering prevents them from applying to any servers. When you select a server for IPAM to manage, that server is added to the security filtering for the GPO, and then is given permission to apply the GPO. If the naming of the GPOs does not match what you specified in the provisioning wizard, this process fails.

Manage DNS zones with IPAM

You can use IP Address Management (IPAM) to manage DNS servers and zones for all servers that the IPAM server manages. During discovery, IPAM discovers all DNS servers in the domains that you specify. You can use IPAM to perform the following DNS management tasks:

- **Examine DNS servers and zones.** You can examine all managed DNS servers, in addition to the forward lookup zones and the reverse lookup zones on those DNS servers. Zone status and health is available for forward lookup zones, but not for reverse lookup zones.
- **Create new zones.** To create DNS zones, on the **navigation** pane, select the **DNS and DHCP Servers** node. Right-click or access the context menu for the DNS server to which you want to add a zone, and then select **Create DNS zone**.

- Create DNS records. You can create DNS records for any zone that IPAM manages. To do this, perform the following steps:
 1. On the IPAM **navigation** pane, select **DNS Zones**, and then select the appropriate zone, for example, `contoso.com`.
 2. Right-click or access the context menu for the zone, and then select **Add DNS resource record**.
 3. Verify that the correct DNS zone name and DNS server name display in the list, and then add a new DNS resource record. For example, select **Resource record type A**, and then add the required information: name, FQDN, and IP address.
- Manage conditional forwarders. To add a conditional forwarder, on the **navigation** pane, select the **DNS and DHCP Servers** node. Right-click or access the context menu for the DNS server to which you want to add a zone, and then select **Create DNS conditional forwarder**.
 - To manage a conditional forwarder after you create it, on the **navigation** pane, under DNS Zones, select **Conditional Forwarders**. You can then manage the conditional forwarding settings in the **details** pane.
- Open the DNS console for any server that IPAM manages. You can open the **Microsoft Management Console (MMC)** for DNS by right-clicking or accessing the context menu for a server on the DNS and DHCP servers page, and then selecting **Launch MMC**.

Configure DHCP servers with IPAM

You can configure Dynamic Host Configuration Protocol (DHCP) servers and DHCP scope information by using the IP Address Management (IPAM) administration interface. IPAM enables you to configure multiple DHCP servers and to use functionality such as DHCP Failover so that the servers work together in your DHCP implementation.

Configuring DHCP servers

You typically perform DHCP configuration for individual servers from the DNS and DHCP servers page. You can perform several configuration tasks on a DHCP server from within the IPAM administration console:

- Examine DHCP scope information across all servers.
- Edit DHCP server properties. You can edit server properties such as DHCP audit logging, DNS dynamic update configuration, and media access control (MAC) address filtering.
- Edit DHCP server options. You can configure and create DHCP server options based on vendor or user classes.
- Configure DHCP vendor or user classes. You can examine and modify user and vendor classes.
- Configure DHCP policy. You can edit DHCP policy properties and conditions.
- Import DHCP policy. You can import DHCP policies by using files that other DHCP servers export.
- Add DHCP MAC address filters. You can add DHCP MAC address filters to allow or deny DHCP address assignments based on MAC addresses.
- Activate and deactivate DHCP policies. You can control the implementation of DHCP policies.
- Replicate DHCP servers. This option replicates the configuration of failover scopes on a server to failover partner servers.

- Launch the **DHCP Management console**. You can open the **DHCP Management console** for the selected server.

Configuring DHCP scopes

You can configure DHCP scope details in IPAM by performing the following tasks:

- Edit DHCP scope properties.
- Duplicate a DHCP scope. Use a DHCP scope as a template for creating a new scope on the same server or another server.
- Create a DHCP reservation.
- Add to a DHCP superscope.
- Configure a DHCP Failover.
- Import a DHCP policy.
- Activate and deactivate DHCP scopes.
- Activate and deactivate DHCP policies for the selected scope.
- Replicate a DHCP scope.
- Remove a DHCP Failover configuration.
- Remove a scope from a DHCP superscope.

Use IPAM to manage IP addressing

You can use IP Address Management (IPAM) to manage, track, audit, and report your organization's IPv4 and IPv6 address spaces. The IPAM **IP ADDRESS SPACE** node provides you with IP address utilization statistics and historical trend data so that you can make informed planning decisions for dynamic, static, and virtual address spaces. IPAM automatically discovers address spaces and utilization data from the Dynamic Host Configuration Protocol (DHCP) servers that IPAM manages. You can also import IP address information from comma-separated value (CSV) files.

IPAM also enables you to detect overlapping IP address ranges that are defined on different DHCP servers to:

- Find free IP addresses within a range.
- Create DHCP reservations.
- Create DNS records.

View and manage IP addressing

The IPAM **Administration Console** provides a number of ways to filter the view of the IP address space. You can customize the available components of the IP address space by using any of the following views:

- IP address blocks
- IP address ranges
- IP addresses
- IP address inventory
- IP address range groups

IP address blocks

IP address blocks are the highest-level entities within an IP address space organization. An IP address block is an IP subnet marked by a start IP address and an end IP address. You can use IP address blocks to create and allocate IP address ranges to DHCP. You can add, import, edit, and delete IP address blocks. IPAM maps IP address ranges to the appropriate IP address block automatically based on the boundaries of the range.

IP address ranges

IP address ranges are the next hierarchical level of IP address space entities after IP address blocks. An IP address range is an IP subnet that is marked by a start IP address and an end IP address. IP address ranges typically correspond to a DHCP scope, a static IPv4 or IPv6 address range, or to an address pool that is used to assign addresses to hosts.

IP addresses

IP addresses are the addresses that make up the IP address range. IPAM enables end-to-end lifecycle management of IPv4 and IPv6 addresses, including record syncing with DHCP and DNS servers. IPAM maps an address to the appropriate range automatically based on the starting and ending address of the IP address range.

IP address inventory

In the **IP Address Inventory** view, there is a list of all IP addresses in the enterprise along with their device names and types. IP address inventory is a logical group within the **IP addresses** view. You can use this group to customize the way the address space displays for managing and tracking IP usage.

IP address range groups

Using IPAM, you can organize IP address ranges into logical groups called **IP address range 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.

Monitor DHCP and DNS servers

IPAM enables automated, periodic service monitoring of DHCP and DNS servers across a single forest or across multiple forests. In the **IPAM console**, monitoring and management of DHCP and DNS servers is organized into the views that the following table lists.

Table 1: IPAM console views used to monitor and manage DHCP and DNS servers

View	Description
DNS and DHCP servers	By default, managed DHCP and DNS servers are arranged by their network interface in /32 subnets for IPv4 and /128 subnets for IPv6. You can select the view so that it displays only DHCP scope properties, only DNS server properties, or both.

View	Description
DHCP scopes	This view enables scope utilization monitoring. Utilization statistics are automatically collected periodically from a managed DHCP server. You can track important scope properties such as Name, ID, Prefix Length, and Status.
DNS zone monitoring	You enable zone monitoring for forward lookup zones. Zone status is based on events that IPAM collects. The status of each zone is summarized.
Server groups	You can organize managed DHCP and DNS servers into logical groups. For example, you might organize servers by business unit or geography. You define groups by selecting the grouping criteria from the built-in fields or user-defined fields.

Test your knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

Which of the following are valid options for storing IP Address Management (IPAM) data? (Choose two.)

- Windows Internal Database
- JET database
- Access database
- Microsoft SQL Server database

Question 2

Which IPAM security groups can manage IP address blocks and IP address inventory? (Choose two.)

- IPAM DHCP Administrator
- IPAM ASM Administrator
- IPAM MSM Administrator
- IPAM Administrator

Module 03 lab and review

Lab: Implementing and configuring network infrastructure services in Windows Server

Scenario

Contoso, Ltd. is a large organization with complex requirements for network services. To help meet these requirements, you will deploy and configure DHCP so that it is highly available to ensure service availability. You will also set up DNS so that Trey Research, a department within Contoso, can have its own DNS server in the testing area. Finally, you will provide remote access to Windows Admin Center and secure it with Web Application Proxy.

Objectives

After completing this lab, you'll be able to:

- Deploy and configure DHCP
- Deploy and configure DNS

Estimated time: 30 minutes

Module review

Use the following questions to check what you've learned in this module.

Question 1

Which network infrastructure service in Windows Server allows you to monitor and manage IP address ranges for the entire organization?

- Domain Name System (DNS)
- NPS
- IP Address Management (IPAM)
- Remote access services

Question 2

Which of the following are true about DHCP Failover? (Select two.)

- IP address ranges must split 80:20 between servers.
- A failover relationship can have up to four partners.
- A failover relationship can have only two partners.
- Load balance mode configures one server as primary to service all requests.
- The necessary firewall rules are configured automatically when the DHCP role is installed.

Question 3

Which of the following options are required when configuring a DHCP reservation? (Select three.)

- MAC address
- Description
- IP address
- Reservation name
- Computer name

Question 4

Which type of DNS zone automatically replicates to all domain controllers in a domain that have the DNS role installed?

- Primary
- Secondary
- Stub
- Active Directory-integrated

Question 5

Which service running on domain controllers creates the SRV records used by clients to locate the domain controller?

- Netlogon
- DNS client
- Workstation
- DHCP Client

Question 6

Which feature of DNS can you use to resolve a host record to different IP addresses depending on user location?

- DNSSEC
- Stub zone
- Conditional forwarder
- DNS policies

Question 7

How do you create the Group Policy Objects (GPOs) used to configure a server that is managed by IPAM?

- Run the Install-WindowsFeature cmdlet
- Run the Invoke-IpamGpoProvisioning cmdlet
- Select Group Policy provisioning in the configuration wizard
- Run the New-GPO cmdlet

Answers

Question 1

If you configure a DHCP scope with a lease length of four days, when will computers attempt to renew the lease for the first time?

- 1 day
- 2 days
- 3 days
- 3.5 days

Explanation

Two (2) days is the correct answer. If you configure a DHCP scope with a lease length of four days, computers will attempt to renew the lease for the first time after two days.

Question 2

Which permissions are required to authorize a DHCP server in a multiple domain AD DS forest?

- Member of "Enterprise Admins" group
- Member of "Domain Admins" group
- Member of local "Administrators" group on the DHCP server
- Member of "DHCP Administrators"

Explanation

Member of "Enterprise Admins" group is the correct answer. In an Active Directory Domain Services (AD DS) forest with multiple domains, you need permissions in all domains to authorize DHCP servers in all the domains. The "Enterprise Admins" group has permissions to authorize DHCP servers in all the domains in an AD DS forest.

Question 1

Which type of resource record is used only for IPv6 addresses?

- PTR record
- TXT record
- AAAA record
- CNAME record

Explanation

The correct answer is AAAA record. An AAAA record is a host record that resolves a name to an IPv6 address. IPv4 uses an A record.

Question 2

Which DNS functionality should you use to direct DNS queries for a single domain to a partner organization through firewalls?

- Forwarder
- Stub zone
- Root hints
- Conditional forwarder

Explanation

The correct answer is conditional forwarder. Partner organizations commonly use a conditional forwarder because it defines settings for a single domain. Also, you can configure specific IP addresses for communication, which simplifies firewall configuration.

Question 1

Which of the following are valid options for storing IP Address Management (IPAM) data? (Choose two.)

- Windows Internal Database
- JET database
- Access database
- Microsoft SQL Server database

Explanation

Windows Internal Database and Microsoft SQL Server database are the correct answers.

Question 2

Which IPAM security groups can manage IP address blocks and IP address inventory? (Choose two.)

- IPAM DHCP Administrator
- IPAM ASM Administrator
- IPAM MSM Administrator
- IPAM Administrator

Explanation

IPAM ASM Administrator and IPAM Administrator are the correct answers.

Question 1

Which network infrastructure service in Windows Server allows you to monitor and manage IP address ranges for the entire organization?

- Domain Name System (DNS)
- NPS
- IP Address Management (IPAM)
- Remote access services

Explanation

IPAM is the correct answer. IPAM is used to centrally monitor and manage DNS, DHCP, and IP address ranges.

Question 2

Which of the following are true about DHCP Failover? (Select two.)

- IP address ranges must split 80:20 between servers.
- A failover relationship can have up to four partners.
- A failover relationship can have only two partners.
- Load balance mode configures one server as primary to service all requests.
- The necessary firewall rules are configured automatically when the DHCP role is installed.

Explanation

The correct answers are "A failover relationship can have only two partners" and "The necessary firewall rules are configured automatically when the DHCP role is installed."

Question 3

Which of the following options are required when configuring a DHCP reservation? (Select three.)

- MAC address
- Description
- IP address
- Reservation name
- Computer name

Explanation

The correct answers are MAC address, IP address, and reservation name.

Question 4

Which type of DNS zone automatically replicates to all domain controllers in a domain that have the DNS role installed?

- Primary
- Secondary
- Stub
- Active Directory-integrated

Explanation

Active Directory-integrated is the correct answer. Active Directory-integrated zones are stored in Active Directory Domain Services (AD DS) and are replicated to domain controllers that have the DNS role installed.

Question 5

Which service running on domain controllers creates the SRV records used by clients to locate the domain controller?

- Netlogon
- DNS client
- Workstation
- DHCP Client

Explanation

Netlogon is the correct answer. When the Netlogon service starts, it dynamically registers the SRV records in DNS.

Question 6

Which feature of DNS can you use to resolve a host record to different IP addresses depending on user location?

- DNSSEC
- Stub zone
- Conditional forwarder
- DNS policies

Explanation

DNS policies is the correct answer. When you create a DNS policy, you can specify conditions that control how a DNS server responds to a request. This includes alternate host records based on the client IP address.

Question 7

How do you create the Group Policy Objects (GPOs) used to configure a server that is managed by IPAM?

- Run the Install-WindowsFeature cmdlet
- Run the Invoke-IpamGpoProvisioning cmdlet
- Select Group Policy provisioning in the configuration wizard
- Run the New-GPO cmdlet

Explanation

Run the Invoke-IpamGpoProvisioning cmdlet is the correct answer. When you run this cmdlet, you specify the prefix to use for the GPO names. The GPOs are created and linked to the root of the domain.



Module 4 File servers and storage management in Windows Server

Volumes and file systems in Windows Server

Lesson overview

A *volume* is a usable area of space on one or more physical disks, formatted with a file system. To format a volume, you must select the file system that the volume should use. Windows Server supports different file systems, including file allocation table (FAT), FAT32, the NTFS file system, and the Resilient File System (ReFS). In Windows Server, you can choose to use several types of volumes to create high-performance storage, fault-tolerant storage, or a combination of both. This lesson explores how select a file system for a Windows Server volume and how to create and manage volumes in Windows Server 2019.

Lesson objectives

After completing this lesson, you'll be able to:

- Provide an overview of file systems in Windows Server.
- Describe the use of ReFS in Windows Server.
- Describe disk volumes.
- Describe how to manage volumes in Windows Server.
- Describe File Server Resource Manager.
- Describe how to manage permissions on volumes.

Overview of file systems in Windows Server

Before you can store data on a volume, you must first format the volume. To do so, you must select the file system that the volume should use. Several file systems are available, each with its own advantages and disadvantages.

Types of file systems

The different types of file systems include:

- File allocation table (FAT), FAT32, and extended file allocation table (exFAT)
- The NT File System (NTFS) file system
- Resilient File System (ReFS)

FAT

The FAT file system is the most simplistic of the file systems that the Windows operating system supports. The FAT file system is characterized by a table that resides at the top of the volume. To protect the volume, two copies of the FAT file system are maintained in case one becomes damaged. Additionally, the file allocation tables and the root directory must be stored in a fixed location, so that the system's boot files can be located.

A disk formatted with the FAT file system is allocated in clusters, and the size of the volume determines the size of the clusters. When you create a file, an entry is created in the directory, and the first cluster number containing data is established. This entry in the table either indicates that this is the last cluster of the file, or points to the next cluster. There is no organization to the FAT directory structure, and files are given the first open location on the drive.

Because of the size limitation with the file allocation table, the original release of FAT could only access partitions that were less than 2 gigabyte (GB) in size. To enable larger disks, Microsoft developed FAT32, which supports partitions of up to 2 terabytes (TB).

FAT doesn't provide any security for files on the partition. You should never use FAT or FAT32 as the file system for disks attached to Windows Server servers. However, you might consider using FAT or FAT32 to format external media such as USB flash media. Note, however, that FAT 32 in Windows 10 now supports encryption through the Encrypting File System (EFS).

The file system designed especially for flash drives is Extended FAT (exFAT). You can use it when FAT32 isn't suitable, such as when you need a disk format that works with a television (TV), which requires a disk that is larger than 2 TB. A number of media devices support exFAT, such as modern flat panel TVs, media centers, and portable media players.

NTFS

NTFS is the standard file system for all Windows operating systems. Unlike FAT, there are no special objects on the disk, and there is no dependence on the underlying hardware, such as 512-byte sectors. In addition, in NTFS there are no special locations on the disk, such as the tables.

NTFS is an improvement over FAT in several ways, including better support for metadata and the use of advanced data structures to improve performance, reliability, and disk space utilization. NTFS also has more extensions such as security access control lists (ACLs), which you can use for auditing, file-system journaling, and encryption. NTFS allows for file system compression or encryption but not on the same volume at the same time.

NTFS is required for a number of Windows Server roles and features such as Active Directory Domain Services (AD DS), Volume Shadow Copy Service (VSS), the Distributed File System (DFS). NTFS also provides a significantly higher level of security than FAT or FAT 32.

ReFS

Windows Server 2012 first introduced ReFS to enhance the capabilities of NTFS. ReFS improves upon NTFS by offering larger maximum sizes for individual files, directories, disk volumes, and other items. Additionally, ReFS offers greater resiliency, meaning better data verification, error correction, and scalability.

You should use ReFS with Windows Server 2019 for very large volumes and file shares to overcome the NTFS limitation of error checking and correction. File compression and encryption aren't available in ReFS. Also, you can't use ReFS for the boot volume.

Sector size

When you format a disk using a particular file system, you must specify the appropriate sector size. In the **Format Partition** dialog box, the sector size is described as the *Allocation unit size*, which refers to the amount of space consumed based on the particular disk's writing algorithm. You can select anywhere from 512 bytes to 64 kilobytes (KB). To improve performance, try to match the allocation unit size as closely as possible to the typical file or record size that will write to the disk.

For example, if you have a database that writes 8,192-byte records, the optimum allocation unit size would be 8 KB. This setting would allow the operating system to write a complete record in a single allocation unit on the disk. By using a 4 KB allocation unit size, the operating system would have to split the record across two allocation units, and then update the disk's master file table with the fact that the allocation units were linked. By using an allocation unit at least as big as the record, you can reduce the workload on the server's disk subsystem.

Note: the smallest writable unit is the allocation unit. If your database records are all 4,096 bytes, and your allocation unit size is 8 KB, then you'll be wasting 4,096 bytes per database write.

Why use ReFS in Windows Server?

As previously mentioned, Resilient File System (ReFS) is the latest file system created by Microsoft for server operating systems, extremely large volumes, and file shares. *ReFS* is a file system that's based on the NTFS file system. It has the following advantages over NTFS:

- Metadata integrity with checksums
- Expanded protection against data corruption
- Maximum reliability, especially during a loss of power (while NTFS has been known to experience corruption in similar circumstances)
- Large volume, file, and directory sizes
- Storage pooling and virtualization, which makes creating and managing file systems easier
- Redundancy for fault tolerance
- Disk scrubbing for protection against latent disk errors
- Resiliency to corruptions with recovery for maximum volume availability
- Shared storage pools across machines for more failure tolerance and load balancing

ReFS inherits some features from NTFS, including:

- BitLocker Drive Encryption
- Access control lists (ACLs) for security

- Update sequence number (USN) journal
- Change notifications
- Symbolic links, junction points, mount points and reparse points
- Volume snapshots
- File IDs

ReFS uses a subset of NTFS features, so it maintains backward compatibility with NTFS. Therefore, programs that run on Windows Server can access files on ReFS just as they would on NTFS. You can use ReFS drives with Windows 10 and Windows 8.1 only when formatting two- or three-way mirrors.

NTFS enables you to change a volume's allocation unit size. However, with ReFS, each volume has a fixed size of 64 KB, which you can't change. ReFS doesn't support Encrypting File System (EFS) for files or file level compression, but as previously mentioned, it does support BitLocker encryption.

As its name implies, the new file system offers greater resiliency, meaning better data verification, error correction, and scalability.

Compared to NTFS, ReFS offers larger maximum sizes for individual files, directories, disk volumes, and other items, which the following table lists.

Table 1: Attributes and limits of ReFs

Attribute	Limit
Maximum size of a single file	Approximately 16 exabytes (EB) (18,446,744,073,709,551,616 bytes)
Maximum number of files in a directory	2 ⁶⁴
Maximum number of directories in a volume	2 ⁶⁴
Maximum file name length	32,000 Unicode characters
Maximum path length	32,000
Maximum size of any storage pool	4 petabytes (PB)
Maximum number of storage pools in a system	No limit
Maximum number of spaces in a storage pool	No limit

When to use ReFS

ReFS is ideal for the following situations:

- Microsoft Hyper-V workloads. ReFS has performance advantages when using both .vhd and .vhdx files.
- Storage Spaces Direct. In Windows Server, nodes in a cluster can share direct-attached storage. In this situation, ReFS provides improved throughput, but also supports higher capacity disks used by the cluster nodes.
- Archive data. The resiliency that ReFS provides means it's a good choice for data that you want to retain for longer periods.

When not to use ReFS

ReFS shouldn't be used in the following situations:

- When you need backwards compatibility when dual booting to an operating system older than Windows 8 and Windows Server 2012

- When used with removable media
- When used as the boot drive
- When you need file-level compression or encryption

Overview of disk volumes

When selecting a type of disk for use in Windows Server, you can choose between basic and dynamic disks.

Basic disk

A basic disk is initialized for simple storage. It contains partitions, such as primary partitions and extended partitions. Basic storage uses partition tables that all versions of the Windows operating system can use. You can subdivide extended partitions into logical volumes.

By default, when you initialize a disk in the Windows operating system, the disk is configured as a basic disk. It's easy to convert basic disks to dynamic disks without any data loss. However, you can't convert a dynamic disk back to basic disk without losing all the data on the disk. Instead, you'll need to back up the data before you do such a conversion.

Converting basic disks to dynamic disks doesn't offer any performance improvements, and some programs can't address data that is stored on dynamic disks. For these reasons, most administrators don't convert basic disks to dynamic disks unless they need to use some of the additional volume-configuration options that dynamic disks provide.

Dynamic disk

Dynamic storage enables you to perform disk and volume management without having to restart computers that are running Windows operating systems. A *dynamic disk* is a disk that you initialize for dynamic storage, and that contains dynamic volumes. Dynamic disks are used for configuring fault-tolerant storage.

When you configure dynamic disks, you create volumes rather than partitions. A *volume* is a storage unit that's made from free space on one or more disks. You can format the volume with a file system and then assign it a drive letter or configure it with a mount point.

Note: Microsoft has deprecated dynamic disks from the Windows operating system and no longer recommends using them. Instead, when you want to pool disks together into larger volumes you should consider using basic disks or the newer Storage Spaces technology. If you want to mirror the volume from which the Windows operating system boots, you might want to use a hardware RAID controller, such as the one included on most motherboards.

Required disk volumes

Regardless of which type of disk you use, you must configure both a system volume and a boot volume on one of the server's hard disks:

- **System volumes.** The system volume contains the hardware-specific files that the Windows operating system needs to load, such as Bootmgr and BOOTSECT.bak. The system volume can be the same as the boot volume, although this isn't required.

- Boot volumes. The boot volume contains the Windows operating system files that are in the %Systemroot% and %Systemroot%\System32 folders. The boot volume can be the same as the system volume, although this isn't required.

Types of dynamic disk volumes

In Windows Server, if you're using dynamic disks you can create a number of disk volumes types, including:

- Simple volumes. A *simple volume* is a volume that uses free space from a single disk. It can be a single region on a disk, or consist of multiple, concatenated regions. You can extend a simple volume within the same disk or extend it to more disks. If you extend a simple volume across multiple disks, it becomes a spanned volume.
- Spanned volumes. A *spanned volume* is a volume that's created from the free disk space accumulated from multiple linked disks. You can extend a spanned volume onto a maximum of 32 disks. You can't mirror a spanned volume, and they aren't fault tolerant. Therefore, if you lose one disk, you'll lose the entire spanned volume.
- Striped volumes. A *striped volume* is a volume that has data spread across two or more physical disks. The data on this type of volume is allocated alternately and evenly to each of the physical disks. A striped volume can't be mirrored or extended, and it isn't fault tolerant. This means that the loss of one disk causes the immediate loss of all the data. Striping also is known as *RAID-0*.
- Mirrored volumes. A *mirrored volume* is a fault-tolerant volume that has all data duplicated onto two physical disks. All the data from one disk is copied to another disk to provide data redundancy. If one of the disks fails, you can access the data from the remaining disk. However, you can't extend a mirrored volume. Mirroring also is known as *RAID-1*.
- RAID-5 volumes. A *RAID-5 volume* is a fault-tolerant volume that has data striped across a minimum of three or more disks. Parity also is striped across the disk array. If a physical disk fails, you can recreate the part of the RAID-5 volume that was on that failed disk by using the remaining data and the parity. You can't mirror or extend a RAID-5 volume.

Demonstration: Manage volumes in Windows Server

In this demonstration, you'll learn how to:

- Create a new mirrored volume with Diskpart.
- Use the Windows Admin Center to identify Shares on the remote server.

Demonstration steps

Create a new mirrored volume with Diskpart

- On **SEA-SVR3**, use diskpart.exe to create a mirrored dynamic disk using disk 1 and 2, with the following settings:
 - On each disk, clear the readonly attribute.
 1. Initialize each disk with no errors.

2. Convert each disk to dynamic.
3. Create a volume mirror using both disks.
4. Format as NTFS file system, with the quick option, label the disk as **Mirrored Volume**, and assign it to drive letter **M**:
5. Close Diskpart when done.

Use the Windows Admin Center to create a share on the remote server

1. On **SEA-ADM1**, in **WAC**, connect to **SEA-SVR3** with the **Contoso\Administrator** credentials.
2. In the **Files** node of **SEA-SVR3**, in the **M: Mirrored Volume** drive, create a new folder named **Corp-Data**.
3. Share the folder, and then ensure the **Contoso\Managers** group have read and write access.

Copy files to Mirrored Volume\Corpdata, break mirror, and note Corpdata still exists

1. On **SEA-ADM1**, open Windows PowerShell and copy the **CreateLabFiles.cmd** file to the Corpdata share using the following command:

```
copy C:\labfiles\mod04\CreateLabFiles.cmd \\sea-SVR3\Corpdata
```
2. Return to **SEA-SVR3**, and then browse to **M:\Corpdata**. Verify that **CreateLabfiles.cmd** is included in the folder.
3. Use Diskpart to review the volumes. and then
4. Select the mirrored volume and remove drive 2 (Break disk=2).
5. Close Diskpart and examine **M:\Corpdata**. Note that the files are still there.
6. Return to Diskpart and examine the list of volumes. Note that drive E: was the removed disk 2 from the mirror.

Overview of File Server Resource Manager

You can use 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 can also define quota templates that you can apply to new volumes or folders.
- **File screening management.** You can use this feature to control the types of files that users can store on a file server. You can limit the file types with specific file extensions that users can store on your file shares. For example, you can create a file screen that doesn't allow users to save files with an.mp3 extension in a file server's personal shared folders. File screening only searches for the file extension. It doesn't examine the file contents.
- **Storage reports.** You can use this feature to identify trends in disk usage and how your data is classified. You can also 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:
 - File location
 - Classification properties
 - File creation date
 - File modification date
 - 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 receive in Windows client operating systems when they don't have access to a file or a folder.

Note: You can access FSRM by using the **File Server Resource Manager Microsoft Management Console (MMC)** console, or by using Windows PowerShell.

You can access all available cmdlets by running the following command at a Windows PowerShell command prompt:

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

Manage permissions on volumes

You can configure file and folder permissions only on NTFS file system and Resilient File System (ReFS) volumes. *Permissions* are rules that determine what operations specific users can perform on a file or a folder. An owner of a file or folder, or anyone with full control access can grant or deny permissions to that file or folder.

Typically, you assign permissions to groups to minimize administrative overhead. However, if you assign permissions to a group, every group member has the same permissions that you assign. You can also assign permissions to individual users and computers. Permissions are cumulative when you assign permissions to a group and to individual group members. This means that a user has the permissions that you assign to them, in addition to those you assign to the group.

Permissions example

When you assign multiple permissions, consider the following example: Anthony is a member of the Marketing group, which has the Read permission added to the **Pictures** folder. You assign the Write permission to Anthony for the **Pictures** folder. Anthony now will have Read and Write permissions because he is a member of the Marketing group, and you assigned the Write permission directly to him.

Types of permissions

There are two types of permissions that you can configure for files and folders on NTFS file systems and ReFS volumes:

- **Basic.** Basic permissions are the most commonly used permissions, such as Read or Write permissions. You typically assign them to groups and users, and each basic permission is built from multiple advanced permissions.
- **Advanced.** Advanced permissions provide an additional level of control. However, advanced permissions are more difficult to document and more complex to manage. For example, the basic Read permission is built from the:
 - List folder/read data
 - Read attributes
 - Read extended attributes
 - Advanced Read permissions

Basic file and folder permissions

The following list describes the basic file and folder permissions, which you can choose to allow or deny:

- **Full control.** Provides users with complete control of the file, folder, and permissions.
- **Modify.** Allows users to read a file, write changes to it, and modify permissions. The advanced permissions that make up the modify permissions are:
 - Traverse folder/execute file
 - List folder/read data
 - Read attributes
 - Read extended attributes
 - Create files/write data
 - Create folders/append data
 - Write attributes
 - Write extended attributes
 - Delete
 - Read permissions
- **Read & execute.** Allows users to access folder content, read files, and start programs. This applies to an object and any of the child objects by default. The advanced permissions that make up the Read & execute permissions are:
 - Traverse folder/execute file
 - List folder/read data
 - Read attributes
 - Read extended attributes

- Read permissions
- Read. Allows users to read a file but doesn't allow them to change it. This applies to an object and any of the child objects by default. The Advanced permissions that make up the Read permissions are:
 - List folder/read data
 - Read attributes
 - Read extended attributes
 - Read permissions
- Write. Allow users to change folder and file content. This applies to an object and any of the child objects by default. The advanced permissions that make up Write permissions are the:
 - Create files/write data
 - Create folders/append data
 - Write attributes
 - Write extended attributes
- Special. This is a custom configuration.

Note: Groups or users that have the Full control permission on a folder can delete any files in that folder regardless of the permissions that protect the file.

To modify permissions, you must have the Full control permission for a folder or file. The one exception is for file and folder owners, who can modify permissions even if they don't have any current permissions. Administrators can exercise a special *right*, or privilege, to take ownership of files and folders to make modifications to permissions.

Important rules for file permissions

There are two important rules for file permissions:

- Explicit versus inherited. Permissions that you explicitly assign take precedence over those that are inherited from a parent folder.
- Deny vs. Allow. Within a set of explicit permissions, Deny permissions override conflicting Allow permissions. Likewise, within a set of implicit, inherited permissions, Deny permissions override conflicting Allow permissions.

Therefore, taking these rules into account, file permissions apply in the following order:

1. Explicit Deny
2. Explicit Allow
3. Inherited Deny
4. Inherited Allow

Note: If set at the same level, the Deny permissions override any Allow permissions. If set at different levels, the permission that is set at the lower level overrides the permission that is set at the higher level. However, if Deny is set two levels up while Allow is set on the parent folder, both permissions will be inherited but Allow would override Deny.

Test your knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

What are the two disk types in Windows 10 Disk Management?

Question 2

What file system do you currently use on your file server and will you continue to use it?

Question 3

If permissions on a file are inherited from a folder, can you modify them on a file?

Question 4

Can you set permissions only on files in NTFS volumes?

Implementing sharing in Windows Server

Lesson overview

Collaboration is an important part of an administrator's job. Your organization might create documents that only certain team members can share, or you might work with a remote team member who needs access to a team's files. Because of collaboration requirements, you must understand how to manage shared folders in a network environment.

Sharing folders enables users to connect to a shared folder over a network, and to access the folders and files that it contains. Shared folders can contain applications, public data, or a user's personal data. Managing shared folders helps you provide a central location for users to access common files, and it simplifies the task of backing up data that those folders contain. This lesson examines various methods of sharing folders, along with the effect this has on file and folder permissions when you create shared folders on an NTFS file system–formatted partition.

The Server Message Block (SMB) protocol is a network file sharing protocol that allows applications on a computer to read and write to files, and to request services from server programs in a computer network. The SMB protocol is used on top of TCP/IP. Using the SMB protocol, an application (or the user of an application) can access files or other resources at a remote server. This allows applications to read, create, and update files on the remote server. It can also communicate with any server program that is set up to receive an SMB client request.

Lesson objectives

After completing this lesson, you'll be able to:

- Describe SMB.
- Describe how to implement and manage SMB shares.
- Describe how to configure SMB shares by using Server Manager and Windows PowerShell.
- List the best practices for sharing Resources.
- Provide an overview of network file system (NFS).

What is SMB?

Server Message Block (SMB) is a network file sharing protocol developed by Microsoft in the 1980s. It was created to be part of the network basic input/output system (NetBIOS). The original specification, SMB 1, was designed to be a very verbose protocol, continuously sending many control and search packets in clear text. However, this was considered to be too noisy, it degraded overall network performance, and was unsecure. Since that time, the extraneous chatter has been substantially reduced and the protocol has become more secure, especially since the SMB 3.0 and higher specifications.

The latest version of SMB is SMB 3.1.1, which supports Advanced Encryption Standard (AES) 128 Galois/Counter Mode (GCM) encryption in addition to the AES 128 Counter with cipher block chaining message authentication code (CBC-MAC, or CCM) encryption that is included in SMB 3.0. SMB 3.1.1 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.

Microsoft Hyper-V supports storing virtual machine (VM) data—such as VM configuration files, checkpoints, and .vhd files—on SMB 3.0 and later file shares.

Note: The recommended 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 VM files on Internet Small Computer System Interface (iSCSI) or Fibre Channel storage area network (SAN) devices. When creating a VM in Hyper-V on Windows Server, you can specify a network share when choosing the VM location and the virtual hard disk location. You can also 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.

Windows Server continues to support the SMB 3.0 enhancements in addition to several advanced functions that you can employ by using SMB 3.1.1. For example, you can store VM 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 Microsoft Hyper-V nodes, but by using file servers that host VM files on their file shares. With this capability, Hyper-V can store all VM files, including configuration files, .vhd files, and checkpoints on highly available SMB file shares.

The SMB 3.0 features include:

- **SMB Transparent Failover.** This feature enables 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 input/output (I/O), through all the nodes in a file server cluster.
- **SMB Multichannel.** This feature enables 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 central processing unit (CPU) processing time.
- **SMB Encryption.** This feature provides the end-to-end encryption of SMB data on untrusted networks, and it 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 also reduces the number of round trips from client to the 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 features in SMB 3.1.1 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, uses a fixed cryptographic algorithm, AES-128-CCM. However, AES-128-GCM performs better with 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 don't perform this negotiation correctly, 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 languages with a nonzero revision number.** Windows Server uses three separate digits to notate the version of SMB. This information is then used to negotiate the highest level of SMB functionality.

SMB versions

Windows Server will negotiate and use the highest SMB version that a client supports. In this regard, the client can be another server, a Windows 10 device, or even an older legacy client or Network-attached storage (NAS) device. This support can go down to SMB 2.2 or 2.0. Older Windows Server versions also include support for SMB 1.0, which is known for its vulnerabilities. Therefore, the use of SMB 1.0 should be avoided for security reasons. In Windows Server version 1709, and Windows 10 version 1709 (and later) support for SMB 1.0 isn't installed by default.

Configure SMB shares

Creating and configuring file shares has long been a core part of network administration. The ability to share files is one of the reasons that computer networks first became popular. Most administrators are aware that you can create shared folders or file shares from within File Explorer. However, in Windows Server you can also create file shares by using the Windows Admin Center, Server Manager, or Windows PowerShell. In Server Manager, the terms *file share* and *SMB share* refer to the same component.

Share and file permissions

The permissions a user must use to access files on a Server Message Block (SMB) share are a combination of share permissions and file permissions. The most restrictive set of permissions always applies. For example, if you give a user the Full Control file permissions, but they have only Read share permissions, the user's access is Read.

Note: In some previous Microsoft literature, file permissions used to be called *NTFS file system permissions*. The term is gradually being replaced with just *file permissions*, as these permissions also apply to Resilient File System (ReFS), and not just NTFS.

To simplify data access, when you use the **Quick** profile for creating an SMB share, the share permission is set to **Everyone Full Control**. Effectively this means that share permissions aren't restricting access to the share, and NTFS permissions are used to control access.

SMB share profiles

You can use Server Manager on Windows Server to create a new share. The built-in **New Share Wizard** offers three SMB file share profiles from which you can choose:

- **Quick.** This is the fastest method for sharing a folder on a network. With this method, you can select a volume or enter a custom path for the shared folder location. You can use the **New Share Wizard** to configure more options, such as access-based enumeration, share caching, encrypted data access, and permissions. You can configure these options and other options manually after you create the share.
- **Advanced.** This profile offers the same configuration options as the quick profile, in addition to more options such as folder owners, default data classification, and quotas. To create an advanced profile,

you must install the **File Server Resource Manager** role service on at least one server that you're managing by using Server Manager.

- **Applications.** This specialized profile has appropriate settings for Microsoft Hyper-V, databases, and other server applications. Unlike the quick and advanced profiles, you can't configure access-based enumeration, share caching, default data classification, or quotas when you're creating an applications profile.

The following table identifies the configuration options that are available for each SMB share profile.

Table 1: Configuration options available for each SMB share profile

Share type	Access-based enumeration	Share caching	Encrypted data access	Default data classification	Quotas	Permissions
Quick	Yes	Yes	Yes	No	No	Yes
Advanced	Yes	Yes	Yes	Yes	Yes	Yes
Applications	No	No	Yes	No	No	Yes

Windows Admin Center

You can use the Windows Admin Center to create SMB shares, but you can't configure them beyond user permissions. To enable access-based enumeration, allow for caching, enable branch cache, or to add or remove encryption, you can use the folder properties, Server Manager, or the Windows PowerShell commands covered in the next section.

Windows PowerShell cmdlets in the SmbShare module

The **SmbShare** module for Windows PowerShell contains 38 cmdlets in Windows Server. This includes commonly used cmdlets such as **New-SmbShare**, **Set-SmbShare**, and **Remove-SmbShare**. If you use the **SmbShare** cmdlets, you can configure any share properties, even those that aren't available in Server Manager.

If you want to identify the shares that exist on a server, or review the properties of those shares, you can use **Get-SmbShare**. The default output displays the **Name**, **ScopeName**, **Path**, and **Description**. **ScopeName** is only relevant when the server is part of a cluster and displays as * for unclustered file servers.

You can use **Get-SmbSession** to identify users that are connected to SMB shares. If the users have open files, then you can use **Get-SmbOpenFile** to identify the open files.

If you're concerned about controlling the bandwidth allocated to SMB shares on a server, you can use **Set-SMBBandwidthLimit** to define a maximum throughput level that is allocated to SMB traffic on a server for different categories. This is useful for Hyper-V hosts to ensure that certain categories of traffic don't overwhelm the host and affect other categories, including:

- **Default.** This refers to all SMB traffic that doesn't relate to Hyper-V or Live Migration, such as standard file shares.
- **Hyper-V.** This refers to SMB traffic that you use for running virtual machines, such as accessing virtual hard drives on an SMB share.
- **Live Migration.** This refers to SMB traffic that generates when you perform a live migration from one Hyper-V host to another.

Note: To explore all of the cmdlets in the **SmbShare** module, run the following command:

```
Get-Command -Module SmbShare
```

Demonstration: Configure SMB shares by using Server Manager and Windows PowerShell

In this demonstration, you'll learn how to:

- Create a Server Message Block (SMB) share by using Server Manager.
- Create an SMB share by using Windows PowerShell.
- Review SMB session information.

Create an SMB share by using Server Manager

1. On **SEA-ADM1**, in **Server Manager**, create an SMB share on the **SEA-SVR3 M:** drive, with the following properties:
 - Profile: **Quick**
 - Name: **SalesShare**
 - Setting: **Enable access-based enumeration**
 - Permission: Add **Contoso\Sales** with the **Modify** permission

Create an SMB share by using Windows PowerShell remote session

1. On **SEA-ADM1**, in **Windows PowerShell**, open a remote session to **SEA-SVR3**.
2. Create an SMB share on **SEA-SVR3** with following properties:
 - Profile: **Quick**
 - Name: **SalesShare2**
 - Setting: **Enable access-based enumeration**
3. Examine the properties of the new share

Review SMB session information

1. Use **File Explorer** on **SEA-ADM1** to open **SalesShare2** on **SEA-SVR3**.
2. Use the remote **Windows PowerShell** session on **SEA-ADM1** to observe the connection to **SEA-SVR3\SalesShare2**.

Best practices for sharing resources

The Server Message Block (SMB) protocol includes the **SMB Direct** and **SMB Multichannel** features that enable you to deploy cost-efficient, continuously available, and high-performance storage for server applications on file servers. Both **SMB Multichannel** and **SMB Direct** are enabled by default on Windows Server. Both features require you to use a Remote Direct Memory Access (RDMA)-enabled network adapter. You can use multiple network connections simultaneously with SMB Multichannel, which enhances overall file-sharing performance. **SMB Direct** ensures that multiple network adapters can

coordinate the transfer of large amounts of data at line speed while using fewer central processing unit (CPU) cycles.

Here are some best practices to use when sharing resources and making a file server highly available:

- Use large physical disks that are resilient against downtime caused by potential file corruption.
- Microsoft SQL Server, Microsoft Hyper-V, and other such server applications should deploy on continuously available file servers.

Note: Continuously available file server concepts will be discussed in Module 6, “High availability in Windows Server.”

- Use Dynamic Host Configuration Protocol (DHCP) failover services to improve network availability.
- Aggregate bandwidth and maximize network reliability by using Load Balancing and Failover.
- Create continuously available block storage for server applications by using Internet Small Computer System Interface (iSCSI) transparent failover.
- Use RDMA and SMB Direct.
- Use multiple network adapters by using SMB Multichannel.
- Use Offloaded Data Transfer to move data quickly between storage devices.
- Create continuously available network file system (NFS) file shares by using NFS transparent failover.
- Use SMB Volume Shadow Copy Service (VSS) for remote file shares to protect your application data.
- In Hyper-V, put virtual machine files on SMB file shares, which will give you flexible storage solutions for future virtual or cloud infrastructure.

Overview of NFS

Network file system (NFS) is a file-system protocol that's based on open standards and allows access to a file system over a network. NFS has been developed actively, and the current version is 4.1. The core releases and characteristics of the NFS protocol are:

- **NFS version 1.** Initially, NFS was used on UNIX operating systems, but was subsequently supported on other operating systems, including Windows.
- **NFS version 2.** This version focused on improving performance. There is a file-size limit of 2 gigabytes (GB), because it was a 32-bit implementation.
- **NFS version 3.** This version introduced support for larger file sizes because it was a 64-bit implementation. It also had performance enhancements such as better protection from unsafe writes, increased transfer sizes, and security enhancements such as over-the-wire permission checks by the server.
- **NFS version 4.** This version provided enhanced security and improved performance.
- **NFS version 4.1.** This version added support for clustering.

In UNIX, NFS works based on exports. Exports are similar to folder shares in Windows because they are shared UNIX file-system paths.

The two components for NFS support in Windows are:

- **Client for NFS.** This component enables a computer running a Windows operating system to access NFS exports on an NFS server, regardless of which platform the server is running on.

- **Server for NFS.** This component enables a Windows-based server to share folders over NFS. Any compatible NFS client can access the folders, regardless of which operating system the client is running on. The vast majority of UNIX and Linux computers have a built-in NFS client.

Support for NFS has been improved and expanded with each iteration of the Windows Server operating system. Support for Kerberos protocol version 5 (v5) authentication is available in Server for NFS. Kerberos protocol v5 authentication provides authentication before granting access to data. It also uses checksums to ensure that no data tampering has occurred. Windows Server also supports NFS version 4.1. This support included improved performance with the default configuration, native Windows PowerShell support, and faster failovers in clustered deployments.

Usage scenarios

You can use NFS in Windows in many scenarios. Some of the most popular uses include:

- VMWare virtual machine (VM) storage. In this scenario, VMWare hosts VMs on NFS exports. You can use Server for NFS to host the data on a Windows Server computer.
- Multiple operating-system environments. In this scenario, your organization uses a variety of operating systems, including Windows, Linux, and Mac. The Windows file-server system can use Server for NFS and the built-in Windows sharing capabilities to ensure that all the operating systems can access shared data.
- Merger or acquisition. In this scenario, two companies are merging. Each company has a different IT infrastructure. Users from one company use Windows 10 client computers, and they must access data that the other company's Linux and NFS-based file servers are hosting. You can deploy Client for NFS to the client computers to enable the users to access the data.

Test your knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

Can any user connect to any shared folder?

Question 2

What could be a reason that a user can't open files on a share?

Question 3

What is the main difference between sharing a folder by using "Network File and Folder Sharing" and by using the "Advanced Sharing" option?

Question 4

What could be a reason that a user doesn't have the "Always available offline" option when they right-click or access the context menu for a file in the share, but when they right-click or access the context menu for a file in another share, the "Always available offline" option is available?

Implementing Storage Spaces in Windows Server

Lesson overview

Managing physical disks that are attached directly to a server can be a tedious task. To address this issue and make more-efficient use of storage, many organizations have implemented storage area networks (SANs). However, SANs require special configurations and specific hardware in some scenarios. They can also 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 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.

Microsoft includes Storage Spaces Direct in the Windows Server Datacenter edition, which enables you to create a highly available storage solution using local storage from multiple Windows Server computers.

Lesson objectives

After completing this lesson, you'll be able to:

- Describe Storage Spaces.
- Describe the components and features of Storage Spaces.
- Describe how the uses of Storage Spaces.
- Explain how to provision a Storage Space.
- Describe how to configure Storage Spaces.
- Describe Storage Spaces Direct.
- Describe how to configure Storage Spaces Direct.

What are Storage Spaces?

A *storage space* is a storage-virtualization capability built into Windows Server 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 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.** *Storage Spaces* 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 storage area network (SAN).

You can manage Storage Spaces by using Windows Storage Management application programming interface (API) in Windows Management Instrumentation (WMI) and Windows PowerShell. You can also use the File and Storage Services role in **Server Manager** to manage Storage Spaces.

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

- Physical disks. *Physical disks* are disks such as SATA or serial-attached SCSI (SAS) 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 are 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 the disks.
 - You can attach disks by using a variety of bus interfaces, including:
 - SCSI
 - SAS
 - SATA
 - NVMe Express

If you want to use failover clustering with storage pools, you can't use SATA, USB, or SCSI disks.

- Storage pool. A *storage pool* is a collection of one or more physical disks that you can use to create virtual disks. You can add available, unformatted physical disks to a storage pool. Note that you can attach a physical disk to only one storage pool. However, several physical disks can be in that storage pool.
- Virtual disk (or storage space). This is like 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. Virtual disks resemble Redundant Array of Independent Disks (RAID) technologies, but Storage Spaces store the data differently than RAID.
- Disk drive. You can make your disk drives available from your Windows operating system by using a drive letter.

You can format a storage space virtual disk with FAT32, NTFS, and Resilient File System (ReFS). However, you must format the virtual disk with NTFS or ReFS to use it with Data Deduplication or with File Server Resource Manager (FSRM).

Other features in Windows Server Storage Spaces include:

- Tiered Storage Spaces. The Tiered Storage Spaces feature allows you to use a combination of disks in a storage space. For example, you could use very fast but small-capacity hard disks such as solid-state drives (SSDs) 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 often to the slower disks.

By default, the Storage Spaces feature moves data once a day at 01:00 AM. You can also 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 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. However, it's also limited to 1 gigabyte (GB) of data.

- Windows Server 2019 added support for persistent memory (PMem). You use PMem as a cache to accelerate the active working set, or as capacity to guarantee consistent low latency on the order of microseconds.

Components and features of Storage Spaces

When configuring Storage Spaces, a key step is planning for implementing virtual disks. Before you implement virtual disks, you should consider the following Storage Spaces features:

- **Storage layout:** Storage layout is one of the characteristics that defines the number of disks from the storage pool that are allocated. Valid options include:
 - A simple space, which has data striping but no redundancy. In data striping, logically sequential data segments across several disks in a way that enables different physical storage drives to access these sequential segments. Striping can improve performance because it's possible to access multiple segments of data simultaneously. To enable data striping, you must deploy at least two disks. The simple storage layout doesn't provide any redundancy, so if one disk in the storage pool fails, you'll lose all data unless you have a backup.
 - *Two-way and three-way mirrors.* Mirroring helps provide protection against the loss of one or more disks. Mirror spaces maintain two or three copies of the data that they host. Specifically, two-way mirrors maintain two data copies, and three-way mirrors maintain three data copies for three-way mirrors. Duplication occurs with every write to ensure that all data copies are always current. Mirror spaces also stripe the data across multiple physical drives. To implement mirroring, you must deploy at least two physical disks. Mirroring provides protection against the loss of one or more disks, so use mirroring when you're storing important data. The disadvantage of using mirroring is that the data duplicates on multiple disks, so disk usage is inefficient.
 - *Parity.* A parity space resembles a simple space because data writes across multiple disks. However, parity information also writes across the disks when you use a parity storage layout. You can use the parity information to calculate data if you lose a disk. Parity enables Storage Spaces to continue to perform read-and-write requests even when a drive has failed. The parity information always rotates across available disks to enable input/output (I/O) optimization. A storage space requires a minimum of three physical drives for parity spaces. Parity spaces have increased resiliency through journaling. The parity storage layout provides redundancy but is more efficient in utilizing disk space than mirroring. **Note:** The number of columns for a given storage space can also impact the number of disks.
- **Disk sector size:** A storage pool's sector size is set the moment it's created. Its default sizes are set as follows:
 - If the list of drives being used contains only 512 and 512e drives, the pool sector size is set 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-kilobyte (KB) drive, the pool sector size is set to 4 KB.
- **Cluster disk requirement:** Failover clustering prevents work interruptions if there is a computer failure. For a pool to support failover clustering, all drives in the pool must support serial attached SCSI (SAS).

- **Drive allocation:** Drive allocation defines how the drive allocates to the pool. Options are:
 - Data-store. This is the default allocation when any drive is added to a pool. Storage Spaces can automatically select available capacity on data-store drives for both storage space creation and just-in-time (JIT) allocation.
 - A manual drive isn't used as part of a storage space unless it's specifically selected when you create that storage space. This drive allocation property lets administrators specify particular types of drives for use only by certain Storage Spaces.
 - Hot spare. These are reserve drives that aren't used in the creation of a storage space but are added to a pool. If a drive that is hosting storage space columns fails, one of these reserve drives is called on to replace the failed drive.
- **Provisioning schemes:** You can provision a virtual disk by using one of two methods:
 - Thin provisioning space. Thin provisioning enables storage to be allocated readily on a just-enough and JIT basis. Storage capacity in the pool is organized into provisioning slabs that aren't allocated until datasets require the storage. Instead of the traditional fixed storage allocation method in which large portions of storage capacity are allocated but might remain unused, thin provisioning optimizes any available storage by reclaiming storage that is no longer needed using a process known as *trim*.
 - Fixed provisioning space. In Storage Spaces, fixed provisioned spaces also use flexible provisioning slabs. The difference is that it allocates the storage capacity up front when you create the space. You can create both thin and fixed provisioning virtual disks within the same storage pool. Having both es in the same storage pool is convenient, especially when they are related to the same workload. For example, you can choose to use a thin provisioning space for a shared folder containing user files, and a fixed provisioning space for a database that requires high disk I/O.
- **Stripe parameters:** You can increase the performance of a virtual disk by striping data across multiple physical disks. When creating a virtual disk, you can configure the *stripe* by using two parameters, *NumberOfColumns* and *Interleave*. A *stripe* represents one pass of data written to a storage space, with data written in multiple stripes, or *passes*. *Columns* correlate to underlying physical disks across which one stripe of data for a storage space is written. *Interleave* represents the amount of data written to a single column per stripe. The *NumberOfColumns* and *Interleave* parameters determine the width of the stripe (e.g., *stripe \ width = NumberOfColumns* and *Interleave*). In the case of parity spaces, the stripe width determines how much data and parity Storage Spaces writes across multiple disks to increase performance available to apps. You can control the number of columns and the stripe interleave when creating a new virtual disk by using the Windows PowerShell cmdlet **New-VirtualDisk** with the **NumberOfColumns** and **Interleave** parameters.

When creating pools, Storage Spaces can use any direct-attached storage (DAS) device. You can use Serial ATA (SATA) and SAS drives (or even older integrated drive electronics (IDE) and small computer system interface (SCSI) drives) that are connected internally to the computer.

When planning your Storage Spaces storage subsystems, you must consider the following factors:

- Fault tolerance. Do you want data to be available if a physical disk fails? If so, you must use multiple physical disks and provision virtual disks by using mirroring or parity.
- Performance. You can improve performance for read and write actions by using a parity layout for virtual disks. You also need to consider the speed of each individual physical disk when determining performance. Alternatively, you can use disks of different types to provide a tiered system for storage.

For example, you can use solid-state drives (SSDs) for data to which you require fast and frequent access, and you can use SATA drives for data that you don't access as frequently.

- Reliability. Virtual disks in parity layout provide some reliability. You can improve that degree of reliability by using hot-spare physical disks in case a physical disk fails.
- Future storage expansion. One of the main advantages of using Storage Spaces is the ability to expand storage in the future by adding physical disks. You can add physical disks to a storage pool any time after you create it to expand its storage capacity or to provide fault tolerance.

Storage Spaces usage scenarios

When considering whether to use Storage Spaces in a given situation, you should weigh its benefits and limitations. The benefits of Storage Spaces are that it can:

- Implement and easily manage scalable, reliable, and inexpensive storage.
- Aggregate individual drives into storage pools, which are managed as a single entity.
- Use inexpensive storage with or without external storage.
- Use different types of storage in the same pool (for example, SATA, SAS, USB, and SCSI).
- Grow storage pools as required.
- Provision storage when required from previously created storage pools.
- Designate specific drives as hot spares.
- Automatically repair pools containing hot spares.
- Delegate administration by pool.
- Use the existing tools for backup and restore, and use Volume Shadow Copy Service (VSS) for snapshots.
- Manage either locally or remotely, by using Microsoft Management Console (MMC) or Windows PowerShell.
- Utilize Storage Spaces with Failover Clusters.

Note: While this list mentions USB as a supported storage medium, using USB in a pool might be more practical on a Windows client or while developing a proof of concept. USB performance also depends on the performance capabilities of the storage you choose to pool together.

Storage Spaces also has some inherent limitations. For example, in Windows Server:

- Storage Spaces volumes aren't supported on boot or system volumes.
- You should add only unformatted, or non-partitioned, drives.
- You must have at least one drive in a simple storage pool.
- Fault tolerant configurations have specific requirements:
 - A mirrored pool requires a minimum of two drives.
 - Three-way mirroring requires a minimum of five drives.
 - Parity requires a minimum of three drives.
- All drives in a pool must use the same sector size.
- Storage layers that abstract the physical disks aren't compatible with Storage Spaces, including:
 - VHDs and pass-through disks in a virtual machine (VM).

- Storage subsystems deployed in a separate RAID layer.
- Fibre Channel and Internet Small Computer System Interface (iSCSI) aren't supported.
- Failover Clusters are limited to SAS as a storage medium.

Note: Microsoft Support provides troubleshooting assistance only in environments where you deploy Storage Spaces on a physical machine, not a VM. In addition, Microsoft must certify just a bunch of disks (JBOD) hardware solutions that you implement.

When planning for the reliability of a workload in your environment, Storage Spaces provide different resiliency types. As a result, some workloads are better suited for specific resilient scenarios. The following table depicts these recommended workload types.

Table 1: Recommended workload types

Resiliency type	Number of data copies maintained	Workload recommendations
Mirror	2 (two-way mirror) 3 (three-way mirror)	Recommended for all workloads
Parity	2 (single parity) 3 (dual parity)	Sequential workloads with large units of read/write, such as archival
Simple	1	Workloads that don't need resiliency, or provide an alternate resiliency mechanism

Provision 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 defaults to 512e. A *512 drive* 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-kilobyte (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 users 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.

Drive allocation

You can configure how a pool allocates drives. 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 just-in-time (JIT) allocation.
- **Manual.** You can specify **Manual** as the usage type for drives that you add to a pool. A storage space won't use a manual drive automatically unless you select it when you create that storage space. This

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 won't use when creating a storage space. However, if a failure occurs on a drive that is hosting columns of a storage space, the hot-spare reserve drive replaces the failed drive.

Provisioning schemes

You can provision a virtual disk by using one of two provisioning 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 but doesn't allocate them until datasets grow to the required storage size. 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 can also save on operating costs such as electricity and floor space, which are required to keep even unused drives in operation. 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 time that you create the storage space.

Demonstration: Configure Storage Spaces

In this demonstration, you'll learn how to:

- Create a storage pool.
- Create a mirrored virtual disk and volume.
- Examine disk properties in Windows Admin Center.

Create a storage pool

1. On **SEA-ADM1**, open **Server Manager**.
2. In Server Manager, in **File and Storage Services**, and then select **Disks**.
3. Set **SEA-SVR3** disks 1 through 4 to **Online**.
4. In Server Manager, on **SEA-SVR3**, create a new storage pool named **SP1**. Use two disks of the four available to make up the pool.

Create a mirrored virtual disk and a volume named Corp Data

1. In Server Manager, in **Storage Pools**, in **SP1**, create a new virtual disk named **Mirrored** that uses a mirror storage layout and thin provisioning. Use **25 GB** for the size.
2. Create a new volume from **Mirrored** named **Corp Data**, and assign it drive letter **M**.
3. Close Server Manager.

Examine disk properties in Windows Admin Center

1. On **SEA-ADM1**, in **WAC**, connect to **SEA-SVR3** with the **Contoso\Administrator** credentials.
2. Open the **Files** node and examine the new Storage Spaces drive named **Corp Data**.

Overview of Storage Spaces Direct

Storage Spaces might not have you need for big-scale installations, depending on the configuration. There's a limit to how many just a bunch of disks (JBODs) you can interconnect to physical servers. However, Storage Spaces Direct in Windows Server overcomes these limitations.

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

Windows Server 2019 improvements to Storage Spaces Direct

Windows Server 2019 includes several new improvements to Storage Spaces Direct. These improvements incorporate many advances in hardware and other technologies related to storage. New improvements include:

- Deduplication and compression for Resilient File System (ReFS) volume. Store up to ten times more data on the same volume with deduplication and compression for the ReFS file system. The variable-size chunk store container with optional compression maximizes savings rates, while the multi-threaded post-processing architecture keeps performance impact minimal. Deduplication supports volumes up to 64 terabytes (TB) and will deduplicate the first 4 TB of each file.
- Native support for persistent memory. Unlock better performance with native Storage Spaces Direct support for persistent memory modules. Use persistent memory as cache to accelerate the active working set, or as capacity to guarantee consistent, low latency on the order of microseconds. Manage persistent memory just as you would any other drive in Windows PowerShell or Windows Admin Center.
- Nested resiliency for two-node hyper-converged infrastructure at the edge. Survive two simultaneous hardware failures with an all-new software resiliency option inspired by RAID 5+1. With nested resiliency, a two-node Storage Spaces Direct cluster can provide continuously accessible storage for apps and VMs even if one server node stops working, a drive fails in the other server node.
- Two-server clusters using a USB flash drive as a witness. Use a low-cost USB flash drive that's plugged into your router to function as a witness in two-server clusters. If a server ceases operation and then comes back up, the USB drive cluster knows which server has the most up-to-date data.
- Windows Admin Center. Manage and monitor Storage Spaces Direct with the new purpose-built Dashboard and experience in Windows Admin Center. Create, open, expand, or delete volumes with just a few selects. Monitor performance such as input/output (I/O) and I/O operations per second (IOPS) latency from the overall cluster down to the individual solid-state drive (SSD) or hard disk drive (HDD).
- Performance history. Get effortless visibility into resource utilization and performance with built-in history. Over 50 essential counters spanning compute, memory, network, and storage are automatically collected and stored in the cluster for up to one year. Best of all, there's nothing to install, config-

ure, or start—it just works. You can visualize in Windows Admin Center or query and process in Windows PowerShell.

- Scale up to 4 petabytes (PB) per cluster. Achieve multi-petabyte scale, which is great for media, backup, and archival use cases. In Windows Server 2019, Storage Spaces Direct supports up to 4 PB (or 4,000 TB) of raw capacity per storage pool. Related capacity guidelines are increased as well. For example, you can create twice as many volumes (64 instead of 32), each twice as large as before (64 TB instead of 32 TB). You can also stitch multiple clusters together into a cluster set for even greater scale within one storage namespace.
- Mirror-accelerated parity is two times faster. With mirror-accelerated parity you can create Storage Spaces Direct volumes that are part mirror and part parity, similar to mixing RAID-1 and RAID-5/6 to get the best of both. In Windows Server 2019, mirror-accelerated parity performance is more than doubled compared to Windows Server 2016 thanks to optimizations.
- Drive latency outlier detection. Easily identify drives with abnormal latency with proactive monitoring and built-in outlier detection, inspired by Microsoft Azure's long-standing and successful approach. Whether it's average latency or something more subtle—such as 99th percentile latency—that stands out, slow drives are automatically labeled in Windows PowerShell and Windows Admin Center with an **Abnormal Latency** status.
- Manually delimit the allocation of volumes to increase fault tolerance. This enables administrators to manually delimit the allocation of volumes in Storage Spaces Direct. Doing so can significantly increase fault tolerance under certain conditions, but it also imposes some added management considerations and complexity.
- Storage-class memory support for VMs. This enables NTFS-formatted direct access volumes to be created on non-volatile dual inline memory modules (DIMMs) and exposed to Microsoft Hyper-V VMs. This enables Hyper-V VMs to take advantage of the low-latency performance benefits of storage-class memory devices.

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 manage this process. 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 Storage Spaces Direct.
- **Enable-ClusterS2D**. This cmdlet configures a cluster for Storage Spaces Direct 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 Storage Spaces Direct.

Note: For more information about clusters and the Failover Clustering feature, refer to Module 6, "High Availability in Windows Server."

Configuring Storage Spaces Direct using Microsoft System Center Virtual Machine Manager

Although there is no graphical user interface (GUI) to configure Storage Spaces Direct in Windows Server, you can simplify new cluster deployments that are Storage Spaces Direct-enabled by using System Center. To do so, when you run the **Create Hyper-V Cluster Wizard**, select the **Enable Storage Spaces Direct** check box. The wizard then performs the following high-level tasks:

1. Installs the relevant Windows Server roles.
2. Runs cluster validation.
3. Installs and configures failover clustering.
4. Enables storage features.

You must create the storage pool and the volumes on the cluster, and then deploy the VMs 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. These include:

- **Network.** Storage Spaces Direct needs a network for the hosts to communicate. Best practice is to have the network adapters be capable of Remote Direct Memory Access (RDMA) or have two network adapters to ensure performance and minimize latency.
- **Internal disks.** Each server or storage node has internal disks or a JBOD that connects externally.
- **Two servers.** There is a requirement to have a minimum of two servers in a Storage Spaces Direct solution. Depending on the resiliency you want to achieve, you might need more servers. Storage Spaces Direct can use up to 16 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 for each node, so they are discoverable to 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 that you create provide resiliency against both disk and server failure because data is stored on disks on different servers.
- **In Windows Server, ReFS is the primary file system to store VM files because of accelerated VHD/VHDX operations, which provide superior performance as compared to NTFS. ReFS also provides error detection and automatic correction. ReFS-formatted disks are a recommended component of Storage Spaces Direct.**
- **Cluster Shared Volumes (CSVs).** CSVs consolidate all volumes into a single namespace that's 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 isn't implemented in hyper-converged configurations in which Hyper-V runs on the same cluster as the Storage Spaces Direct feature.

Scale-Out File Server or Hyper-V scenarios

When you use Storage Spaces Direct, you determine whether you want to separate the virtualization and storage layers. You can use Storage Spaces Direct in two different scenarios, namely hyper converged solution and disaggregated solution.

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 would use this solution for small and medium-sized 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 for Hyper-V, and one for 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 independently. You use this solution for large-scale deployments. This is known as a *disaggregated solution*.

Other uses for Storage Spaces Direct are for storage of Hyper-V Replica files, or as backup or archival of VM files. You can also deploy Storage Spaces Direct in support of Microsoft SQL Server 2012 or later, which can store both system and user database files.

Demonstration: Configure Storage Spaces Direct

In this demonstration, you'll learn how to:

- Install the Windows Server roles and features for Failover Clustering.
- Validate cluster configuration and create a cluster.
- Enable the Storage Spaces Direct feature, create a storage pool, virtual disk, file server and file share.
- Test high availability for the storage.

Demonstration steps

Install the Windows Server roles and features

1. On **SEA-ADM1**, in **Server Manager**, verify that **SEA-SVR1**, **SEA-SVR2**, and **SEA-SVR3** have a **Manageability** of **Online – Performance counters not started** before continuing.
2. Start **Windows PowerShell ISE**, and then open **C:\Labfiles\Mod04\Implement-StorageSpacesDirect.ps1**.
3. Run the commands in Step 1. The first command in Step 1 installs the Failover Clustering role service on **SEA-SVR1**, **SEA-SVR2**, and **SEA-SVR3**. The next command restarts the three servers, which is required to complete the install, and the third installs the Failover Cluster Manager console on **SEA-ADM1**. When you start the second command to restart the servers, you can go ahead and run the third command to install the console without waiting for the restarts to finish.

Validate cluster configuration and create a cluster

1. On **SEA-ADM1**, select the **Windows** key, and from the **Start** menu, select **Server Manager**.
2. In **Server Manager**, select **Tools**, and then select **Failover Cluster Manager**. This is to confirm it's installed.
3. In the **Administrator: Windows PowerShell ISE** window, select the line in step 2, starting with **Test-Cluster**, and then select the **F8** key. Wait until the installation finishes.

4. Verify that the output of the command only includes warnings and that the last line is a validation report in html format.
5. In the **Administrator: Windows PowerShell ISE** window, select the line in step 3 starting with **New-Cluster**, and then select **F8**. Wait until the installation finishes.
6. Verify that the output of the command only includes warnings, and that the last line has a **Name** column with the value **S2DCluster**.
7. Switch to the **Failover Cluster Manager** window.
8. Select **Connect to Cluster**, enter **S2DCluster**, and then select **OK**.

Enable the Storage Spaces Direct feature, create a storage pool, virtual disk, file server, and file share

1. In the **Administrator: Windows PowerShell ISE** window, select the line in step 4 starting with **Invoke-Command**, and then select **F8**. Wait until the installation finishes.
2. If a **Confirm** dialog box opens, select **Yes**.
3. There should be no output from the command, but a warning message might open.
4. In the **Administrator: Windows PowerShell ISE** window, select the line in step 5 starting with **Invoke-Command**, and then select **F8**. Wait until the installation finishes.
5. In the output of the command, verify that the **FriendlyName** attribute has a value of **S2DStorage-Pool**.
6. In the **Failover Cluster Manager** window, expand **S2DCluster.Contoso.com**, expand **Storage**, and then select **Pools**.
7. Verify the existence of **Cluster Pool 1**.
8. In the **Administrator: Windows PowerShell ISE** window, select the line in step 6 starting with **Invoke-Command**, and then select **F8**. Wait until the installation finishes.
9. Verify that in the output of the command has an attribute **FileSystemLabel**, with a value of **CSV**.
10. In the **Failover Cluster Manager** window, select **Disks**.
11. Verify the existence of **Cluster Virtual Disk (CSV)**.
12. In the **Administrator: Windows PowerShell ISE** window, select the line in step 7 starting with **Invoke-Command**, and then select **F8**. Wait until the installation finishes.
13. 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.
14. In the **Failover Cluster Manager** window, select **Roles**.
15. Verify that **S2D-SOFS** is in the list. This confirms that the command was successful.
16. In the **Administrator: Windows PowerShell ISE** window, select the three lines in step 8 starting with **Invoke-Command**, and then select **F8**. Wait until the installation finishes.
17. Verify that the command output has an attribute **Path** with a value of **C:\ClusterStorage\CSV\VM01**. This validates that the command was successful.
18. In the **Failover Cluster Manager** window, select **S2D-SOFS**, and then select the **Shares** tab.
19. Verify that **VM01** is included in the list. This verifies that the command was successful.

Test high availability for the storage

1. On **SEA-ADM1**, open **File Explorer**, and browse to **\s2d-sofs\VM01**.
2. Create a new folder named **VMFolder**, and then open it.
3. Switch to the **Administrator: Windows PowerShell ISE** window.
4. At the Windows PowerShell command prompt, enter the following command, and then select Enter:
`Stop-Computer -ComputerName SEA-SVR3`
5. Switch to the **Server Manager** window, and then select **All Servers**.
6. In the **Servers** list, select **SEA-SVR3**.
7. Verify that **Manageability** has changed to **Target computer not accessible**.
8. Switch back to the **File Explorer** window.
9. Create a new text document and save it in the **VMFolder**.
10. In **Failover Cluster Manager**, select **Disks**, and then select **Cluster Virtual Disk (CSV)**.
11. Verify that for the **Cluster Virtual Disk (CSV)**, the **Health Status** is **Warning**, and **Operational Status** is **Degraded**. (**Operational Status** might also display as **Incomplete**.)

Test your knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

What are the advantages of using Storage Spaces compared to using system area networks (SANs) or a network access server (NAS)?

Question 2

What are the disadvantages of using Storage Spaces compared to using SANs or NAS?

Implementing Data Deduplication

Lesson overview

Data Deduplication is a role service of Windows Server that 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 storage.

Lesson objectives

After completing this lesson, you'll be able to:

- Describe the Data Deduplication components.
- Describe how to deploy Data Deduplication.
- Identify Data Deduplication usage scenarios.
- Describe how to implement Data Deduplication.
- Describe the backup and restore considerations with Data Deduplication.

Data Deduplication components

To reduce disk utilization, Data Deduplication scans files, then divides those files into chunks, and retains only one copy of each chunk. After deduplication, files are no longer stored as independent streams of data. Instead, Data Deduplication replaces the files with stubs that point to data blocks that it stores in a common chunk store. The process of accessing deduplicated data is completely transparent to users and apps. You might find that Data Duplication increases overall disk performance. Multiple files can share one chunk cached in memory; therefore, that chunk is read from disk less often.

To avoid disk performance issues, Data Deduplication runs as a scheduled task rather than in real time. By default, optimization runs once per hour as a background task. However, depending on the configured usage type, the minimum file age might be three days.

The Data Deduplication role service consists of several components, including:

- **Filter driver.** This component monitors local or remote input/output (I/O) and manages the chunks of data on the file system by interacting with the various jobs. There is one filter driver for every volume.
- **Deduplication service.** This component manages the following job types:
 - Consisting of multiple jobs, they perform 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 then modified and meets the data deduplication policy threshold for optimization, the file will be optimized again.
 - Garbage collection. Data Deduplication includes garbage collection jobs to process deleted or modified data on the volume so that any data chunks no longer being referenced are cleaned up. This job processes previously deleted or logically overwritten optimized content to create usable volume free space. When an optimized file is deleted or overwritten by new data, the old data in the chunk store isn't deleted immediately. While garbage collection is scheduled to run weekly, you might consider running garbage collection only after large deletions have occurred.
 - 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 the most popular

data chunks. As data is accessed or deduplication jobs process data, if these features encounter corruption, they 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 the following three sources of redundant data:

- Backup copies. Deduplication keeps backup copies of popular chunks (chunks referenced over 100 times) in an area called the *hotspot*. If the working copy suffers a soft damage such as bit flips or torn writes, deduplication uses its redundant copy.
- Mirror image. If using mirrored Storage Spaces, deduplication can use the mirror image of the redundant chunk to serve the I/O and fix the corruption.
- New chunk. If a file is processed with a chunk that is corrupted, the corrupted chunk is eliminated, and the new incoming chunk is used to fix the corruption.

Note: Because of the additional validations that 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 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 doesn't support Data Deduplication. Before you start 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 it doesn't deduplicate the new data. Afterwards, you would use the unoptimization job to undo the existing deduplicated data on a volume. At the end of a successful de-optimization job, all the data deduplication metadata is deleted from the volume.

Note: Be cautious when using the de-optimization job because all the deduplicated data will return to the original logical file size. As such, you should verify the volume has enough free space for this activity, or you should move or delete some of the data to allow the job to complete successfully.

Data Deduplication process

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

- Processes the files on the volume by using optimization jobs, which are background tasks, run with low priority on the server.
- Uses an algorithm to segments all file data on the volume into small, variable-sized chunks that range from 32 kilobytes (KB) to 128 KB.
- Identifies chunks that have one or more duplicates on the volume.
- Inserts chunks into a common chunk store.
- Replaces all duplicate chunks with a reference (or *stub*) to a single copy of the chunk in the chunk store.
- Replaces the original files with a reparse point, which contains references to its data chunks.
- Compresses chunks and organizes them in container files in the **System Volume Information** folder.

- Removes 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. More information about this is discussed later in the module.

Data deduplication doesn't have any write-performance impact because the data isn't deduplicated while the file is being written. Windows Server uses post-process deduplication, which ensures that the deduplication potential is maximized. Another advantage to this type of deduplication process is that your application servers and client computers offload all processing, 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 can potentially process all the data on a selected volume, except for files that are less than 32 KB in size, and files in folders that are excluded. 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 will contain the following elements:

- **Unoptimized files.** Unoptimized files include:
 - Files that don't meet the selected file-age policy setting
 - System state files
 - Alternate data streams
 - Encrypted files
 - Files with extended attributes
 - Files smaller than 32 KB
 - Other reparse point files
- **Optimized files.** Optimized files includes files that are stored as reparse points, and that contain pointers to a map of the respective chunks in the chunk store that are needed to restore the file when it's requested.
- **Chunk store.** This is the location for the optimized file data.
- **Additional free space.** As a result of the data optimization, the optimized files and chunk store occupy much less space than they did prior to optimization.

Resilient File System (ReFS) now supports data deduplication in Windows Server 2019. It includes a new store that can contain up to ten times more data on the same volume when deduplication is applied. ReFS supports volumes up to 64 terabytes (TB), and deduplicates the first 4 TB of each file. It uses a variable-size chunk store that includes optional compression to maximize savings rates, while the multi-threaded post-processing architecture keeps performance impact minimal.

Deploy Data Deduplication

Preparation for Data Deduplication

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

Target deployments

Data Deduplication is designed to be applied on primary—and not to logically extended—data volumes without adding any more dedicated hardware. 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. You should consider using deduplication for the following data types:

- General file shares. These include group content publication and sharing, user home folders, and Folder Redirection/Offline Files.
- Software deployment shares. These are software binaries, images, and updates.
- VHD libraries. These are Virtual hard disk (VHD) file storage for provisioning to hypervisors.
- VDI deployments. These are Virtual Desktop Infrastructure (VDI) deployments using Microsoft Hyper-V.
- Virtualized backup. These include backup applications running as Hyper-V guests and saving backup data to mounted VHDs.

Deduplication can be extremely effective for optimizing storage and reducing the amount of disk space consumed, usually saving 50 to 90 percent of a system's storage space when applied to the correct data. Use the following questions 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 .vhd files tend to have plenty of duplication and yield higher storage savings from deduplication. More information on the deployment candidates for deduplication and the supported/unsupported scenarios are discussed later in this module.
- Does the data access pattern allow for enough time for deduplication?
Files that often change and are accessed often by users or applications aren't good candidates for deduplication. In these situations, 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 deduplication candidates allow time for deduplication of the files.
- Does the server have enough resources and time to run deduplication?

Deduplication requires reading, processing, and writing large amounts of data, which consumes server resources. Therefore, deduplication works more efficiently when it occurs outside of a server's busy times. A server that is constantly at maximum resource capacity might not be an ideal candidate for deduplication.

Evaluate savings with the Deduplication Evaluation Tool

You can use the Deduplication Evaluation Tool, **DDPEval.exe**, to determine the expected savings that you'll get if you enable deduplication on a particular volume. DDPEval.exe supports evaluating local drives and mapped or unmapped remote shares.

Note: When you install the deduplication feature, the Deduplication Evaluation Tool (DDPEval.exe) is automatically installed to the \Windows\System32\ directory.

Additional reading: For more information on planning to deploy Data Deduplication, refer to **Plan to Deploy Data Deduplication**¹.

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, you might need to 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 to process files earlier.
- You have directories that you don't want to deduplicate. Add a directory to the exclusion list.
- You have file types that you don't want to deduplicate. Add a file type to the exclusion list.
- The server has different off-peak hours than the default setting, and you want to change the Garbage Collection and Scrubbing schedules. Update the schedules using Windows PowerShell.

Installing and configuring Data Deduplication

After completing your planning, you can use the following steps to deploy Data Deduplication to a server in your environment:

1. Install Data Deduplication components on the server. To install deduplication components on the server, you can use one of the following options:
 - Server Manager. In Server Manager, you can install Data Deduplication by navigating to **Add Roles and Features Wizard**. Under **Server Roles**, select **File and Storage Services**, select the **File Services** check box, select the **Data Deduplication** check box, and then select **Install**.
 - Windows PowerShell. You can use the following Windows PowerShell command to install Data Deduplication:

```
Add-WindowsFeature -Name FS-Data-Deduplication
```
2. Enable Data Deduplication. Use the following options to enable Data Deduplication on the server:
 - Server Manager. From the Server Manager dashboard:
 1. Right-click or access the context menu for a data volume, and then select **Configure Data Deduplication**.
 2. In the **Data deduplication** box, select the workload 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 (VMs).
 3. Enter the minimum number of days that should elapse from the date of file creation before files are deduplicated.
 4. Enter the extensions of any file types that shouldn't be deduplicated.
 5. Select **Add** to browse to any folders with files that shouldn't be deduplicated.
 6. Select **Apply** to apply these settings and return to the Server Manager dashboard, or select the **Set Deduplication Schedule** button to continue to set up a schedule for deduplication.

¹ <https://aka.ms/sxzd2l>

- 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:

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

You can also use the Windows PowerShell cmdlet **Set-DedupVolume** to configure more 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 shouldn't be deduplicated.
 - The folders that should be excluded from deduplication.
3. Configure Data Deduplication jobs. You can run Data Deduplication jobs manually, on demand, or use a schedule. The following list are the types of jobs which you can perform on a volume:

- Optimization. Optimization includes built-in jobs that are scheduled automatically for optimizing the volumes on a periodic basis. Optimization jobs deduplicate data and compress file chunks on a volume per the policy settings. You can also 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 can also 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 processing-intensive operation, you should consider waiting until after the deletion load reaches a threshold to run this job on demand, or you should schedule the job for after hours. You can also 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 aren't scheduled automatically. However, you can use the following command to trigger an unoptimization job on demand:

```
Start-DedupJob -Volume _VolumeLetter_ -Type Unoptimization
```

Additional reading: For additional information on Set-DedupVolume, refer to **Set-DedupVolume²**.

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 access the schedules by using this Windows PowerShell cmdlet **Get-DedupSchedule**.

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

² <https://aka.ms/Set-DedupVolume>

the **Deduplication Settings** page in Server Manager, or by using the Windows PowerShell cmdlets: **New-DedupSchedule**, **Set-DedupSchedule**, or **Remove-DedupSchedule**.

Note: Data Deduplication jobs support—at most—weekly job schedules. If you need to create a schedule for a monthly job or for any other custom period, you'll need to use Windows Task Scheduler. However, you won't be able to use the **Get-DedupSchedule** cmdlet to access these custom job schedules that you create in Windows Task Scheduler.

Usage scenarios for Data Deduplication

Data Deduplication savings

Your Data Deduplication savings will vary by data type, the mix of data, the size of the volume, and the files that the volume contains. To evaluate the savings by volume, use the Deduplication Evaluation Tool before you enable deduplication.

The following list highlights typical deduplication savings for various content types.

- User documents. These include group content publication or sharing, user home folders (or MyDocs), and profile redirection for accessing offline files. Applying Data Deduplication to these shares might save upwards of 30 to 50 percent of your system's storage space.
- Software deployment shares. This includes software binaries, cab files, symbols files, images, and updates. Applying Data Deduplication to these shares could save 70 to 80 percent of your system's storage space.
- Virtualization libraries. This includes virtual hard disk files (i.e., .vhd and .vhdx files) storage for provisioning to hypervisors. Applying Data Deduplication to these libraries could save 80 to 95 percent of your system's storage space.
- General file share. This includes a mix of all the previously identified data types. Applying Data Deduplication to these shares might save 50 to 60 percent of your system's storage space.

Data Deduplication deployment candidates

Based on potentials savings and typical resource usage in Windows Server, deployment candidates for deduplication are ranked as follows:

- **Ideal candidates for deduplication**
 - Folder redirection servers
 - Virtualization depot or provisioning library
 - Software deployment shares
 - Microsoft SQL Server and Microsoft Exchange Server backup volumes
 - Scale-out File Servers (SOFS) Cluster Shared Volumes (CSVs)
 - Virtualized backup VHDs (for example, Microsoft System Center Data Protection Manager)
 - Virtualized Desktop Infrastructure VDI VHDs (only personal VDIs)

Note: In most VDI deployments, special planning is required for the *boot storm*, a name given to the phenomenon of large numbers of users trying to simultaneously sign in to their VDI, typically upon arriving to work in the morning. A boot storm hammers the VDI storage system and can cause long delays for VDI users. However, in Windows Server, when chunks read from the on-disk deduplication store during startup of a virtual machine (VM), they're cached in memory. As a result, subsequent

reads don't require frequent access to the chunk store because the cache intercepts them. This results in the minimization of boot storm effects because the memory is much faster than disk.

- **Should be evaluated based on content**
 - Line-of-business (LOB) servers
 - Static content providers
 - Web servers
 - High-performance computing (HPC)
- **Not ideal candidates for deduplication**
 - Microsoft Hyper-V hosts
 - Windows Server Update Service (WSUS)
 - SQL Server and Exchange Server database volumes

Data Deduplication interoperability

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

Windows BranchCache

You can optimize access to data over the network by enabling BranchCache on Windows Server and Windows client operating systems. When a BranchCache-enabled system communicates over a wide area network (WAN) with a remote file server that's enabled for Data Deduplication, all 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.

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, a copy of the content that is retrieved from the web server or file server is cached within the branch office. If another client in the branch requests the same content, the client can download it directly from the local branch network instead of again having to use the WAN to retrieve the content from the central office.

Failover Clusters

Windows Server fully supports failover clusters, which means deduplicated volumes will fail over gracefully between nodes in the cluster. Effectively, a deduplicated volume is a self-contained and portable unit (it has all the data and configuration information that the volume contains) but requires that each node in the cluster that accesses it must be running the Data Deduplication feature. This is because 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, the scheduled jobs will be applied on the next scheduled interval by the new node.

FSRM quotas

Although you shouldn't create a *hard* quota on a volume root folder enabled for deduplication, you can use File Server Resource Manager (FSRM) to create a *soft* quota on a volume root that's 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) doesn't 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.

Note:FSRM is a suite of tools for Windows Server that enables you to identify, control, and manage the type and quantity of data stored on your servers. FSRM enables you to configure hard or soft quotas on folders and volumes. A *hard quota* prevents users from saving files after the quota limit is reached; whereas, a *soft quota* doesn't enforce the quota limit, but generates a notification when the data on the volume reaches a threshold. When a hard quota is enabled on a volume root folder that's enabled for deduplication, the actual free space on the volume and the quota-restricted space on the volume aren't the same; this might cause deduplication optimization jobs to fail.

DFS Replication

Data Deduplication is compatible with Distributed File System (DFS) Replication. Optimizing or unoptimizing a file won't trigger a replication because the file doesn't change. DFS Replication uses remote differential compression (RDC) (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.

Demonstration: Implement Data Deduplication

In this demonstration, you'll learn how to:

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

Demonstration steps

To perform this demonstration, complete the following tasks.

Install the Data Deduplication role service

- On **LON-ADM1**, in **Server Manager**, add the **Data Deduplication** role.

Enable Data Deduplication

To enable Data Deduplication:

1. In **Server Manager**, select **File and Storage Services**, and then select **Disks**.
2. Select the **1** disk, and then select the **M** volume.
3. Enable Data Deduplication, and then select the **General purpose file server** setting.
4. Configure the following settings:
 - Deduplicate files older than (in days): **5**
 - **Enable throughput optimization**
5. In File Explorer share the **C:\Labfiles** folder.

6. On **SEA-SVR3**, map drive *X* to **\SEA-ADM1\Labfiles**.
7. On the **M** drive, make a directory named **Data**, and then enter the following command:
`copy x:\mod04\createlabfiles.cmd M:`
8. In the **Command** window, enter **CreateLabFiles.cmd**.
9. Do a directory listing and notice that **M:\Data** free space.

Check the status of Data Deduplication

To check the Data Deduplication status:

1. On **SEA-ADM1**, open **WAC**, connect to **SEA-SVR3**, and open the **PowerShell** node.
2. Execute the following commands to verify Data Deduplication status, selecting Enter after each command:
`Get-DedupStatus
Get-DedupStatus | fl
Get-DedupVolume
Get-DedupVolume | fl
Start-DedupJob M: -Type Optimization -Memory 50`
3. Repeat the commands **Get-DedupStatus** and **Get-DedupVolume**. Notice that the available free space increases on drive **M**.
4. Close all open windows.

Backup and restore considerations with Data Deduplication

One of the benefits of using Data Deduplication is that backup and restore operations are faster. This is because you have reduced the used space on a volume, meaning 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.

Backup and restore

Note: Many block-based backup systems should work with Data Deduplication, maintaining the optimization on the backup media. File-based backup operations that don't use deduplication usually copy the files in their original format.

With deduplication in Windows Server, you can back up and restore individual files and full volumes. You can also create optimized file-level backup/restore using Volume Shadow Copy Service (VSS) writer.

However, the backup or restore of only the reparse points and the backup or restore of only the chunk store aren't supported with deduplication in Windows Server.

A backup application can perform an incrementally optimized backup as follows:

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

Note: New chunks are appended to the current chunk store container. When the container's size reaches approximately 1 gigabyte (GB), that container file is sealed, and a new container file is created.

Restore operations

Restore operations can also benefit from Data Deduplication. Any file-level, full-volume restore operations can benefit because they're essentially a reverse of the backup procedure, and less data means quicker operations. The full volume restore process occurs in the following order:

1. The complete set of Data Deduplication metadata and container files are restored.
2. The complete set of Data Deduplication reparse points are restored.
3. All non-deduplicated files are restored.

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 third-party vendor, you should verify whether the backup solution supports Data Deduplication in Windows Server. Unsupported backup solutions should be avoided as they might introduce corruptions after a restore. Some common methods on solutions that support Data Deduplication in Windows Server are as follows:

- Some backup vendors support an *unoptimized backup*, which rehydrates the deduplicated files upon backup; in other words, it backs up the files as normal, full-size files.
- Some backup vendors support *optimized backup* for a full volume backup, which backs up the deduplicated files as-is; for example, as a reparse point stub with the chunk store.
- Some backup vendors support both.

The backup vendor should have comments on what their product supports, the method it uses, and with which version.

Test your knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

Can you configure data deduplication on a boot volume?

Question 2

Can I change the Data Deduplication settings for my selected usage type?

Question 3

Is Data Deduplication allowed on Resilient File System (ReFS)-formatted drives?

Implementing iSCSI

Lesson overview

Internet Small Computer Systems Interface (iSCSI) is a TCP/IP–based storage networking standard for connecting data storage services. It allows for block-level access to storage by transporting iSCSI commands over a network.

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. iSCSI meets the requirements for redundant remote storage connections. For this purpose, you'll learn how to create a connection between servers and iSCSI storage. You will also learn how to create both single and redundant connections to an iSCSI target by using the iSCSI initiator software that's available in Windows Server.

Lesson objectives

After completing this lesson, you'll be able to:

- Describe iSCSI.
- Describe the components of iSCSI.
- Identify the considerations for implementing iSCSI.
- Identify iSCSI usage scenarios.
- Describe how to configure and connect to an iSCSI target.

What is iSCSI?

Internet Small Computer System Interface (iSCSI) is a protocol that supports access to remote, small computer system interface (SCSI)–based storage devices over a TCP/IP network. iSCSI carries standard SCSI commands over IP networks to facilitate data transfers, and to manage storage over a network. You can use iSCSI to transmit data over local area networks (LANs), wide area networks (WANs), an intranet, or the internet.

iSCSI relies on standard Ethernet networking architecture; specialized hardware such as a host bus adapter (HBA) or network switches, is optional. iSCSI uses TCP/IP (typically, Transmission Control Protocol (TCP) port 3260). This means that iSCSI enables two hosts to negotiate (for example, session establishment, flow control, and packet size), and then exchange small computer system interface (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 storage area network (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 items:

- 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. Recommendations are that you use a dedicated physical and logical network to achieve faster, more 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 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 operating system versions include the iSCSI initiator, and they 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.

iSCSI components

This topic discusses the two main components of Internet Small Computer System Interface (iSCSI): the iSCSI Target Server, and the iSCSI initiator. We'll also learn about the Internet Storage Name Service (iSNS) and Data Center Bridging (DCB).

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 both iSCSI targets and iSCSI virtual disks, and then use Server Manager to manage these iSCSI targets and virtual disks. In Windows Server, the iSCSI Target Server is available as a role service under the File and Storage Services role in Server Manager.

The functionality of the iSCSI Target Server in Windows Server includes:

- Network or diskless boot. You can rapidly deploy diskless servers by using either boot-capable network adapters or a software loader. You can also 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 (VMs) that are running Microsoft Hyper-V or in high performance computing (HPC) clusters.
- Server application storage. Some applications such as Microsoft Exchange Server require block storage. The iSCSI Target Server can provide these applications with continuously available block storage. Because the storage is accessible remotely, it can also combine block storage for central or branch office locations.
- Heterogeneous storage. iSCSI Target Server supports iSCSI initiators that aren't based on Windows operating systems, so you can share storage on servers that are running Windows operating systems in mixed environments.
- Lab environments. The iSCSI Target Server role enables your Windows Server computer to be a network-accessible block-storage device. This is useful in situations when you want to test applications prior to deploying them on storage area network (SAN) storage.

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

- CHAP. You can enable Challenge Handshake Authentication Protocol (CHAP) to authenticate initiator connections, or you can enable reverse CHAP to enable the initiator to authenticate the iSCSI target.

- Query initiator computer for ID. This enables you to select an available Initiator ID from the list of cached IDs on the Target server. To use this, you must use a supported version of the Windows or Windows Server operating system.
- Virtual hard-disk support. You create iSCSI virtual disks as virtual hard disks. Windows Server supports both VHD and VHDX files. VHDX supports up to 64 terabyte (TB) capacity. You create new iSCSI virtual disks as VHDX files, but you can import VHD files as well.
- You manage the iSCSI Target Server by using either Server Manager or Windows PowerShell. Windows Server uses the Storage Management Initiative Specification provider with Microsoft System Center Virtual Machine Manager, and to manage an iSCSI Target Server in a hosted and private cloud.
- 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 **iSCSI Target Server Scalability Limits**³.

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.Micro-
soft:SQL1.Contoso.com"
Add-IscsiVirtualDiskTargetMapping SQLTarget E:\iSCSIVirtualHardDisk\1.vhdx
```

Additional reading: For more information on iSCSI Target cmdlets in Windows PowerShell, refer to **IscsiTarget**⁴.

When you enable the iSCSI Target Server to provide block storage, it capitalizes on your existing Ethernet network. You need either a dedicated network for iSCSI to ensure performance or Quality of Service (QoS) 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'll need shared storage for the cluster. This storage can be either hardware Fibre Channel storage or a serial-attached small computer system interface (SCSI) storage array. You configure the iSCSI Target Server as a cluster role in the failover cluster. You can also use Storage Spaces Direct for storing and providing highly available iSCSI targets.

iSCSI initiator

The iSCSI initiator is installed by default in all supported versions of Windows operating systems. To connect your computer to an iSCSI target, you only need to start the service and configure it.

The following Windows PowerShell cmdlets are 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”**

³ <https://aka.ms/iscsi-target-server-limits>

⁴ <https://aka.ms/iscsi-target>

iSNS

You can use the iSNS protocol when the iSCSI initiator attempts to discover iSCSI targets. The **iSNS Server service** feature in Windows Server 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:

- It contains a repository of active iSCSI nodes.
- It contains iSCSI nodes that can be initiators, targets, or management nodes.
- It allows initiators and targets to register with the iSNS server, and the initiators then query the iSNS server for the list of available targets.
- It contains a dynamic database of the iSCSI nodes. The database provides the iSCSI initiators with iSCSI target discovery functionality. The database updates 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 it enables iSNS to delete entries in the database that are no longer active.
- It 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.
- It provides Discovery Domain Service. You can divide iSCSI nodes into one or more groups called *discovery domains*, which provide zoning so that an iSCSI initiator can only refer 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 manage the traffic itself.

Windows Server supports Data Center Bridging by installing the **Data Center Bridging** feature. The advantage of using Data Center Bridging it can run all Ethernet traffic, including traffic going to and from your Fibre Channel or iSCSI SANs. This saves on datacenter cabling, network equipment, server space, and power. Data Center Bridging is also referred to as *Data Center Ethernet*, *Converged Enhanced Ethernet*, or *converged networking*.

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

Considerations for implementing iSCSI

Before embarking on an Internet Small Computer System Interface (iSCSI deployment), you should consider your infrastructure, staff, and customer requirements to ensure that you select the appropriate solution. The primary considerations that you should take into account are:

- Network speed and performance. The network speed should be at least 1 Gigabyte per second (Gbps) but in many cases, iSCSI networks in a datacenter now are 10 Gbps, 40 Gbps, or even 100 Gbps.

- High availability. The network infrastructure must be highly available because data is sent from the servers to the iSCSI storage over network devices and components.
- Security. The iSCSI solution should have an appropriate level of security. In situations where you need high security, you can use a dedicated network and with iSCSI authentication. In situations with lower security requirements, you might not have to use a dedicated network and with iSCSI authentication.
- Vendor information. Read the vendor-specific recommendations for different types of deployments and applications that use iSCSI storage, such as Microsoft Exchange Server and Microsoft SQL Server.
- Infrastructure staff. IT personnel who will design, configure, and administer the iSCSI storage must include IT administrators with different areas of specialization, such as Windows Server administrators, network administrators, storage administrators, and security administrators. This will help you design an iSCSI storage solution that has optimal performance and security. It also will help you create consistent management and operations procedures.
- Application teams. The design team for an iSCSI storage solution should include application-specific administrators, such as Exchange Server and SQL Server administrators, so that you can implement the optimal configuration for the specific technology or solution.

In addition to reviewing the infrastructure and teams, you also need to investigate competitive solutions to determine if they better meet your business requirements. Other alternatives to iSCSI include Fibre Channel, Fibre Channel over Ethernet, and InfiniBand.

iSCSI usage scenarios

In addition to configuring the basic Internet Small Computer System Interface (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 doesn't make that storage highly available. If iSCSI loses the connection, the server loses access to its storage. Therefore, most iSCSI storage connections are made redundant through one of two high availability technologies: Multiple Connected Session (MCS), and Multipath I/O (MPIO).

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

MCS is an iSCSI 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 storage area network (SAN) devices.

MPIO provides redundancy differently:

- 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) for when you want to connect to a third-party SAN device that's connected to the iSCSI initiator. The Windows operating system includes a default MPIO DSM that's installed as the MPIO feature within Server Manager.
- MPIO is widely supported. Many SANs can use the default DSM without any more software, while others require a specialized DSM from the manufacturer.
- MPIO is more complex to configure, and it isn't as fully automated during failover as MCS.

Demonstration: Configure and connect to an iSCSI target

In this demonstration, you'll learn how to:

- Add an Internet Small Computer System Interface (iSCSI) Target Server role service.
- Create iSCSI virtual disks and an iSCSI target.
- Connect to an iSCSI target.
- Verify the presence of the iSCSI drive.

Demonstration steps

Add iSCSI Target Server role services on SEA-ADM1 and SEA-SVR3 and prepare disks on SEA-SVR3

To add the iSCSI Target Server role and prepare disks:

1. On **SEA-ADM1**, open a Windows PowerShell window.
2. Enter the following command, and then select Enter:

```
Invoke-Command -ComputerName SEA-SVR3 -ScriptBlock {Install-WindowsFeature -Name FS-iSCSITarget-Server -IncludeManagementTools}
```
3. Open a remote **Windows PowerShell** session to **SEA-SVR3**.
4. Use Windows PowerShell to initialize the three offline disks on **SEA-SVR3**. Also, create a volume on the disks and format as Resilient File System (ReFS) using the following commands (where X refers to the drive number):

```
Initialize-Disk -Number XNew-Partition -DiskNumber 1 -UseMaximumSize -AssignDriveLetterFormat-Volume -DriveLetter _X_ -FileSystem ReFS
```
5. Use Windows PowerShell to create an inbound and outbound Firewall exception for port 3260:

```
New-NetFirewallRule -DisplayName 'iSCSITargetIn' -Profile 'Any' -Direction Inbound -Action Allow -Protocol TCP -LocalPort 3260New-NetFirewallRule -DisplayName 'iSCSITargetOut' -Profile 'Any' -Direction Outbound -Action Allow -Protocol TCP -LocalPort 3260
```
6. Close the remote session but keep Windows PowerShell open.

Create two iSCSI virtual disks and an iSCSI target on SEA-ADM1

To create the iSCSI virtual disks and iSCSI target:

1. On **SEA-ADM1**, in the **Server Manager** window, in **File and Storage Services**, under **Disks**, select the **SEA-DC1** server. Note that it only contains the boot and system volume on drive C.
2. In the **Server Manager** window, in **File and Storage Services**, under **iSCSI**, select the **SEA-SVR3** server.
3. Create a new iSCSI virtual disk with the following settings:
 - Storage Location: **E:**

- Name: **iSCSIDisk1**
 - Disk size: **5 Gigabyte (GB), Dynamically Expanding**
 - iSCSI target: **New**
 - Target name: **iSCSIFarm**
 - Access servers: **SEA-DC1** (Use **Browse** and **Check names**.)
4. Create a second iSCSI virtual disk with the following settings:
- Storage Location: **F:**
 - Name: **iSCSIDisk2**
 - Disk size: **5 GB, Dynamically Expanding**
 - iSCSI target: **iSCSIFarm**

Connect SEA-DC1 to the iSCSI target

To connect **SEA-DC1** to the iSCSI target:

1. On **SEA-DC1**, open Windows PowerShell, enter the following commands, and then select Enter:
`Start-Service msiscsiiscsicpl`

Note: The **iscsicpl** command will bring up an **iSCSI Initiator Properties** dialog box.

2. Connect to the following iSCSI target:

- Name: **SEA-DC1**
- Target name: **iqn.1991-05.com.microsoft:SEA-dc1-fileserver-target**

Verify the presence of the iSCSI disks

To verify the presence of the iSCSI disks:

1. In **Server Manager** on **SEA-ADM1**, in the tree pane, select **File and Storage Services**, and then select **Disks**.
2. Notice the new two **5 GB** disks on the **SEA-DC1** server that are offline. Notice that the bus entry is **iSCSI**. (If you're in the **File and Storage Services** section of Server Manager, you might need to select the refresh button to find the two new disks.)

Test your knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

What are the required components of an Internet Small Computer System Interface (iSCSI) solution? Select all that apply.

- IP network
- iSCSI targets
- iSCSI initiators
- iSCSI qualified name
- Domain Name System (DNS)

Question 2

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

- True
- False

Deploying DFS

Lesson overview

Providing files across multiple locations can be a challenging task. You must consider how to maintain easily accessible files and balance that access with file consistency between locations. You can use Distributed File System (DFS) to provide highly available, easily accessible files to branch offices. DFS performs wide area network (WAN)-friendly replication between multiple locations and can maintain consistency between file locations.

Lesson objectives

After completing this lesson, you'll be able to:

- Describe DFS.
- Understand how to deploy DFS.
- Describe how to implement DFS replication.
- Identify DFS namespace and replication.
- Describe how to manage a DFS database.
- Provide an overview of Microsoft Azure Files and storage utilities for the datacenter.

Overview of DFS

Distributed File System (DFS) functions provide the ability to logically group shares on multiple servers and to transparently link shares into a single hierarchical namespace. DFS organizes shared resources on a network in a tree-like structure. It has two components in its service: Location transparency via the namespace component, and redundancy via the file replication component. Together, these components improve data availability in the case of failure or heavy load by allowing shares in multiple locations to be logically grouped under one folder referred to as the *DFS root*.

You can implement DFS to provide the following efficiencies to different network file usage scenarios in branch offices:

- Sharing files across branch offices
- Data collection from branch offices
- Data distribution to branch offices

Sharing files across branch offices

Large organizations that have many branch offices often must share files or collaborate between these locations. DFS can help replicate files between branch offices or from a branch office to a hub site. Having files in multiple branch offices also benefits users who travel from one branch office to another. The changes that users make to their files in a branch office replicate to other branch offices.

Note: Recommend this scenario only if users can tolerate some file inconsistencies, because changes are replicated throughout the branch servers. Also note that DFS replicates a file only after it's closed. Therefore, DFS isn't recommended for replicating database files or any files that are kept open for prolonged periods.

Data collection from branch offices

DFS technologies can collect files from a branch office and replicate them to a hub site, thus allowing the files to be used for a number of specific purposes. Critical data can replicate to a hub site by using DFS, and then backed up at the hub site by using standard backup procedures. This increases data recoverability at the branch office if a server fails, because files will be available and backed up in two separate locations. Additionally, companies can reduce branch office costs by eliminating backup hardware and onsite IT personnel expertise. Replicated data can also be used to make branch office file shares fault tolerant. If the branch office server fails, clients in the branch office can access the replicated data at the hub site.

Data distribution to branch offices

You can use DFS to publish and replicate documents, software, and other line-of-business (LOB) data throughout your organization. DFS can also increase data availability and distribute client load across various file servers.

DFS-Replication (DFSR)

Distributed File System Replication (DFSR) is a technology that syncs data between two or more file shares. You can use it to make redundant copies of data available in a single location or in multiple locations. You can use DFSR across wide area network (WAN) links because it's very efficient. When a folder in a DFS Namespace has multiple targets, you should use DFSR to replicate the data between the targets.

DFSR is very efficient over networks because after an update it replicates only changes to files rather than entire files. DFSR uses the remote differential compression technology to identify only the changes to an updated file.

Deploy DFS

When implementing Distributed File System (DFS), you must have a general understanding of the overall topology of your DFS implementation. In general, DFS topology functions as follows:

1. The user accesses a folder in the virtual namespace. When a user attempts to access a folder in a namespace, the client computer contacts the server that's hosting the namespace root. The host server can be a standalone server that's hosting a standalone namespace, or a domain-based configuration that's stored in Active Directory Domain Services (AD DS) and then replicated to various locations to provide high availability. The namespace server sends back to the client computer a referral containing a list of servers that host the shared folders (called *folder targets*), and are associated with the folder being accessed. DFS is a site-aware technology, so to ensure the most reliable access, client computers are configured to access the namespaces that arrive within their site first.
2. The client computer accesses the first server in the referral. A *referral* is a list of targets that a client computer receives from a namespace server when the user accesses a root or folder with namespace targets. The client computer caches the referral information and then contacts the first server in the referral. This referral typically is a server within the client's own site unless no server is located within the client's site. In this case, the administrator can configure the target priority.

For example, the Marketing folder that's published within the namespace actually contains two folder targets: one share is located on a file server in New York, and the other share is located on a file server in London. The shared folders are kept synced by Distributed File System Replication (DFSR). Even though multiple servers host the source folders, this fact is transparent to users, who access only a single folder in

the namespace. If one of the target folders becomes unavailable, users are redirected to the remaining targets within the namespace.

Permissions required to create and manage a DFS namespace

To perform DFS namespace management tasks, a user either must be a member of an administrative group or has to be delegated specific permission to perform the task. To delegate the required permissions, right-click or access the context menu for the namespace, and then select **Delegate Management Permissions**.

The following table describes the groups that can perform DFS administration by default and the method for delegating the ability to perform DFS management tasks.

Table 1: The groups that can perform DFS administration by default and the method for delegating the ability to perform DFS management tasks

Task	Groups that can perform the task by default	Delegation method
Create a domain-based namespace	Domain admins	Select Delegate Management Permissions.
Add a namespace server to a domain-based namespace	Domain admins	Add users to local administrators group on the namespace server.
Manage a domain-based namespace	Local administrators on each namespace server	Select Delegate Management Permissions.
Create a standalone namespace	Local administrators on each namespace server	Add users to local administrators group on the namespace server.
Manage a standalone namespace	Local administrators on each namespace server	Select Delegate Management Permissions.
Create a replication group or enable DFSR on a folder	Domain admins	Add users to local administrators group on the namespace server.

Using Data Deduplication

Data Deduplication can help provide a more robust and efficient DFS environment when combined with DFSR:

- **Capacity optimization.** Data Deduplication enables a server to store more data in less physical disk space.
- **Scale and performance.** Data Deduplication is highly scalable in Windows Server. It can run on multiple volumes without affecting other services and applications running on the server. Data Deduplication can be throttled to accommodate other heavy workloads on the server so that no performance degradation occurs for important server tasks.
- **Reliability data integrity.** Windows Server uses checksum consistency and validation to ensure that the integrity of data affected by Data Deduplication remains intact. Data Deduplication also maintains redundant copies of the most frequently used data on a volume to protect against data corruption.
- **Bandwidth efficiency.** In combination with DFSR, Data Deduplication can greatly reduce the bandwidth consumed when replicating file data, if replication partners are also running Windows Server.

- Simple optimization management. Windows Server and Windows PowerShell contain integrated support for Data Deduplication. Implementation and management within Windows Server are accomplished with familiar tools.

When you want to configure Data Deduplication for use with DFS, you enable it on the volume or volumes that are hosted in the replicated DFS folders. You must enable Data Deduplication for volumes on all Windows Server-based computers that are participating in the DFSR topology.

Implement DFS Replication

Distributed File System Replication (DFSR) provides a way to keep folders synced between servers across well-connected and limited bandwidth connections. It primarily uses remote differential compression (RDC), or hidden staging folders.

DFSR and RDC

DFSR uses *remote differential compression (RDC)*, which is a client-server protocol that's used to efficiently update files over a limited bandwidth network. RDC detects data insertions, removals, and rearrangements in files, enabling DFSR to replicate only the changed file blocks when files update. By default, RDC is used only for files that are 64 kilobytes (KB) or larger.

DFSR also supports cross-file RDC, which allows DFSR to use RDC, even when a file with the same name doesn't exist at the client. Cross-file RDC can determine files that are similar to the file that needs to be replicated. It uses blocks of similar files that are identical to the replicating file to minimize the amount of data that needs to replicate.

DFSR and hidden staging folders

DFSR uses a hidden staging folder to stage a file before sending or receiving it. Staging folders function as caches for new and changed files to be replicated from sending members to receiving members. The sending member begins staging a file when it receives a request from the receiving member. The process involves reading the file from the replicated folder and building a compressed representation of the file in the staging folder. After it's constructed, the staged file is sent to the receiving member. (If RDC is used, only a fraction of the staging file might be replicated.) The receiving member downloads the data and builds the file in its staging folder. After the file download completes on the receiving member, DFSR decompresses the file and installs it into the replicated folder. Each replicated folder has its own staging folder, which by default is in the local path of the replicated folder in the **DfsrPrivate\Staging** folder.

- DFSR detects volume changes by monitoring the file system update sequence number (USN) journal and replicates changes only after the file is closed.
- DFSR uses a version vector exchange protocol to determine which files must be synced. The protocol sends less than 1 KB per file across the network to sync the metadata associated with changed files on the sending and receiving members.
- DFSR uses a conflict resolution heuristic of *last writer wins* for files that are in conflict (that is, a file that's updated on multiple servers simultaneously) and *earliest creator wins* for name conflicts. Files and folders that lose the conflict resolution are moved to a folder known as the *Conflict and Deleted* folder. If the file or folder is deleted, you can also configure the service to move deleted files to the **Conflict and Deleted** folder for retrieval. Each replicated folder has its own hidden **Conflict and Deleted** folder, which is in the local path of the replicated folder in the **DfsrPrivate\ConflictandDeleted** folder.

- DFSR is self-healing and can automatically recover from USN journal wraps, USN journal loss, or DFSR database loss.
- DFSR uses a Windows Management Instrumentation (WMI) provider that provides interfaces to obtain configuration and monitoring information from the DFSR service.

Deploying and configuring DFSR

After a Distributed File System (DFS) namespace and folder target are created, you can enable DFSR by configuring a replication group and enabling replication between group members.

Other DFSR functionalities

The DFSR functionality also includes:

- WMI provider. This provider enables the latest WMI-based methods for managing DFS.
- Database prestaging for initial sync. When prestaging DFSR data, you can bypass the initial replication phase when you create new replicated folders
- Database corruption recovery. This DFSR feature enables you to rebuild corrupt databases without data loss resulting from nonauthoritative initial sync.
- File staging tuning. You can configure variable file staging sizes on individual servers.

DFSR module for Windows PowerShell

You can configure and manage DFSR by using the cmdlets from the DFSR module for Windows PowerShell. The DFSR module supplies new cmdlets for managing all facets of DFSR functionality. For example, the following cmdlet creates a new replication folder named **Promotions** and adds it to the replication group named **Contoso-Marketing**:

```
New-DfsReplicatedFolder -GroupName "Contoso\_Marketing" -FolderName "Promotions"
```

This example retrieves the members of the Contoso_Marketing DFSR replication group on **SEA-SVR3**:

```
Get-DfsrMember -GroupName "Contoso\_Marketing" -ComputerName "SEA-SVR3"
```

To access all the available cmdlets in the DFSR module for Windows PowerShell, use the following Windows PowerShell cmdlet:

```
Get-Command -Module DFSR
```

DFS namespaces and replication

Each folder that you create in a Distributed File System (DFS) namespace has one or more targets. Each target is a file share containing files. If a folder has multiple targets, you should be using Distributed File System Replication (DFSR) to replicate the data between the file shares. This ensures that clients access the same data regardless of which target they use. In many cases, you'll want to optimize how clients access the targets to avoid replication conflicts.

Minimizing replication conflicts

DFSR doesn't replicate file locking to multiple copies of data. As a result, if there are two replicated copies of data, different users can open and modify each separate file copy. This creates a replication conflict.

The file most recently modified is retained and the losing file moves to the **ConflictAndDeleted** folder. There's no automated method for resolving replication conflicts. You must manually review the contents of both files.

One way to minimize replication conflicts is to set a single copy of the data with the highest priority for all clients. Then all clients access the same files and a file lock will be in place if two users try to edit the same file simultaneously. To make a target the highest priority, you should override referral ordering and select **First among all targets** in the target's properties.

When you use DFS across multiple locations and want users to access local copies of the data, you should segregate the data based on the location that uses it. For example, you would have separate shares for London and Vancouver. This allows you to set each local copy as having the highest priority.

To ensure that all clients use a consistent target after recovering from a failover, you should select the option **Clients fail back to preferred targets** on the folder. If you don't enable this option, some clients could continue to use a lower priority target even after recovery of the preferred target.

Speeding up distribution of changes

When clients retrieve information about folders and DFS roots, they cache that information to reduce network utilization. By default, DFS root information caches for five minutes, and information about targets caches for 30 minutes. It's unlikely that you'll need to reduce the cache for DFS roots, because five minutes is a short time. However, you might want to change the cache time used for a folder when you're making changes to the targets. You can set the cache time for a folder in its properties.

Clients get DFS information from namespace servers. You can optimize namespace servers for consistency or scalability. When you have less than 16 namespace servers, you should optimize for consistency. This way, the namespace servers all retrieve information from the primary domain controller (PDC) emulator for the domain. If you optimize for scalability, the namespace servers retrieve information from any available domain controller. In an environment with multiple Active Directory Domain Services (AD DS) sites, optimizing for scalability can result in significantly slower configuration syncing.

Hiding inaccessible folders

When clients access a namespace, they might not have permission to read data from all folders. To simplify user access, you should hide the inaccessible folders by using access-based enumeration. To hide a folder for a group:

1. Enable **Access-Based Enumeration** on the namespace.
2. Set **Deny view permissions** on the folder for the group.

Replication processing

Because DFSR syncs data between two or more folders on different servers, you can use it to make redundant copies of data available in a single location or in multiple locations. You can use DFSR across wide area network (WAN) links because it's very efficient. When a folder in a DFS Namespace has multiple targets, you should use DFSR to replicate the data between the targets.

Common scenarios for using DFSR include:

- Sharing files across branches.
- Data collection from branch offices.
- Data distribution to branch offices.

DFSR is quite efficient over networks because after an update it replicates only changes to files rather than entire files. DFSR uses the remote differential compression (RDC) technology to identify only the changes to an updated file.

During replication, DFSR uses a hidden staging folder as part of the replication process. The staging folder contains data that's either being replicated out to or in from replication partners. The path for the staging folder is **DfsrPrivate\Staging**.

When two users simultaneously modify a file on two replicas, there's a replication conflict. If there are replication conflicts for a file, the most recently modified file is kept. The earlier file copy moves to the **DfsrPrivate\ConflictAndDeleted** folder.

You can control replication for each replication group by setting the replication topology in the properties of the replication group. You can also set bandwidth limits and a replication schedule in the replication group properties.

There are three options for the replication topology:

- **Hub and spoke.** In this configuration, one replica is the central point with and through which all other replicas communicate. This is useful when your WAN also has a hub-and-spoke topology.
- **Full mesh.** In this configuration, all replicas communicate with all other replicas. This option is highly resilient because there's no central node that can fail. However, it becomes quite complex when there are more than a few replicas.
- **No topology.** In this configuration, you must manually create the connections between replicas.

Manage DFS databases

Distributed File System (DFS) includes database management tasks that use database cloning to help administrators perform initial database replication. DFS also includes tasks that can recover the DFS database in case of data corruption.

DFS database cloning

Windows Server 2019 provides a feature that clones the database for the initial replication. To create a clone of the database, you use the **Export-DfsrClone** cmdlet to export the Distributed File System Replication (DFSR) database and volume configuration XML file settings for a given local computer volume. On a large dataset, exports could take a long time to complete. You can use the **Get-DfsrCloneState** cmdlet to determine the status of the export operation.

After you clone the data and copy the exported database and XML file to the new DFS member server, you use the **Import-DfsrClone** cmdlet to inject the database onto a volume and validate the files on the file system. This provides dramatic performance improvements during the initial sync.

The following cmdlet exports a database and creates a clone of the database in a folder named **Dfsr-clone**:

```
Export-DfsrClone -Volume C: -Path "C:\Dfsrclone"
```

After copying the cloned database to the **C:\Dfsrclone** folder on the new DFS member server, use the following cmdlet to import the cloned database:

```
Import-DfsrClone -Volume C: -Path "C:\Dfsrclone"
```

DFS database recovery

When DFS Replication detects database corruption, it rebuilds the database and then resumes replication normally, with no files arbitrarily losing conflicts. When replicating with a read-only partner, DFS Replication resumes replication without waiting indefinitely for an administrator to set the primary flag manually. The database corruption recovery feature rebuilds the database by using local file and USN information, and then marks each file with a normal replicated state. You can't recover files from the **ConflictAndDeleted** and **Preexisting** folders except from backup. Use the Windows PowerShell cmdlets **Get-DfsrPreservedFiles** and **Restore-DfsrPreservedFiles** to allow the recovery of files from these folders. You can restore these files and folders into their previous location or to a new location. You can choose to move or copy the files, and you can keep all versions of a file or only the latest version.

Overview of Azure Files and storage utilities for the datacenter

You can extend on-premises deployments of Windows Server to the cloud by using Microsoft Azure hybrid services. These cloud services provide an array of useful functions, both for extending on-premises servers into Azure and for centrally managing servers from Azure. You can further extend your Windows Server 2019 storage features with various Azure storage-related services. Three important Azure hybrid services related to storage that you can use are:

- **Azure Files.** This service provides connectivity to share files on-premises and in the cloud. Azure Files include Azure file shares and file share snapshots.
- **Azure File Sync.** This service copies files from your on-premises Windows Server to an Azure file share, which enables you to centralize your file services in Azure while still having access to local data. File Sync also provides cloud tiering.
- **Storage Replica.** This Windows Server technology enables replication of volumes between servers or clusters for disaster recovery. Storage Replica provides support for synchronous and asynchronous replication in Azure.

Many other storage-related Azure services and technologies are available to Windows Server 2019 integration. To learn more about these technologies and services, refer to *Course WS-012T00 Windows Server 2019 Hybrid and Azure IaaS*.

Azure Files

Azure Files is a cloud service that anyone with internet access and appropriate permissions can access by default. Most users connect to Azure Files by using the Server Message Block (SMB) protocol, although Network File System (NFS) is also supported. SMB uses TCP port 445 for establishing connections. Many organizations and internet service providers block that port, which is a common reason why users can't access Azure Files. If unblocking port 455 isn't an option, you can still access Azure Files by first establishing a point-to-site virtual private network (VPN), a site-to-site VPN, or by using an Azure ExpressRoute connection to Azure. Alternatively, an organization can use File Sync to sync an Azure file share to an on-premises file server, which users can always access.

Azure Storage firewalls and virtual networks

Azure Storage, the platform which Azure Files is a part of, provides a layered security model. This model enables you to help secure and control the level of access to your storage accounts based on the type and subset of networks from where the request originates. By default, a storage account firewall allows access from all networks, but you can modify its configuration to allow access only from specified IP

addresses, IP ranges, or from a list of subnets in an Azure virtual network. Firewall configuration also enables you to select trusted Azure platform services to access a storage account securely.

In addition to the default public endpoint, storage accounts that include Azure Files provide the option to have one or more private endpoints. A *private endpoint* is an endpoint that's only accessible within an Azure virtual network. When you create a private endpoint for a storage account, the storage account gets a private IP address from within the address space of the virtual network, similar to how an on-premises file server or a network-attached storage device receives an IP address from the dedicated address space of an on-premises network. This secures all traffic between the virtual network and the storage account over a private link. You can also use the storage account firewall to block all access through a public endpoint when using private endpoints.

Additional reading: For more information on configuring connectivity, refer to **Azure Files networking considerations**⁵.

Connect to an Azure file share

To use an Azure file share with the Windows operating system, you must either mount it, which means assigning it a drive letter or mount point path, or access it through its Universal Naming Convention (UNC) path. The UNC path includes the Azure storage account name, the `file.core.windows.net` domain suffix, and the share name. For example, if the Azure storage account is named **storage1** and the share name is **share1**, the UNC path would be `\\storage1.file.core.windows.net\share1`.

If user-identity authentication is enabled for the storage account and you're connecting to an Azure file share from a domain-joined Windows device, you don't need to manually provide any credentials. Otherwise, you must provide credentials. You can use `(AZURE<storage account name>)` as the username and the storage access key as the password. The same credentials are used if you connect to an Azure file share by using the script that the Azure portal provides.

Note: Be aware that a storage access key provides unrestricted access to an Azure file share and that user-identity authentication should be used instead if possible.

Additional reading: For more information on connecting to Azure Files, refer to **Troubleshoot Azure Files problems in Windows**⁶.

Azure file share snapshots

In Windows Server, you can create a shadow copy of a volume, which captures the volume state at that point in time. You can later access the shadow copy over a network by using the File Explorer **Previous Version** feature. Similar functionality is available with Azure file share snapshots. A share snapshot is a point-in-time, read-only copy of the Azure file share data.

You create a share snapshot at the file share level. You can then restore individual files in the Azure portal or in File Explorer, where you can also restore a whole share. You can have up to 200 snapshots per share, which enables you to restore files to different point-in-time versions. If you delete a share, all its snapshots are also deleted.

Share snapshots are incremental. Only the data that has changed after your most recent share snapshot is saved. This minimizes the time required to create the share snapshot and saves on storage costs.

⁵ <https://aka.ms/storage-files-networking-overview>

⁶ <https://aka.ms/storage-troubleshoot-file-connection>

When to use share snapshots

- Use share snapshots as protection against accidental deletions or unintended changes. A share snapshot contains point-in-time copies of the share's files. If share files are unintentionally modified, you can use share snapshots to review and restore previous versions of the files.
- Use share snapshots for general backup purposes. After you create a file share, you can periodically create a share snapshot. This enables you to maintain previous versions of data that can be used for future audit requirements or disaster recovery.

Azure File Sync

Using File Sync, administrators can continue to manage on-premises file servers while seamlessly accessing file share data and reducing overall administrative effort.

Use File Sync to centralize your organization's file shares in Azure Files while keeping the flexibility, performance, and compatibility of an on-premises file server. File Sync transforms Windows Server into a quick cache of your Azure file share. You can use any protocol that's available on Windows Server to access your data locally, including SMB, NFS, and FTP over SSL (FTPS). You can have as many caches as you need across the world.

Using File Sync has many uses and advantages:

- Lift and shift. This is the ability to move applications that require access between Azure and on-premises systems. It provides write access to the same data across Windows Server computers and Azure Files. This lets organizations with multiple offices share files with all offices.
- Branch offices. Branch offices must back up files, or you must set up a new server that will connect to Azure Storage.
- Backup and disaster recovery. After implementing File Sync, Azure Backup will back up your on-premises data. You can also restore file metadata immediately and recall data as needed for rapid disaster recovery.
- File archiving. Only recently accessed data is located on local servers. Unused data moves to Azure in what's called *cloud tiering*.

Cloud tiering

Cloud tiering is an optional feature of File Sync in which frequently accessed files cache locally on the server while all other files are tiered to Azure Files based on policy settings. When a file is tiered, the File Sync file system replaces the file locally with a pointer or reparse point. The reparse point represents a URL to the file in Azure Files. When a user opens a tiered file, File Sync seamlessly recalls the file data from Azure Files without the user knowing that the file is stored in Azure. Cloud tiering files will have dimmed icons with an offline **O** file attribute to let the user know the file is only in Azure.

Additional reading: For more information, refer to **Planning for an Azure File Sync deployment**⁷.

File Sync moves file data and metadata exclusively over HTTPS and requires port 443 to be open out-bound. Based on policies in your datacenter, branch, or region, further restricting traffic over port 443 to specific domains might be desired or required.

There's a lot to consider when syncing large amounts of files. For example, you might want to copy the server files to the Azure file share before you configure File Sync.

⁷ <https://aka.ms/storage-sync-files-planning>

Storage Replica

Storage Replica enables replication of volumes between servers or clusters for disaster recovery. These volumes can also replicate between on-premises servers or clusters and Azure Virtual Machines (VM) servers and clusters. It also enables you to create stretch failover clusters that span two sites, with all nodes staying in sync.

Storage Replica supports synchronous and asynchronous replication:

- *Synchronous replication* mirrors data within a low-latency network site with crash-consistent volumes to ensure zero data loss at the file system level during a failure.
- *Asynchronous replication* mirrors data across sites beyond metropolitan ranges over network links with higher latencies, but without a guarantee that both sites have identical copies of the data at the time of a failure.

Why use Storage Replica?

Storage Replica allows storage-agnostic replication between clusters or servers. You can choose synchronous or asynchronous replication depending on the network latency and distance between the servers. All Storage Replica scenarios require GUID partition table initialized (GPT-initialized) disks.

You can configure Storage Replica to replicate storage among servers and clusters, and clusters in different datacenters. Storage Replica supports three scenarios:

- Stretch cluster
- Server-to-server
- Cluster-to-cluster
 - You can configure cluster-to-cluster Storage Replicas for cross-regions or within the same region for applications in Azure.

If you're replicating between on-premises servers and Azure VMs, you must create a network link between the on-premises servers and the Azure VMs. To do so, use ExpressRoute, a site-to-site VPN gateway connection, or install VPN software in your Azure VMs to connect them with your on-premises network.

Storage Replica offers disaster recovery and preparedness capabilities in Windows Server. Storage Replica technology helps Windows Server offer zero data loss with the ability to synchronously protect data on different racks, floors, buildings, campuses, counties, and cities. After a disaster, all data exists elsewhere without any possibility of loss, and the same applies before a disaster. Storage Replica offers the ability to switch workloads to safe locations prior to catastrophes, again with no data loss.

Storage Replica can do one-to-one replication only. Storage Replica can't create a third replica, such as Hyper-V Replica or Distributed File System (DFS). Storage Replica isn't suited for branch-office scenarios, because network connections to branch offices tend to be highly latent, and with lower bandwidth, synchronous replication becomes difficult. Storage Replica replicates data on the block level. Other replication engines, such as Distributed File System Replication (DFSR), can't replicate open files. With Storage Replica, changes in these files will replicate without issues.

Storage Replica synchronous replication

You might have to invest in expensive networking equipment to ensure that your network can perform synchronous replication. Synchronous replication has the following workflow:

1. The application writes data to the storage.
2. Log data writes in the primary site and the data replicates to the remote site.
3. Log data writes at the remote site.
4. The remote site sends an acknowledgement.
5. The primary site acknowledges the application write.

Storage Replica asynchronous replication

Server-to-server is the only scenario that supports asynchronous replication. Asynchronous replication has the following workflow:

1. The application writes data to the storage.
2. Log data writes in the primary site.
3. The primary site acknowledges the application write.
4. Data replicates to the remote site.
5. Log data writes at the remote site.
6. The remote site acknowledges the application write.

What's new in Windows Server 2019 Storage Replica

There are a number of improvements to Storage Replica in this release:

- Use in Windows Server Standard Edition. You can now use Storage Replica with Windows Server Standard Edition and Datacenter Edition. Storage Replica on Windows Server Standard Edition has the following limitations:
 - Storage Replica replicates a single volume instead of an unlimited number of volumes.
 - Volumes can have a size of up to 2 terabytes (TB) instead of an unlimited size.
- Log performance improvements. In this edition, you'll find upgrades on how the Storage Replica log tracks replication, improving replication throughput and latency, especially on all-flash storage and Storage Spaces Direct clusters that replicate between each other.

To gain the increased performance, all members of the replication group must run Windows Server 2019.
- Test failover. You can now temporarily mount a snapshot of the replicated storage on a destination server for testing or backup purposes.
- Windows Admin Center support. Support for graphical management of replication is now available in Windows Admin Center via Server Manager. This includes server-to-server replication, cluster-to-cluster, and stretch cluster replication.
- Miscellaneous improvements. Storage Replica also contains the following improvements:
 - Alters asynchronous stretch cluster behaviors so that automatic failovers now occur
 - Multiple bug fixes

Test your knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

What kinds of Distributed File System (DFS) namespaces are there and how do you ensure their availability?

Question 2

Is DFS Replication compatible with Data Deduplication?

Question 3

Can you use the Volume Shadow Copy Service (VSS) with DFS Replication?

Question 4

Is DFS Replication cluster aware?

Module 04 lab and review

Lab: Implementing storage solutions in Windows Server

Scenario

At Contoso, Ltd, you need to implement the Storage Spaces feature on the Windows Server 2019 servers to simplify storage access and provide redundancy at the storage level. Management wants you to test Data Deduplication to save storage. They also want you to implement Internet Small Computer System Interface (iSCSI) storage to provide a simpler solution for deploying storage in the organization. Additionally, the organization is exploring options for making storage highly available and researching the requirements that it must meet for high availability. You want to test the feasibility of using highly available storage, specifically Storage Spaces Direct.

Objectives

After completing this lab, you'll be able to:

- Implement Data Deduplication.
- Configure Internet Small Computer System Interface iSCSI storage.
- Configure Storage Spaces.
- Implement Storage Spaces Direct.

Estimated time: 90 minutes

Module review

Use the following questions to check what you've learned in this module.

Question 1

You attach five 2-terabyte (TB) disks to your Windows Server 2012 computer. You want to simplify the process of managing the disks. In addition, you want to ensure that if one disk fails, the failed disk's data isn't lost. What feature can you implement to accomplish these goals?

Question 2

Your manager has asked you to consider using Data Deduplication within your storage architecture. In what scenarios are the Data Deduplication role service particularly useful?

Question 3

Can you use both local and shared storage with Storage Spaces Direct?

Answers

Question 1

What are the two disk types in Windows 10 Disk Management?

The two types of disks are basic and dynamic.

Question 2

What file system do you currently use on your file server and will you continue to use it?

Answers could vary. A common answer is NT File System (NTFS), because NTFS should be the basis for any file system used on a Windows Server operating system. If you use FAT32 or Extended FAT (exFAT), you should be able to support your decision, because these file systems don't support security access control lists (ACLs) on files and folders.

Question 3

If permissions on a file are inherited from a folder, can you modify them on a file?

No, you can't modify inherited permissions. You can modify them on the folder where they were set explicitly, and then your modified permissions will be inherited with a file. Conversely, you can disable inheritance on a file, select or convert inherited permissions to explicit permissions, and then modify explicit permissions on it.

Question 4

Can you set permissions only on files in NTFS volumes?

No. You can set permissions on folders and entire volumes, including the root folder. Permissions that you set on folders or volumes are inherited to all content on that volume or in that folder, by default. You can set permissions on NTFS volumes and on Resilient File System (ReFS) volumes.

Question 1

Can any user connect to any shared folder?

No. Only users with the appropriate permissions can connect to shared folders. You configure permissions on shared folders when you share a folder, and you can modify permissions.

Question 2

What could be a reason that a user can't open files on a share?

There can be many reasons why a user can't open files on a share, including network connectivity issues, authentication problems, and issues with share and file permissions.

Question 3

What is the main difference between sharing a folder by using "Network File and Folder Sharing" and by using the "Advanced Sharing" option?

If you share a folder by using "Network File and Folder Sharing", you can set share and file permissions in a single step. If you share a folder by using the "Advanced Sharing" option, you can set only share folder permissions. You can't modify file permissions by using the "Advanced Sharing" option in a single step.

Question 4

What could be a reason that a user doesn't have the "Always available offline" option when they right-click or access the context menu for a file in the share, but when they right-click or access the context menu for a file in another share, the "Always available offline" option is available?

The most probable reason for such behavior is that the share doesn't allow offline files, and it has been configured with the "No files or Programs from the shared folder are available offline" option.

Question 1

What are the advantages of using Storage Spaces compared to using system area networks (SANs) or a network access server (NAS)?

Storage Spaces provides an inexpensive way to manage storage on servers. With Storage Spaces, you don't need to buy specialized storage or network devices. You can attach almost any kind of disk to a server and manage all the disks on your server as a block. You can provide redundancy by configuring mirroring or parity on the disks. Storage Spaces also are easy to expand by adding more disks. By using Storage Spaces tiering, you can also optimize the use of fast and slow disks in your storage space.

Question 2

What are the disadvantages of using Storage Spaces compared to using SANs or NAS?

Most SAN and NAS devices provide many of the same features as Storage Spaces. These storage devices also provide redundancy, data tiering, and easier capacity expansion. Additionally, they improve performance by removing all the storage-related calculations from the server and performing these tasks on dedicated hardware devices. This means that NAS and SAN devices (SAN devices in particular), are likely to provide better performance than using Storage Spaces.

Question 1

Can you configure data deduplication on a boot volume?

No, you can't configure data deduplication on a boot volume. You can configure data deduplication only on volumes that aren't system or boot volumes.

Question 2

Can I change the Data Deduplication settings for my selected usage type?

Yes. Although Data Deduplication provides reasonable defaults for recommended workloads, you might still want to tweak Data Deduplication settings to get the most out of your storage. Additionally, other workloads will require some tweaking as well, to ensure that Data Deduplication doesn't interfere with the workload.

Question 3

Is Data Deduplication allowed on Resilient File System (ReFS)-formatted drives?

With Windows Server 2016, Data Deduplication wasn't available for ReFS, and only available for NTFS file system. Now, with Windows Server 2019, Data Deduplication is available for both ReFS and NTFS file systems.

Question 1

What are the required components of an Internet Small Computer System Interface (iSCSI) solution?

Select all that apply.

- IP network
- iSCSI targets
- iSCSI initiators
- iSCSI qualified name
- Domain Name System (DNS)

Explanation

*If you access the iSCSI target through IP addresses, DNS isn't a required part of an iSCSI solution. iSCSI has its own name service, *internet Storage Name Service (iSNS)*. DNS is required only if you want to use fully qualified domain names (FQDN) to access your iSCSI storage.*

Question 2

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

- True
- False

Explanation

You can configure the iSCSI Target Server by using Server Manager and Windows PowerShell. However, you can't configure the iSCSI initiator by using Server Manager; you can only configure the iSCSI initiator through its own interface, or through Windows PowerShell.

Question 1

What kinds of Distributed File System (DFS) namespaces are there and how do you ensure their availability?

There are two kinds of DFS Namespaces: Standalone, and domain-based. For standalone DFS namespaces, you ensure the availability of a standalone DFS root by creating it on the cluster storage of a clustered file server by using the Cluster Administrator snap-in. For domain-based DFS namespaces, you ensure the availability of domain-based DFS roots by creating multiple root targets on non-clustered file servers or on

the local storage of the nodes of server clusters. (Domain-based DFS roots can't be created on cluster storage.) All root targets must belong to the same domain. To create root targets, use the DFS snap-in or the Dfsutil.exe command-line tool.

Question 2

Is DFS Replication compatible with Data Deduplication?

Yes, DFS Replication can replicate folders on volumes that use Data Deduplication in Windows Server.

Question 3

Can you use the Volume Shadow Copy Service (VSS) with DFS Replication?

Yes. DFS Replication is supported on VSS volumes, and you can restore previous snapshots successfully with the previous version's client.

Question 4

Is DFS Replication cluster aware?

Yes, DFS Replication is cluster aware. DFS Replication in Windows Server 2008 R2 through Windows Server 2019 includes the ability to add a failover cluster as a member of a replication group.

Question 1

You attach five 2-terabyte (TB) disks to your Windows Server 2012 computer. You want to simplify the process of managing the disks. In addition, you want to ensure that if one disk fails, the failed disk's data isn't lost. What feature can you implement to accomplish these goals?

You can use Storage Spaces to create a storage pool of all five disks, and then create a virtual disk with parity or mirroring to make it highly available.

Question 2

Your manager has asked you to consider using Data Deduplication within your storage architecture. In what scenarios are the Data Deduplication role service particularly useful?

You should consider using deduplication for file shares, software deployment shares, and VHD and VHDX file libraries. For file shares, include group content publication or sharing, user home folders, and profile redirection for accessing offline files. With the release to manufacturing (RTM) version of Windows Server 2012, you could save approximately 30 to 50 percent of your system's disk space. With the Cluster Shared Volume (CSV) support in Windows Server 2012 R2, the disk savings can increase up to 90 percent in certain scenarios. Software deployment shares include software binaries, images, and updates. You might be able to save approximately 70 to 80 percent of your disk space. VHD and VHDX file libraries include VHD and VHDX file storage for provisioning to hypervisors. You might be able to save disk space of approximately 80 to 95 percent.

Question 3

Can you use both local and shared storage with Storage Spaces Direct?

No. Storage Spaces Direct can use only local storage. A standard storage space can use shared storage.



Module 5 Hyper-V virtualization and containers in Windows Server

Hyper-V in Windows Server

Lesson overview

The Hyper-V server role provides you with a virtualized computing environment to create, configure, and manage virtual machines (VMs). To effectively plan and implement a VM environment, it is important to understand key features provided by Hyper-V in Windows Server.

In this lesson, you learn how to use Hyper-V to implement virtualization. You also learn best practices for configuring Windows server hosts and considerations related to deployment scenarios such as nested virtualization. Finally, you learn considerations, requirements, and processes for migrating on-premises Hyper-V VMs to Microsoft Azure.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe how Hyper-V provides virtualization capabilities to Windows Server.
- Use Hyper-V Manager to manage virtualization.
- Identify best practices for configuring Hyper-V hosts.
- Explain nested virtualization.
- Describe how to migrate on-premises Hyper-V VMs to Azure.

Overview of Hyper-V

The Hyper-V server role in Windows Server provides virtualization capabilities to support a virtual network environment. Hyper-V allows you to subdivide the hardware capacity of a single physical host computer and allocate the capacity to multiple virtual machines (VMs). Each VM has its own operating system that runs independently of the Hyper-V host and other VMs.

When you install the Hyper-V server role, a software layer known as the *hypervisor* is inserted into the boot process. The hypervisor is responsible for controlling access to the physical hardware. Hardware drivers are installed only in the host operating system (also known as the *parent partition*). All the VMs communicate only with virtualized hardware.

The operating systems running in VMs are referred to as guest operating systems. Hyper-V in Windows Server 2019 supports the following guest operating systems:

- All supported Windows versions
- Linux editions: CentOS, Red Hat Enterprise Linux, Debian, Oracle Linux, SUSE, and Ubuntu
- FreeBSD

Note: Hyper-V is also available as a feature in some 64-bit versions of Windows and as a downloadable, standalone server product called *Microsoft Hyper-V Server*. Although most features are the same, this module focuses on Hyper-V installed as a server role on Windows Server 2019.

Why use Hyper-V

Hyper-V is used to support various scenarios from simple VMs to complex software-defined infrastructures. You can use Hyper-V to:

- **Consolidate your server infrastructure.** Hyper-V can provide the foundation to consolidate multiple physical servers on to fewer, more powerful computers to use less space and consume less energy.
- **Provide a virtual development or test environment.** Virtualization provides the means to reproduce various development or test environments without having to purchase or maintain physical hardware or isolated network systems. Virtualized development or test environments can be quickly configured and reverted as needed without affecting production systems.
- **Establish a virtual desktop infrastructure (VDI):** Combining Hyper-V and Remote Desktop Virtualization with Windows Server can provide a centralized desktop management solution using VDI. This scenario can help you provide secure and agile personal virtual desktops or virtual desktop pools to your users.
- **Implement a private cloud infrastructure.** You can use Hyper-V as a foundation to provide flexible, on-demand services that function much like public cloud services. Azure Stack hyperconverged infrastructure (HCI) is an example of how Hyper-V can be integrated with other technologies such as Storage Spaces Direct (S2D) and Software Defined Networking (SDN) to run virtualized workloads on-premises.

General features of Hyper-V on Windows Server

In new releases of Windows Server, Hyper-V is often updated with new features and functionality. These new features typically provide options for supporting new workloads, increasing performance, and enhancing security. General features can be grouped as follows:

- **Management and connectivity.** You can manage your Hyper-V environment using the Hyper-V Manager, Hyper-V module for Windows PowerShell, Virtual Machine Connection (also referred to as VMConnect), and Windows PowerShell Direct. These tools can be installed on the same computer on which the Hyper-V server role has been installed, or you can install the tools on a remote management computer.
- **Portability.** To make it easier to move or distribute a VM, Hyper-V provides features such as live migration, storage migration, and standard import/export functionality.

- **Disaster recovery and backup.** Hyper-V supports Hyper-V Replica, which creates copies of VMs in another physical location. These copies can be used to restore VM instances as needed. Other features such as Production checkpoints and support for Volume Shadow Copy Service (VSS) both provide the ability to make application-consistent backups of the state of a VM.
- **Security:** Hyper-V supports security features such as Secure boot and shielded VMs. Secure boot verifies digital signatures on files during the boot process to prevent against malware. Shielded VMs help to secure access to VMs by encrypting the files and only allowing the VM to be run from specific protected virtualization host machines.
- **Optimization.** Hyper-V includes a set of customized services and drivers called Integration Services. These services are available for all supported guest operating systems, which include Time Synchronization, Operating System Shutdown, Data Exchange, Heartbeat, Backup, Guest Services, and Power-Shell Direct. Updates for Integration Services are obtained and delivered through Windows Update.

System requirements for Hyper-V on Windows Server

To support VMs running production applications or services, you must carefully assess the capacity required for your VMs and plan your Hyper-V hosts accordingly. You also need to consider needs such as high availability. The following are some basic hardware requirements for a Hyper-V host:

- A 64-bit processor with second-level address translation (SLAT)
- A processor with VM Monitor Mode extensions
- Sufficient memory for itself and for guest VMs
- Intel Virtualization Technology (Intel VT) or AMD Virtualization (AMD-V) enabled
- Hardware-enforced Data Execution Prevention (DEP) enabled (Intel XD bit, AMD NX bit)

To verify that a system meets the requirements for Hyper-V, you can run **Systeminfo.exe**. The output has a Hyper-V section that identifies whether the requirements are met.

In addition to the hardware requirements for Hyper-V, you should ensure that the Hyper-V hosts have sufficient hardware resources for the VMs. The following is a list of necessary resources:

- **Processor.** Ensure that there are enough physical processor cores to support the VMs you plan to run.
- **Memory.** Ensure that there is enough memory in the Hyper-V host to support the number of VMs that you intend to run.
- **Storage.** Ensure that you have enough storage for the VHDs used by your VMs. Also, ensure that the storage subsystem has high throughput to support multiple VMs accessing the storage simultaneously.
- **Network.** Ensure that there is enough network capacity in the Hyper-V host to allocate to the VMs. In some cases, you might need to allocate network adapters in the host for different purposes.

Methods used for installing the Hyper-V Server role on Windows Server

To install the Hyper-V server role, you can use **Server Manager** or the **Install-WindowsFeature** cmdlet in Windows PowerShell.

Tip: You can also use the **Windows Admin Center** to install the Hyper-V Server role on remote servers.

Install Hyper-V by using Server Manager

1. In **Server Manager**, select **Add Roles and Features**.
2. On the **Select installation type** page, select **Role-based or feature-based installation**.
3. On the **Select destination server** page, select the intended server from the server list.
4. On the **Select server roles** page, select **Hyper-V**. Select **Add Features** when prompted.
5. On the **Create Virtual Switches** page, **Virtual Machine Migration** page, and **Default Stores** page, select the appropriate options.
6. On the **Confirm installation selections** page, select **Restart the destination server automatically if required**, and then select **Install**. The server will restart as required.

Install Hyper-V by using the Install-WindowsFeature cmdlet

1. Open Windows PowerShell with administrator credentials.
2. Enter the following command and replace `<computer_name>` with the name of the server.

```
Install-WindowsFeature -Name Hyper-V -ComputerName <computer_name> -IncludeManagementTools -Restart
```

3. After the server restarts, you can verify that Hyper-V has been installed by entering the following command:

```
Get-WindowsFeature -ComputerName <computer_name>
```

Note: If you install the Hyper-V server role on a server that contains the Server Core installation option, the `-IncludeManagementTools` parameter will only install the **Hyper-V Module for Windows PowerShell**. You can still use the Hyper-V Manager from another computer to remotely manage a Hyper-V host that is installed with a Server Core installation.

Overview of Hyper-V Manager

The Hyper-V Manager is often used as a graphical user interface (GUI) to manage both local and remote Hyper-V hosts. Hyper-V Manager is available when you install the Hyper-V Management Tools, which are included with a full Hyper-V server role installation or installed as a tools-only installation.

Hyper-V Manager supports the following general features:

- **Previous version support.** When using Hyper-V Manager on Windows Server 2019 or Windows 10, you can still manage hosts installed with previous operating systems such as Windows Server 2016, Windows Server 2012, or Windows Server 2012 R2.
- **Support for WS-Management protocol.** Hyper-V Manager supports connections to Hyper-V hosts over the Web Services Management Protocol (WS-Management Protocol). This allows Hyper-V Manager to communicate by using the Kerberos protocol NTLM or Credential Security Support Provider (CredSSP). When using CredSSP, 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 either over port 80 or port 443, the default open ports.
- **Alternate credential support.** Communicating over the WS-Management Protocol makes it possible to use a different set of credentials in Hyper-V Manager and to save the credentials for ease of management. However, alternative credentials only work with Windows 10 and Windows Server 2016 or later hosts. Older servers installed with Hyper-V do not support the WS-Management Protocol for Hyper-V Manager communication.

Other methods for managing Hyper-V on Windows Server

Hyper-V Manager is the most common interface for managing virtual machines (VMs) in Hyper-V. You might also choose to use other tools that provide similar features for specific management scenarios. These tools include:

- **Windows PowerShell**

- The Hyper-V module for Windows PowerShell provides PowerShell cmdlets that can be used for scripting or command-line administrative scenarios.
- You can use PowerShell to manage the configuration, display the status, and perform general management tasks for Hyper-V hosts and their guest VMs.

- **PowerShell Direct**

- PowerShell Direct allows you to use Windows PowerShell inside a VM regardless of the network configuration or remote management settings on either the Hyper-V host machine or the VM.
- You can use the **Enter-PSSession** cmdlet to connect to a VM, which then allows you to perform PowerShell cmdlets directly on the VM with which you created the session.
- To use PowerShell Direct, the VM must be started and you must be signed on to the host computer as a Hyper-V administrator. The host operating system and the target VM must be installed with at least Windows 10 or Windows Server 2016.

- **Windows Admin Center**

- The Windows Admin Center is a browser-based application used for remotely managing Windows Servers, clusters, and Windows 10 PCs.
- When you connect to a Hyper-V host using the Windows Admin Center, you can manage VMs and virtual switches with functionality similar to the Hyper-V Manager.
- Windows Admin Center also provides Summary and Status information on events, CPU utilization, and Memory usage.

Best practices for configuring Hyper-V hosts

You should consider the following best practices when provisioning Windows Server to function as a Hyper-V host:

- Provision the host with adequate hardware.
- Deploy virtual machines (VMs) on separate disks, solid state drives (SSDs), or Cluster Shared Volumes (CSVs) if using shared storage.
- Do not collocate other server roles.
- Manage Hyper-V remotely.
- Run Hyper-V by using a Server Core configuration.
- Run the Best Practices Analyzer and resource metering.
- Use generation 2 VMs if the guest operating system supports them.

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 configured as a team. Inadequately provisioning the Hyper-V host with hardware affects the performance of all VMs that the server hosts.

Deploy VMs on separate disks

You should use separate disks to host VM files rather than store VM files on the same disk as the host operating system files. Doing this minimizes contention and ensures that read/write operations that occur on VM files do not conflict with read/write operations that occur at the host operating system level.

It also minimizes the chance that VM hard disks will grow to consume all available space on an operating system volume. Using SSDs, which have much faster read-write speed and consume less power than standard hard disk drives, is also recommended.

You can reduce the effect on performance if you deploy to a disk that uses striping, such as a RAID 1+0 array. When using shared storage, you can provision multiple VMs on the same logical unit number (LUN) if you use CSVs. However, choosing between separate LUNs for each VM or a shared LUN depends heavily on VM workload and host configuration.

Do not collocate other server roles

You should ensure that Hyper-V is the only server role installed on the host server. Do not collocate the Hyper-V role with other roles such as the domain controller or the file server role. Each role that you deploy on a server requires resources, and when deploying Hyper-V, you want to ensure that VMs 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 VMs 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 VMs. You should also restrict access to the Hyper-V server so that only administrators who are responsible for VM management can make connections. A configuration error on a Hyper-V host can cause downtime for all guest VMs.

Run Hyper-V by using a Server Core configuration

You should manage Hyper-V by using the Server Core configuration. Doing so provides the following benefits:

- The Server Core configuration of Windows Server minimizes hardware-resource utilization for the host operating system. As a result, the hosted VMs have more hardware resources.
- The Server Core requires fewer software updates, which in turn require fewer restarts.
- When you restart a Hyper-V host and it is unavailable, all VMs that the server hosts also become unavailable. To prevent this issue, consider using a clustered environment. Because a Hyper-V host can host many critical servers as VMs, 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 also use resource metering to monitor how hosted VMs use server resources and to determine if specific VMs are using a disproportionate amount of a host server's resources. If the performance characteristics of one VM erode the performance of other VMs that are hosted on the same server, consider migrating that VM to another Hyper-V host.

Use generation 2 VMs if supported by the guest operating system

Generation 2 VMs have slightly faster start times than generation 1 VMs. Generation 2 VMs use a simplified hardware model and allow advanced features such as:

- Pre-Boot Execution Environment (PXE) boot from a standard network adapter
- Hot add/removal of virtual network adapters
- SCSI controller boot
- Secure boot
- Shielded VMs
- Support for Storage spaces direct

Overview of nested virtualization

Nested virtualization is a Hyper-V feature that allows you to install and run Hyper-V inside a guest virtual machine (VM). Nested virtualization enables a guest VM to be a Hyper-V host, which can then host other guest VMs. This can be extremely useful for implementing testing scenarios that previously would have required more physical hardware to run.

To enable nested virtualization, you need to meet the following prerequisites:

- Both the Hyper-V host and the guest VM must be Windows Server 2016 or later.
- A sufficient amount of static RAM.
- Hyper-V VMs must have a configuration version of 8.0 or greater.
- The physical host computer must have an Intel processor with VT-x and EPT technology.

To enable nested virtualization, while the VM is in the OFF state, run the following command from the physical Hyper-V host machine:

```
Set-VMProcessor -VMName <VMName> -ExposeVirtualizationExtensions $true
```

To enable network packets to be routed through two virtual switches, MAC address spoofing must be enabled on the physical Hyper-V host. To enable MAC address spoofing, run the following command from the physical Hyper-V host machine:

```
Set-VMNetworkAdapter -VMName <VMName> | Set-VMNetworkAdapter -MacAddressSpoofing On
```

After enabling nested virtualization, you can install Hyper-V on the guest VM 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:

- Virtualization Based Security (VBS) cannot expose virtualization extensions to guests. You must first disable VBS before enabling nested virtualization.
- Device Guard cannot expose virtualization extensions to guests. You must first disable Device Guard on the host before enabling nested virtualization.
- Dynamic Memory is not supported, and runtime memory resize will fail.

Migration to Azure VMs

Many organizations decide to move some or all of their server infrastructure to cloud-based platforms such as Azure. These organizations realize the benefits and take advantage of the decreased cost of infrastructure maintenance, increased scalability, and high availability that a cloud-based infrastructure provides.

You can discover, assess, and migrate many of your on-premises workloads, apps, and VMs to Azure by using the **Azure Migrate** service.

Azure Migrate is a service included within Microsoft Azure that provides the following benefits:

- **A single migration platform.** Azure Migrate provides a single portal used to start, run, and track your migration to Azure.
- **Assessment and migration tools.** You have several tools to assist in your migration tasks, including **Azure Migrate: Server Assessment** and **Azure Migrate: Server Migration**.
- **Assess and migrate multiple object types.** The Azure Migrate hub portal allows you to assess and migrate:
 - Servers
 - Databases
 - Web applications
 - Virtual desktops
 - Data

How does Hyper-V migration to Azure work

The **Azure Migrate: Server Migration** tool is used to provide agentless replication to Azure for on-premises Hyper-V VMs. Software agent components are only installed on the Hyper-V hosts or cluster nodes; however, no agents are required to be installed on the Hyper-V guest VMs.

The **Azure Migrate: Server Migration** tool shares common technology with the **Microsoft Azure Site Recovery** tool, which had been the primary solution for replicating VMs to Azure before the Azure Migrate service was released. The components used for replication include:

- **Replication provider.** Installed on Hyper-V hosts and registered with Azure Migration Server Migration. Used to orchestrate replication for Hyper-V VMs.
- **Recovery Services agent.** Works with the provider to replicate data from Hyper-V VMs to Azure. Replicated data is migrated to a storage account in your Azure subscription. The Server Migration tool then processes the replicated data and applies the data to replica disks used to create the Azure VMs during the migration.

Test your knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

What is the correct term for the virtualization layer that is inserted into the boot process of the host machine that controls access to the physical hardware?

Question 2

Name four methods for managing Hyper-V virtual machines.

Question 3

What is the PowerShell command for enabling nested virtualization?

Configuring VMs

Lesson overview

After installing the Hyper-V server role on your host server, your next step is to configure the intended virtual infrastructure. The virtual infrastructure consists of several components and tasks, including the configuration of virtual networks, creating virtual disks, and creating and managing virtual machines (VMs) containing supported operating systems.

In this lesson, you learn the concepts related to VM configurations and generation versions. You also learn VM settings, storage options, and virtual disk types. Finally, you learn about the types of virtual networks and how to create and manage a VM.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe VM configuration and generation versions.
- Explain VM settings.
- Describe storage options for Hyper-V.
- Identify virtual hard disk (VHD) formats and types.
- Describe shared VHDX and Hyper-V VHD sets.
- Explain Hyper-V networking.
- Describe the types of virtual networks available for Hyper-V.
- Manage VM states and checkpoints.
- Import and export VMs.
- Create and manage a VM.

VM configuration and generation versions

Your organization might support multiple Hyper-V host machines that contain various Windows or Windows Server versions or semi-annual channel releases. To ensure that you can easily move and start virtual machines (VMs) between Hyper-V hosts, it is important to understand and identify the VM configuration versions used in your virtual environment.

VM configuration versions

A VM configuration version identifies the compatibility of VM components with the version of Hyper-V installed on the host machine. These components include the following:

- **Configuration.** The VM configuration information such as processor, memory, attached storage, and so on.
- **Runtime state.** The runtime state of the VM such as Off, Starting, Running, Stopping, and so on.
- **Virtual hard disk (VHD).** VHD or VHDX files that represent the virtual hard disks VHDs attached to the VM.
- **Automatic virtual hard disk.** Differencing disk files used for VM checkpoints.

- **Checkpoint.** Files representing configuration files and runtime state files used when checkpoints are created.

From the Hyper-V Manager console, you can display the configuration version of a specific VM by referring to the **Configuration Version** entry displayed on the **Summary** tab.

You can also use the following PowerShell cmdlet to get the versions of the VMs stored on the host machine:

```
Get-VM * | Format-Table Name, Version
```

To identify the VM configuration versions your Hyper-V host supports, run the following PowerShell cmdlet:

```
Get-VMHostSupportedVersion
```

When you create a new VM on a Hyper-V host, a default configuration version is used. To determine the default version for your Hyper-V host, run the following PowerShell cmdlet:

```
Get-VMHostSupportedVersion -Default
```

Supported VM configuration versions

If you have VMs created with an earlier version of Hyper-V, some features that are introduced and available on a newer Hyper-V host operating system might not work with those VMs. It is important to understand which operating system versions support which VM configuration versions.

The following table displays the VM configuration versions that are supported on hosts running various versions of the Windows operating system.

Hyper-V host Windows version	Configuration versions
Windows Server 2019	9.0, 8.3, 8.2, 8.1, 8.0, 7.1, 7.0, 6.2, 5.0
Windows 10 Enterprise LTSC 2019	9.0, 8.3, 8.2, 8.1, 8.0, 7.1, 7.0, 6.2, 5.0
Windows Server 2016	8.0, 7.1, 7.0, 6.2, 5.0
Windows 10 Enterprise 2016 LTSC	8.0, 7.1, 7.0, 6.2, 5.0
Windows 10 Enterprise 2015 LTSC	6.2, 5.0
Windows Server 2012 R2	5.0
Windows 8.1	5.0

Note: Hosts that run semi-annual channel versions of Windows, such as Windows Server 1903 and Windows 10 1903 or later, contain and support Hyper-V configuration version 9.0 and 9.1.

Version requirement for Hyper-V features

Feature	Minimum configuration version
Hot Add/Remove Memory	6.2
Production Checkpoints	6.2
PowerShell Direct	6.2
Virtual Machine Grouping	6.2
Secure Boot for Linux VMs	6.2
Virtual Trusted Platform Module (vTPM)	7.0
Virtual machine multi queues (VMMQ)	7.1
XSAVE support	8.0

Feature	Minimum configuration version
Key storage drive	8.0
Guest virtualization-based security support (VBS)	8.0
Nested virtualization	8.0
Virtual processor count	8.0
Large memory VMs	8.0
Increase the default maximum virtual devices to 64 per device	8.3
Allow additional processor features for Perfmon	9.0
Automatically expose simultaneous multithreading configuration for VMs running on hosts using the Core Scheduler	9.0
Hibernation support	9.0

Updating the VM configuration version

As you move or import a VM between Hyper-V hosts, it is important that the host supports the configuration version of the VM. If you move the VM to a newer version of Windows or Windows Server, the VM configuration version does not update automatically. However, you can manually update the VM configuration version to a newer version.

To update the VM configuration version, from the Hyper-V host machine, run the following PowerShell cmdlet:

```
Update-VMVersion <vmname>
```

When you update the configuration version of a VM, it is updated to the highest configuration level supported by the Hyper-V host on which it is running. For example, if you update the configuration version of a VM on a Hyper-V host running Windows Server 2019, the configuration version is updated to version 9.0.

Note: You cannot downgrade a VM configuration version after it has been upgraded. Also, once the VM configuration version is updated, you will be able to import the VM; however, the VM cannot start on hosts that don't support the updated configuration version.

VM generation versions

When you create a new VM, one of the options presented is whether to create a generation 1 or generation 2 VM.

Note: It is important to understand the impact and considerations of your generation selection, as you cannot change the generation after you have created it.

Generation 1 VMs

When you create a generation 1 VM, the following features are supported:

- **Guest operating systems.** Generation 1 VMs support both 32-bit and 64-bit Windows versions. Generation 1 also supports CentOS/Red Hat Linux, Debian, FreeBSD, Oracle Linux, SUSE Linux, and Ubuntu guest operating systems.

- **VM boot.** Generation 1 VMs can boot from a virtual floppy disk (.VFD), integrated drive electronics (IDE) Controller VHD, IDE Controller virtual DVD, or PXE boot by using a legacy network adapter. Generation 1 boot volumes only support a maximum of 2 TB with four partitions.
- **Firmware support.** Legacy BIOS.

Generation 2 VMs

When you create a generation 2 VM, the following features are supported:

- **Guest operating systems.** Generation 2 VMs support only 64-bit Windows versions (excluding Windows Server 2008 and Windows 7). Generation 2 also supports current versions of CentOS/Red Hat Linux, Debian, Oracle Linux, SUSE Linux, and Ubuntu guest operating systems.
- **Virtual machine boot.** Generation 2 VMs can only boot from a SCSI Controller VHD, SCSI Controller virtual DVD, or PXE boot by using a standard network adapter.
- **Secure boot.** Generation 2 VMs support Secure boot and are enabled by default. This feature verifies that the boot loader is signed by a trusted authority in the UEFI database.
- **Shielded virtual machines.** Generation 2 VMs support shielded VMs.
- **Larger boot volume.** Generation 2 VMs support a maximum boot volume size of 64 TB.
- **Firmware support.** UEFI.

Note: Generation 2 VMs do not have a DVD drive by default, but you can add a DVD drive after you create the VM. Also, generation 2 VMs don't support a virtual floppy disk controller.

VM settings

In the Hyper-V Manager, the virtual machine (VM) settings are grouped into two main sections; **Hardware** and **Management**. The configuration files that store hardware and management information are separated into two formats: .vmcx and .vmrs. The .vmcx format is used for configuring VMs, and the .vmrs format is used for runtime data. This helps decrease the chance of data corruption during a storage failure.

Hardware

VMs 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.

Generation 1 VMs have the following hardware by default:

- **BIOS.** Virtual hardware simulates a computer's BIOS. You can configure a VM to switch Num Lock on or off. You can also choose the startup order for a VM's virtual hardware. You can start a VM from a DVD drive, an integrated drive electronics (IDE) device, a legacy network adapter, or a floppy disk.
- **Memory.** You can allocate memory resources to a VM. An individual VM can allocate as much as 1 TB of memory. You can also configure Dynamic Memory to allow for dynamic memory allocation based upon resource requirements.
- **Processor.** You can allocate processor resources to a VM. You can allocate up to 64 virtual processors to a single VM.
- **IDE controller.** A VM can support only two IDE controllers and, by default, allocates two IDE controllers to a VM. These are IDE controller 0 and IDE controller 1. Each IDE controller can support two devices. You can connect virtual hard disks (VHDs) 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 VHDs and DVD-ROMs to VMs that use operating systems that don't support integration services.

- **SCSI controller.** You can use SCSI controllers only on VMs that you deploy with operating systems that support integration services. SCSI controllers allow you to support up to 256 disks by using four controllers with a maximum of 64 connected disks each. You can add and remove virtual SCSI disks while a VM is running.
- **Network adapter.** Hyper-V–specific network adapters represent virtualized network adapters. You can only use network adapters with supported VM guest operating systems that support integration services.
- **COM port.** A COM port enables connections to a simulated serial port on a VM.
- **Diskette drive.** You can map a .vfd floppy disk file to a virtual floppy drive.

Generation 2 VMs have the following hardware by default:

- **Firmware.** UEFI allows all the features of the BIOS in generation 1 VMs. However, it also allows secure boot, which is enabled by default.
- **Memory.** Same as generation 1 VMs.
- **Processor.** Same as generation 1 VMs.
- **SCSI controller.** Generation 2 VMs can use a SCSI controller for a boot device.
- **Network adapter.** Generation 2 VMs support hot add/removal of virtual network adapters.

You can add the following hardware to a VM by editing the VM's properties and then selecting **Add Hardware**:

- **SCSI controller.** You can add up to four virtual SCSI devices. Each controller supports up to 64 disks.
- **Network adapter.** A single VM can have a maximum of eight Hyper-V–specific network adapters.
- **Fibre Channel adapter.** This adapter allows a VM to connect directly to a Fibre Channel SAN. For this adapter, the Hyper-V host should have a Fibre Channel HBA that also has a Windows Server driver that supports virtual Fibre Channels.

Management

Use management settings to configure how a VM behaves on a Hyper-V host. The following VM management settings are configurable:

- **Name.** Use this setting to configure the display name of the VM on a Hyper-V host. Doing this does not alter the actual VM's computer name.
- **Integration Services.** Use this setting to configure which VM integration settings are enabled.
- **Checkpoints.** Use this setting to specify a location for storing VM checkpoints.
- **Smart Paging File Location.** This is the location you use when Smart Paging is required to start a VM.
- **Automatic Start Action.** Use this setting to handle how a VM responds when a Hyper-V host is powered on.
- **Automatic Stop Action.** Use this setting to handle how a VM responds when a Hyper-V host shuts down gracefully.

Storage options in Hyper-V

Just as a physical computer has a hard disk for storage, virtual machines (VMs) also require storage. Hyper-V provides many different VM storage options. If you know which option is appropriate for a given situation, you can ensure that a VM performs well and does not consume unnecessary space or place an unnecessary performance burden on the Hyper-V host server. You need to understand the various options for storing virtual hard disks (VHDs) so that you can select a storage option that meets your requirements for performance and high availability.

A key factor when provisioning VMs is to ensure correct placement and storage of the VHDs. Servers that otherwise are well provisioned with RAM and processor capacity can still experience poor performance if the storage system is overwhelmed or inadequate. You can store VHDs on local disks, a SAN, or Server Message Block (SMB) version 3.0 file shares.

Consider the following factors when you plan the storage location of VHD files:

- **High-performance connection to storage.** You can locate VHD 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. Slow network connections to storage or connections where there is latency result in poor VM performance.
- **Redundant storage.** The volume on which the VHD files are stored should be fault tolerant whether the VHD is stored on a local disk or on a remote NAS or SAN device. Hard disks often fail; therefore, the VM and the Hyper-V host should remain in operation after a disk failure. Replacing failed disks should not affect the operation of the Hyper-V host or VMs.
- **High-performance storage.** The storage device on which you store VHD files should have excellent input/output (I/O) characteristics. Many enterprises use hybrid solid state drives (SSDs) in RAID 1+0 arrays to achieve maximum performance and redundancy. Multiple VMs that are running simultaneously on the same storage can place a tremendous I/O burden on a disk subsystem. Therefore, you must ensure that you choose high-performance storage. If you don't, VM performance suffers.
- **Adequate growth space.** If you have configured VHDs to grow automatically, ensure that there is adequate space into which the files can grow. You should carefully monitor growth so that you are not surprised when a VHD fills the volume that you allocated to host it.

Fibre Channel support for SANs

Hyper-V virtual Fibre Channel adapter is a virtual hardware component that you can add to a VM to enable access to Fibre Channel storage on storage area networks (SANs). To deploy a virtual Fibre Channel:

- You must configure the Hyper-V host with a Fibre Channel host bus adapter (HBA) or Fibre Channel over Ethernet (FCoE) network adapter.
- The Fibre Channel HBA must have a driver that supports virtual Fibre Channel.

Virtual Fibre Channel adapters support port virtualization by exposing HBA ports in the guest operating system. Doing so allows a VM to access a SAN by using a standard World Wide Name that is associated with the VM.

You can deploy up to four virtual Fibre Channel adapters on each VM.

Storing VMs on SMB 3.0 file shares

Hyper-V supports storing VM data, such as VM configuration files, checkpoints, and VHD files on SMB 3.0 file shares. The file share must support SMB 3.0.

Note: The recommended bandwidth for network connectivity to an SMB file share is 1 gigabit per second (Gbps) or more.

SMB 3.0 file shares provide an alternative to storing VM files on iSCSI or Fibre Channel SAN devices. When creating a VM in Hyper-V you can specify a network share as the VM location and the VHD location. You can also attach disks that are stored on SMB 3.0 file shares. You can use .vhd, .vhdx, and .vhds files with SMB 3.0 file shares.

When you use SMB 3.0 file shares, you should separate network traffic to the file shares that contain the VM files. Client network traffic should not be on the same virtual LAN (VLAN) as SMB traffic.

To provide high availability for file shares storing VM files, you can use Scale-Out File Server (SOFS). SOFS provides redundant servers for accessing a file share. This also provides faster performance than when you are accessing files through a single share, because all servers in the SOFS are active at the same time. Windows Server 2016 and later can use Storage QoS to manage QoS policies for Hyper-V and SOFS. This allows deployment of QoS policies for SMB 3.0 storage.

Virtual hard disk formats and types

A *virtual hard disk* (VHD) is a file format that represents a traditional hard disk drive that VMs use for storage. You can configure a VHD with partitions, files, and folders.

You can create and manage VHDs by using:

- The **Hyper-V Manager** console.
- The **Disk Management** console.
- The **Diskpart** command-line tool.
- The **New-VHD** Windows PowerShell cmdlet.
- The **Windows Admin Center** console.

VHD formats

When you create a new VHD, you often choose between the following formats:

- **VHD**
 - Supports virtual disks up to 2,040 GB in size.
 - Select this option if you need to support older Hyper-V versions.
- **VHDX**
 - Supports virtual disks up to 64 TB in size.
 - This format also contains a file structure that minimizes corruption issues if the host server suffers from an unexpected power outage.
 - The .vhdx format supports larger block size for dynamically expanding and differencing disks, which provides increased performance.
 - If the VHDX disk is connected to an SCSI controller, you can extend or shrink the disk size while the VM is running.

You can convert between VHD formats using Hyper-V Manager's **Edit Virtual Hard Disk Wizard** or by using the **Convert-VHD** PowerShell cmdlet. When you do so, a new VHD is created, and the contents of the existing VHD are copied into it. Therefore, ensure that you have sufficient disk space to perform the conversion.

VHD types

Hyper-V supports multiple VHD types, which have varying benefits and drawbacks. The type of hard disk you select will vary depending on your needs. The VHD types are:

- **Fixed size.**
 - This type of VHD allocates all of the space immediately. This minimizes fragmentation, which in turn enhances performance.
- **Dynamically expanding**
 - When you create a dynamically expanding VHD, 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.
 - This type of disk provides better use of physical storage space and only uses space as needed.
- **Differencing**
 - This type of disk is associated with another virtual disk in a parent-child configuration.
 - The goal of a differencing disk is to use a parent disk that contains a base installation and configuration. Any changes made to the differencing disk do not affect the parent disk.
 - Differencing disks are typically used to reduce data storage requirements for child disks that might use the same parent configuration. For example, you might have 10 differencing disks based on the same parent disk that contains a sysprepped image of Windows Server 2019. You then could use the 10 differencing disks to create 10 different VMs. Any changes to the 10 individual VMs will not affect the parent disk. Only the differencing disks are changed.

You can use the **Edit Virtual Hard Disk Wizard** to convert between disk types; however, the target disk format should be the same as the source disk format.

Using physical hard disks

A VM can also use a physical hard disk attached to the host computer rather than using a VHD. Also known as a *pass-through disk*, you can use this type of hard disk configuration to connect directly to an Internet SCSI (iSCSI) LUN or a physical disk attached on the host machine. When you use pass-through disks, the VM must have exclusive access to the target disk. To use 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 VM's disk controllers.

You can attach a pass-through disk to a VM by performing the following steps:

1. Ensure that the target hard disk is offline.
2. Use Hyper-V Manager to edit an existing VM's properties.
3. Select an integrated drive electronics (IDE) or SCSI controller, select **Add**, and then select **Hard Drive**.
4. On the **Hard Drive** page, select **Physical hard disk**. In the drop-down list, select the disk that you want to use as the pass-through disk.

Shared VHDX and VHD Set files

In some scenarios, you might need to share a single virtual hard disk (VHD) between multiple VMs. This is often the case when you incorporate high availability using VMs configured to support failover clustering. Hyper-V in Windows Server 2019 supports two methods for sharing a VHD between multiple VMs, shared VHDXs, and VHD Sets.

Shared VHDXs

Windows Server 2012 R2 and newer operating systems support the ability to create a VHD in the .VHDX format and connect the file to the SCSI controllers of multiple VMs. The shared VHDXs must be stored on cluster shared volumes or a file server with Server Message Block (SMB) version 3.0 file-based storage.

Using a shared VHD with guest failover clustering does pose limitations, such as:

- The .VHDX disk format does not support resizing of the file while the cluster is running. You must shut down the cluster to resize the disk.
- Shared VHDXs do not support Hyper-V Replica to replicate the VM failover cluster.
- Backing up the VMs from the host machine is not supported.

To address these limitations, Windows Server 2016 and later provides the ability to create VHD Sets.

VHD Sets

A **VHD Set** provides the next evolution for sharing virtual disk files with multiple VMs. Consider the following factors involved in using a VHD Set:

- Requires Windows Server 2016, Windows 10 or later.
- Uses the .VHDS file format for the shared disk along with an .AVHDX file that is used as a checkpoint file.
- Supports both fixed size and dynamically expanding disk type.
- Supports dynamic resizing, backup at the host level, and the ability to use Hyper-V replica.

Tip: You can convert from a Shared VHDX to a VHD Set by using the **Convert-VHD** PowerShell cmdlet.

Overview of Hyper-V networking

Networking in Hyper-V basically consists of two primary components: a **virtual network adapter** and a **virtual switch**. The virtual network adapter is configured on the virtual machine (VM) and connects to a port on the virtual switch to communicate on a network.

Virtual network adapter types

Hyper-V supports the following virtual network adapter types:

- **Legacy network adapter.** Emulates an Intel 21140-based PCI Fast Ethernet Adapter. Can be used to boot to a network for operating system installation tasks. Available only in generation 1 VMs.
- **Network adapter.** Also known as a “synthetic” network adapter. This type is faster than the legacy network adapter; however, it does not support booting to a network for generation 1 VMs. This network adapter can be used with both generation 1 and generation 2 VMs.

Virtual switch types

A virtual switch is used to control how network traffic flows between VMs that are hosted on a Hyper-V server, in addition to how network traffic flows between VMs and the rest of the organizational network.

Hyper-V supports three types of virtual switches:

Type	Description
External	This type of switch is used to map a network to a specific network adapter or network adapter team. Hyper-V also supports mapping an external network to a wireless network adapter if you have installed the Wireless local area network (LAN) service on the host Hyper-V server and if the Hyper-V server has a compatible network adapter.
Internal	The internal virtual switch is used to communicate between the VMs on a Hyper-V host and to communicate between the VMs and the Hyper-V host itself.
Private	A private switch is used only to communicate between VMs on a Hyper-V host. You cannot use private switches to communicate between VMs and the Hyper-V host.

When configuring a virtual network, you can also 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.

To create and manage a virtual switch for Hyper-V, you can use the following tools:

- Hyper-V Manager
- New-VMSwitch PowerShell cmdlet
- Windows Admin Center

Networking features for Hyper-V

Several features in Windows Server Hyper-V networking improve network performance and the flexibility of virtual machines (VMs) in private and public cloud environments. The following table provides a summary of features that Windows Server Hyper-V networking supports:

Feature	Description
Network virtualization	This feature allows IP address virtualization in hosting environments so that VMs that migrate to the host can keep their original IP addresses rather than being allocated IP addresses on the Hyper-V server's network. This feature does require Windows Server Datacenter edition.

Feature	Description
Bandwidth management	You can use this feature to specify a minimum and maximum bandwidth that Hyper-V allocates to an adapter. Hyper-V reserves the minimum bandwidth allocation for the network adapter even when other virtual network adapters on VMs that are hosted on the Hyper-V host are functioning at capacity.
Dynamic Host Configuration Protocol (DHCP) guard	This feature drops DHCP messages from VMs that are functioning as unauthorized DHCP servers. This might be necessary in scenarios where you are managing a Hyper-V server that hosts VMs for others but in which you don't have direct control over the virtual VMs' configuration.
Router guard	This feature drops router advertisement and redirection messages from VMs that are configured as unauthorized routers. This might be necessary in scenarios where you don't have direct control over the configuration of VMs.
Port mirroring	You can use this feature to copy incoming and outgoing packets from a network adapter to another VM 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 copy from the host operating system to the VM. 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 are installed on the guest operating system. SR-IOV enables multiple VMs 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.

Feature	Description
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.

Windows Server 2016 and later provides additional networking features to support Software Defined Networking (SDN) infrastructures. These improvements include:

- **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. Another advantage of SET is that you can add multiple Remote Direct Memory Access (RDMA) network adapters, which was not available with traditional teams.
- **RDMA with Hyper-V.** Also known as Server Message Block (SMB) Direct, this is a feature that requires hardware support in the network adapter. A network adapter with RDMA functions at full speed with low resource utilization. Effectively, this means that there is higher throughput, which is important for busy servers with high-speed network adapters such as 10 Gbps. RDMA services can now use Hyper-V switches. You can enable this feature with or without SET.
- **Virtual machine multi queues (VMMQ).** VMMQ improves on VMQ by allocating multiple queues per VM and spreading traffic across the queues.
- **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 VM traffic. This reduces the number of specialized adapters that each host needs.
- **Network Address Translation (NAT) object.** NAT is often useful to control the use of IP addresses. This is particularly true if there are many VMs that require access to the Internet. However, there is no requirement for communication to be initiated from the Internet back to the internal VMs. Windows Server includes a NAT object that translates an internal network address to an external address. You can use the **New-NetNat** PowerShell cmdlet to create a NAT object.

Manage VM states and checkpoints

Managing virtual machine (VM) state

It is important to understand how the state of a VM impacts the resources that it is using. This ensures that your Hyper-V host has sufficient resources to support the VMs that reside on it.

The following table summarizes VM states

State	Description
Off	A VM that is off does not use any memory or processing resources.

State	Description
Starting	A VM that is starting verifies that resources are available before allocating those resources.
Running	A VM that is running uses the memory that has been allocated to it. It can also use the processing capacity that has been allocated to it.
Paused	A paused VM does not consume any processing capacity, but it does still retain the memory that has been allocated to it.
Saved	A saved VM does not consume any memory or processing resources. The memory state for the VM is saved as a file and is read when the VM is started again.

Managing checkpoints

A Checkpoint allows you to make a snapshot of a VM at a specific time. Windows Server Hyper-V supports two types of checkpoints: *production checkpoints* and *standard checkpoints*. Production checkpoints is the default. It is important to identify when to use a standard checkpoint and when to use a production checkpoint.

Caution: Ensure that you only use checkpoints with server applications that support the use of checkpoints. Reverting to a previous checkpoint of a VM that contains an application that does not support VM checkpoints might lead to data corruption or loss. Depending on the application, it might support either a standard checkpoint or a production checkpoint. Most applications must be stopped to use a standard checkpoint. You can use production checkpoints in cases where backup software is supported.

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. You can also use the **Windows Admin Center** or **PowerShell** to create and manage checkpoints. Each VM can have a maximum of 50 checkpoints.

When creating checkpoints for multiple VMs 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 revert a computer back to a point before it performed a computer password change with a domain controller, you must rejoin that computer to the domain.

Checkpoints are not a replacement for backups. If the volume that hosts these files fails, both the checkpoint and the virtual hard disk (VHD) files are lost. You can create a backup from a checkpoint by performing a VM export of a checkpoint. When you export the checkpoint, Hyper-V creates full VHDs that represent the state of the VM when you created the checkpoint. If you choose to export an entire VM,

Standard checkpoints

When you create a standard checkpoint, Hyper-V creates an .avhd file (differencing disk) that stores the data that differentiates the checkpoint from either the previous checkpoint or the parent VHD. When you delete standard checkpoints, this data is either discarded or merged into the previous checkpoint or parent VHD. For example, if you delete the most recent checkpoint of a VM, the data is discarded. If you

delete the second to last checkpoint of a VM, the content of the differencing VHD merges with its parent, so that the earlier and latter checkpoint states of the VM retain their integrity.

Production checkpoints

When you create a production checkpoint, Windows Server uses Volume Shadow Copy Service (VSS) (or File System Freeze for Linux). This places the VM in a safe state to create a checkpoint that can be recovered in the same way as any VSS or application backup. Unlike standard checkpoints that save all memory and processing in the checkpoint, production checkpoints are closer to a state backup. Production checkpoints require a VM to start from an offline state.

Applying checkpoints

When you apply a checkpoint, the VM reverts to the configuration that existed at the time it took the checkpoint. Reverting to a checkpoint does not delete any existing checkpoints. If you revert to a checkpoint after making a configuration change, you receive a prompt to create a checkpoint. Creating a new checkpoint is necessary only if you want to return to that current configuration.

You can create checkpoint trees that have different branches. For example, if you create a checkpoint of a VM on Monday, Tuesday, and Wednesday, and then apply the Tuesday checkpoint, and then make changes to the VM's configuration, you create a new branch that diverts from the original Tuesday checkpoint. You can have multiple branches if you don't exceed the 50-checkpoint limit per VM.

Import and export VMs

You can use Hyper-V import and export functionalities to transfer VMs between Hyper-V hosts and to create point-in-time backups of VMs.

Importing VMs

The VM import functionality in Hyper-V can identify configuration problems such as missing hard disks or virtual switches. In Hyper-V for Windows Server 2016 and later, you can import VMs from copies of VM configurations, checkpoints, and virtual hard disk (VHD) files, rather than specifically exported VMs. This is beneficial in recovery situations where an operating system volume might have failed but the VM files remain intact.

When importing a VM, you have three options:

- **Register the VM in-place (use the existing unique ID).** This option creates a VM by using the files in the existing location.
- **Restore the VM (use the existing unique ID).** This option copies the VM files back to the location from which they were exported and then creates a VM by using the copied files. This option effectively functions as a restore from backup.
- **Copy the VM (create a new unique ID).** This option copies the VM files to a new location that you can specify and then creates a new VM by using the copied files.

Exporting VMs

When exporting a VM, you have two options:

- **Export a specific checkpoint.** This enables you to create an exported VM as it existed at the point of checkpoint creation. The exported VM will have no checkpoints. Select the checkpoint to be exported, and then select **Export**.
- **Export a VM with all checkpoints.** This exports the VM and all checkpoints that are associated with the VM. From the Virtual Machine list in Hyper-V Manager, select the VM to be exported and then select **Export**.

Note: Hyper-V in Windows Server 2016 and later supports exporting VMs and checkpoints while a VM is running.

Demonstration: Create and manage a VM

This demonstration illustrates how to:

- Configure a Hyper-V virtual switch.
- Create a virtual hard disk (VHD).
- Create a VM.
- Manage VMs using Windows Admin Center.

Demonstration Steps

Configure a Hyper-V virtual switch

1. On SEA-ADM1, select **Start**, and then select **Server Manager**.
2. In Server Manager, select **All Servers**.
3. In the Servers list, select and hold **SEA-SVR1**, and then select **Hyper-V Manager**.
4. Use the **Virtual Switch Manager** to create the following switch:
 - Name: **Contoso Private Switch**
 - Connection type: **Private network**

Create a VHD

1. On SEA-ADM1, in Hyper-V Manager, select **New**, and then select **Hard Disk**. The **New Virtual Hard Disk Wizard** starts.
2. Create a new hard disk as follows:
 - Disk Format: **VHD**
 - Disk Type: **Differencing**
 - Name: **SEA-VM1**
 - Parent Disk: **C:\Base\Baselimage.vhd**

Create a VM

1. On SEA-ADM1, in Hyper-V Manager, create a new VM as follows:
 - Name: **SEA-VM1**
 - Generation: **Generation 1**
 - Memory: **4096**
 - Networking: **Contoso Private Switch**
 - Hard disk: **SEA-VM1.vhd**
2. Select **SEA-VM1**, and then in the Actions pane, under SEA-VM1, select **Settings**.
3. Close Hyper-V Manager.

Manage VMs using Windows Admin Center

1. On SEA-ADM1, on the taskbar, select **Microsoft Edge**.
2. In Microsoft Edge, on the Favorites Bar, select **Windows Admin Center**.
3. In the Windows Security box, enter **Contoso\Administrator** by using **Pa55w.rd** as the password, and then select **OK**.
4. In the **All connections** list, select **SEA-SVR1**.
5. In the **Specify your credentials** page, select **Use another account for this connection**, and then enter **Contoso\Administrator** by using **Pa55w.rd** as the password.
6. In the **Tools** list, select **Virtual Machines**. Display the Summary pane.
7. Create a new disk, 5 GB in size.
8. Start **SEA-VM1**. Scroll down and then observe the statistics for the running VM.
9. Refresh the page, and then select **Shut down the VM**.
10. In the **Tools** list, select **Virtual switches** and identify the existing switches.

Test Your Knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

You need to create a virtual machine (VM) that supports Secure boot. Which generation would you choose when you create the VM?

Question 2

Which virtual hard disk (VHD) type only uses the amount of space that needs to be allocated and grows in size as more space is necessary?

Question 3

Which Hyper-V virtual switch allows communication between the VMs on a host computer and also between the VMs and the host itself only?

Question 4

You need to preserve the state and configuration of a VM at a set time period. What can you do?

Securing virtualization in Windows Server

Lesson overview

Most organizations use Hyper-V to virtualize network infrastructure services that traditionally have been hosted on physical servers. Virtualization provides many benefits related to consolidation, portability, and ease of management; however, these benefits also introduce unique security concerns. Unlike a physical server that might be protected in a secure data center, virtual machine (VM) files can simply be exported, copied offsite, and then imported to run on any Hyper-V host server. These VMs, which might consist of services such as domain controllers, HR systems, or sensitive file servers, often contain confidential information that must be protected from malicious use.

Hyper-V supports the concept of a *guarded fabric* to provide a more secure environment for VMs. In this lesson, you are introduced to the concept of implementing a guarded fabric, including the Host Guardian Service (HGS), guarded host servers, and shielded VMs.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the guarded fabric.
- Describe attestation modes in a guarded fabric.
- Explain the HGS.
- Explain the types of protected VMs in a guarded fabric.
- Describe the general process for creating a shielded VM.
- Describe the process for when a shielded VM is powered on in a guarded fabric.

Guarded fabric

A *guarded fabric* in Hyper-V is a security solution that is used to protect virtual machines (VMs) against inspection, theft, and tampering from either malware or malicious system administrators. The VMs that are part of a guarded fabric are called *shielded VMs* and are protected both at rest and during runtime. A shielded VM is encrypted and can only run on healthy and approved hosts within the guarded fabric infrastructure.

The guarded fabric provides a number of security benefits to the virtual infrastructure, such as:

- **Secure and authorized Hyper-V hosts.** Hyper-V hosts that are part of a guarded fabric are called *guarded hosts*. A guarded host is allowed to run a shielded VM only if it can prove that it is in a known, trusted state. A guarded host provides health information and requests permission to start a shielded VM from an external authority called the *Host Guardian Service* (HGS).
- **Verification that a host is in a healthy state.** The HGS performs *attestation*, which means that the service can measure the health of a guarded host and provide certification that the host is healthy and authorized to run the shielded VM.
- **Providing a secure method to release keys to healthy hosts.** Once a guarded host has been verified as healthy and authorized, the HGS releases a secure key, which is used to unlock and start a shielded VM.

To summarize, a guarded fabric is made up of the following components:

- **Guarded Hyper-V hosts.** You might have one or more guarded Hyper-V hosts running Windows Server Datacenter.
- **Host Guardian Service.** Typically, a three-node cluster running the HGS server role.
- **Shielded virtual machines.** A VM that has a virtual trusted platform module (TPM) and is encrypted using BitLocker.

Note: A guarded fabric is capable of running a normal VM with no protection, similar to a standard Hyper-V environment. You can also implement Encryption-supported VMs that are secured by encryption but don't have the same restrictions in place as shielded VMs.

Tools used to automate and manage a guarded fabric

To help manage and automate the processes within the guarded fabric infrastructure, the following tools are typically used:

- **System Center Virtual Machine Manager (VMM).** Provides a unified management solution for virtualization hosts, virtual networking, and storage for effectively maintaining an enterprise virtual environment.
- **Windows Azure Pack.** Provides a graphical web portal interface to manage a guarded fabric and shielded VMs.
- **PowerShell.** If you don't have a fabric manager such as VMM, you can use PowerShell to manually create the components needed for a new shielded VM. You can also use PowerShell to provision the shielded VM to a guarded host.

Attestation modes for guarded fabric

In the guarded fabric infrastructure, Hyper-V hosts can only run protected virtual machines (VMs) after they have been validated by the Host Guardian Service (HGS). The process of evaluating and validating the Hyper-V host is called *attestation*. Protected VMs can't be powered-on or live migrated to a Hyper-V host that has not yet attested or has failed attestation.

The HGS supports the following attestation modes:

- **Trusted platform module (TPM)-trusted attestation.** A hardware-based attestation method. This method offers the strongest protections, but does require a more complex configuration and higher host hardware requirements. These requirements include TPM 2.0 and UEFI 2.3.1 with Secure Boot enabled. A guarded Hyper-V host is approved and validated based upon its TPM identity, Measured Boot sequence, and code integrity policies to ensure that it only runs approved code.
- **Host key attestation.** This mode is based upon asymmetric key pairs and is used when existing Hyper-V host machines don't support TPM 2.0. A guarded Hyper-V host is approved and validated based upon possession of the key.

Note: Windows Server 2016 included another mode called **Admin-trusted attestation**. This method has been deprecated beginning with Windows Server 2019.

Host Guardian Service

The Host Guardian Service (HGS) contains two important components that ensure validity of guarded hosts and protection for virtual machines (VMs). These two components include the following:

- **Attestation Service.** Ensures that only trusted Hyper-V hosts can run protected VMs.

- **Key Protection Service (KPS).** Provides the keys necessary to power-on protected VMs and to permit live migration to other guarded Hyper-V hosts.

The HGS provides authority for the guarded fabric and helps to enforce and assure the following:

- **Protected VMs contain BitLocker encrypted disks**
 - Protected VMs use BitLocker to protect both the operating system disk and data disks.
 - The virtual trusted platform module (TPM) protects the BitLocker keys needed to boot the VM and decrypt the disks.
- **Shielded VMs are deployed from trusted template disks and images**
 - When deploying new shielded VMs, the VM owner is able to specify which template disks they trust.
 - Shielded template disks have signatures that are computed when their content is deemed trustworthy. The disk signatures are then stored in a signature catalog, which is securely provided to the fabric when creating shielded VMs.
 - During provisioning of shielded VMs, the signature of the disk is computed again and compared to the trusted signatures in the catalog. If the signatures match, the shielded VM is deployed. If the signatures don't match, the shielded template disk is deemed untrustworthy, and deployment fails.
- **Passwords and other secrets are protected when a shielded VM is created**
 - When creating VMs, it is necessary to ensure that VM secrets, such as the trusted disk signatures, Remote Desktop Protocol (RDP) certificates, and the password of the VM's local Administrator account, are not divulged to the fabric. These secrets are stored in an encrypted file called a *shielding data file* (a *.PDK* file), which is protected by certificate keys and uploaded to the fabric.
 - When a shielded VM is created, the VM owner selects which shielding data file to use so that these secrets can only be provided to the trusted components within the guarded fabric.
- **Control of where the shielded VM can be started**
 - The shielding data file also contains a list of the guarded fabrics on which a particular shielded VM is permitted to run. This is useful in cases where a shielded VM typically resides in an on-premises private cloud but might need to be migrated to another public or private cloud.
 - The target cloud or fabric must support shielded VMs, and the shielded VM's shielding data file must permit that fabric to run it.

Considerations when planning to implement the HGS

Consider the following factors when planning to implement the HGS:

- **Hardware requirements.** The HGS can be run on both physical and virtual machines; however, if you want to run HGS as a three-node cluster, three physical servers are highly recommended.
- **Operating System requirements.** If you plan to use TPM trusted attestation, the HGS can run on Windows Server 2019 or Windows Server 2016 Standard or Datacenter edition. Host key attestation requires Windows Server 2019 Standard or Datacenter edition operating with v2 attestation, which requires a TPM certificate to be present to add the TPM identifier information to the HGS.
- **Certificate requirements.** When you deploy HGS, you will be asked to provide signing and encryption certificates that are used by the KPS to protect the sensitive information needed to start up a shielded VM. These certificates never leave the HGS and are only used to decrypt shielded VM keys when the host on which they're running has proven it is healthy. Although not required, you should obtain the certificates from a trusted certificate authority.

- **Server Roles required.** The HGS and supporting server roles are installed on the server. The default installation sets up the server in a new Active Directory forest dedicated for HGS. This is to ensure that all sensitive key information is as secure as possible from the rest of the organizational network environment.

Types of protected VMs in a guarded fabric

A guarded fabric can run the following types of virtual machines (VMs):

- **A shielded VM.** Used in fabrics where the data and state of the VM must be protected from both fabric administrators and untrusted software that might be running on the Hyper-V hosts.
- **An encryption-supported VM.** Used where the fabric administrators are fully trusted. For example, an enterprise might deploy a guarded fabric to ensure VM disks are encrypted at rest for compliance purposes. Fabric administrators can continue to use convenient management features, such as VM console connections, PowerShell Direct, and other day-to-day management and troubleshooting tools.
- **A normal VM.** Unprotected VMs can be run in the fabric if needed. However, a normal VM does not offer any of the security benefits of a shielded or encryption-supported VM.

Differences between encryption-supported and shielded VMs

The following table provides a summary of the differences between encryption-supported and shielded VMs:

Capability	Encryption-supported	Shielded
Secure boot	Yes, required but configurable	Yes, required and enforced
Virtual TPM	Yes, required but configurable	Yes, required and enforced
Encrypt VM state and live migration traffic	Yes, required but configurable	Yes, required but enforced
Integration components	Configurable by fabric admin	Certain components blocked such as PowerShell Direct (enabled for Windows Server v1803 and later) and data exchange
Virtual Machine Connection, HID devices (keyboard, mouse)	On, can't be disabled	Enabled for hosts beginning with Windows Server v1803; Disabled on earlier hosts
COM/Serial ports	Supported	Disabled (can't be enabled)
Attach a debugger to the VM process	Supported	Disabled (can't be enabled)

Note: Both encryption-supported and shielded VMs must be configured as generation 2 VMs.

General process for creating shielded VMs

Creating a shielded virtual machine (VM) requires the following general steps:

1. Create a shielded VM template disk
2. Create a shielded data file

3. Deploy the shielded VM

1. Creating a shielded VM template disk

A VM template disk is used as a basis for all future VMs deployed within the fabric. A shielded VM template disk provides additional security by hiding all confidential information such as passwords, certificates, and so on, from the fabric administrators.

Preparing the operating system VHDX

To create the template disk, the first step is to prepare an operating system disk using your standard processes. The goal is to set up a VM with a blank VHDX and install an operating system onto that disk.

The disk you set up must meet the following requirements:

VHDX requirement	Description
Must be a GUID Partition Table (GPT) disk	Needed for generation 2 VMs to support UEFI
Disk type must be Basic	Disks can't be Dynamic as BitLocker doesn't support dynamic disks.
Disk must have at least two partitions	One partition must include the drive that you will install Windows on. BitLocker encrypts this drive. The other partition is the active partition, which contains the bootloader and remains unencrypted so that the computer can start.
File system must be NTFS	Required for BitLocker
Supported operating system	Must be Windows 10 or Windows Server 2012 or later. Must be able to support generation 2 VMs and the Microsoft Secure Boot process.
Operating system must be generalized	The template disk must be generalized by running sysprep.exe

Note: After you install the operating system, be sure to verify and install all the latest Windows updates. Updates are often released to improve the reliability of the end-to-end shielded process.

Use the Shielded Template Disk Creation Wizard

Before you can use the template disk to create shielded VMs, the generalized disk must be prepared and encrypted with BitLocker. You perform this task by using the **Shielded Template Disk Creation Wizard**. The wizard generates a hash for the disk and adds it to a volume signature catalog (VSC). The VSC is signed using a certificate you specify and is used during the provisioning process to ensure the disk being deployed in the fabric has not been altered or replaced with a disk that is not trusted. Finally, BitLocker is installed on the disk's operating system to prepare the disk for encryption during VM provisioning.

The **Shielded Template Disk Creation Wizard** is part of the **Shielded VM Tools** available from the **Remote Server Administration Tools** feature.

2. Creating a shielded data file

A shielded data file (also called a provisioning data file or PDK file) is an encrypted file that a VM owner creates to protect important VM configuration information such as the administrator password, Remote

Desktop Protocol (RDP) and other identity-related certificates, domain-join credentials, and so on. Before creating this file, you need to create or obtain a shielded template disk as described previously.

You will use the **Shielding Data File Wizard**, which is another tool provided by the **Shielded VM Tools** included with the **Remote Server Administration Tools** feature.

3. Deploying a shielded VM

The method used to deploy a shielded VM depends on your management process for the guarded fabric. Common methods to deploy a shielded VM include:

- Deploying using System Center Virtual Machine Manager (VMM).
- Using the Windows Azure Pack to provide a web-based portal to simplify shielded VM deployments.
- Using Windows PowerShell.

Process for powering-on shielded VMs

When you power-on a shielded virtual machine (VM) within a guarded fabric, several processes take place to validate the guarded host, unlock the shielded VM, and then allow the protected VM to start. The general process is described as follows:

1. **User requests to start a shielded VM.** When a fabric administrator attempts to start a shielded VM, the guarded host can't power on the VM until it has attested that it is healthy.
2. **Host requests attestation.** To prove that it is healthy, the guarded host requests attestation. The mode of attestation is determined by the *Host Guardian Service* (HGS). The mode to be used might be either *TPM-trusted attestation* or *Host key attestation*.
3. **Attestation succeeds or fails.** The attestation mode determines which checks are needed to successfully attest the host is healthy. With TPM-trusted attestation, the host's TPM identity, boot measurements, and code integrity policy are validated. With host key attestation, only registration of the host key is validated.
4. **Attestation certificate sent to host.** If attestation is successful, a health certificate is sent to the host, and the host is considered "guarded" and authorized to run shielded VMs. The host uses the health certificate to request the *Key Protection Service* (KPS) to securely release the keys needed to work with shielded VMs.
5. **Host requests VM key.** Guarded hosts need to request the necessary keys from the KPS on the HGS. To obtain the necessary keys, the guarded host provides the current health certificate and an encrypted secret file that can only be decrypted by the KPS.
6. **Key is released.** The KPS examines the health certificate provided by the guarded host to determine its validity. The certificate must not have expired and KPS must trust the attestation service that issued it.
7. **Key is returned to host.** If the health certificate is valid, the KPS attempts to decrypt the provided secret from the guarded host and then securely returns the keys needed to power on the VM.
8. **Host powers on shielded VM.** The guarded host can now unlock and power-on the shielded VM.

Test Your Knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

Describe three main benefits of running protected virtual machines (VMs) in a guarded fabric.

Question 2

Which component in a guarded fabric is used to enforce security and manage the keys to start protected VMs?

Question 3

Describe three types of VMs that can be run in a guarded fabric.

Question 4

Which tool is used to prepare and encrypt a VM template disk?

Containers in Windows Server

Lesson overview

Windows Server 2019 supports the development, packaging, and deployment of apps and their dependencies in Windows containers. By using container technology, you can package, provision, and run applications across diverse environments located on-premises or in the cloud. Windows containers provide a complete lightweight and isolated operating system-level virtualization environment to make apps easier to develop, deploy, and manage.

In this lesson, you are introduced to the concept of preparing and using Windows containers.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe containers and how they work.
- Explain the difference between containers and virtual machines.
- Describe the difference between Process Isolation and Hyper-V isolation modes.
- Describe Docker and how it is used to manage Windows containers.
- Identify the container base images available from the Microsoft Container Registry.
- Understand the process for running a Windows container.
- Explain how to manage containers using the Windows Admin Center.
- Deploy Windows containers by using Docker.

What are containers?

Traditionally, a software application is developed to run only on a supported processor, hardware, and operating system platform. Software applications typically cannot move from one computing platform to another without extensive recoding to provide support for the intended platform. With so many diverse computing systems, a more efficient software development and management platform was needed to support portability between multiple computing environments.

A *Container* is used to package an application along with all of its dependencies and abstract it from the host operating system in which it is to run. Not only is a container isolated from the host operating system, it is also isolated from other containers. Isolated containers provide a virtual runtime, which also improves the security and reliability of the apps that run within them.

Benefits of using containers include the following:

- **The ability to run anywhere.** Containers can run on various platforms such as Linux, Windows, and Mac operating systems. They can be hosted on a local workstation, on servers in on-premises data-centers, or provisioned in the cloud.
- **Isolation.** To an application, a container appears to be a complete operating system. The CPU, memory, storage, and network resources are virtualized within the container isolated from the host platform and other applications.
- **Increased efficiency.** Containers can be quickly deployed, updated, and scaled to support a more agile development, test, and production life cycle.

- **A consistent development environment.** Developers have a consistent and predictable development environment that supports various development languages such as Java, .NET, Python, and Node. Developers know that no matter where the application is deployed, the container will ensure that the application runs as intended.

How containers work

The processor in a standard Windows computer has two different modes: a *user mode* and a *kernel mode*. Core operating system components and most device drivers run in kernel mode, whereas applications run in user mode.

When you install container technology on a computer, each container creates an isolated, lightweight silo used for running an app on the host operating system. A container builds upon and shares most of the host operating system's kernel to gain access to the file system and registry.

Each container has its own copy of the user mode system files, which are isolated from other containers and the host's own user mode environment. The ability to isolate user mode is provided by a container base image, also referred to as a template, which consists of the user mode system files needed to support a packaged app. Container base image templates provide a foundational layer of operating system services used by the containerized app that are not provided (or restricted) from the host's kernel mode layer.

Containers vs. virtual machines

Both virtual machines (VMs) and containers are virtualization technology used to provide isolated and portable computing environments for applications and services.

As previously described, containers build upon the host operating system's kernel and contain an isolated user mode process for the packaged app. This helps to make containers very lightweight and fast to launch.

VMs simulate an entire computer, including the virtualized hardware, operating system, user mode, and its own kernel mode. A VM is quite agile and provides tremendous support for applications; however, VMs tend to be large and can take up a lot of resources from the host machine.

The following table summarizes the similarities and differences of containers and VMs:

Feature	Virtual machine	Container
Isolation	Provides complete isolation from the host operating system and other VMs	Provides lightweight isolation from the host and other containers
Operating system	Runs a complete operating system including the kernel	Only runs the user mode portion of an operating system
Guest compatibility	Able to run any supported operating system inside the VM	If running in Process isolation mode, you must run on the same type of kernel (OS) as the host; Hyper-V isolation mode provides more flexibility on running non-Windows containers on Windows hosts

Feature	Virtual machine	Container
Deployment	Deployed using Hyper-V manager or other VM management tools	Deployed and managed using Docker. Multiple containers can be deployed using an orchestrator such as Azure Kubernetes Service
Persistent storage	Uses virtual hard disk files or Server Message Block (SMB) share	Azure disks for local storage; Azure files (SMB shares) for storage shared by multiple containers
Load balancing	Uses Windows failover cluster to move VMs as needed	Uses an orchestrator to automatically start and stop containers
Networking	Uses virtual network adapters	Creates a default NAT network, which uses an internal vSwitch and a Windows component named WinNAT

When to choose VMs vs. containers

Use a VM when you:

- Need to manage a number of operating systems.
- Need to run an app that requires all the resources and services of a full operating system, such as a graphical user interface.
- Need an environment that preserves changes and is persistent.
- Require full isolation and security.

Use a container when you:

- Need a lightweight application package that quickly starts.
- Need to deploy multiple instances of a single app.
- Need to run an app or process that is nonpersistent in an on-demand basis.
- Need to deploy an app that can run on any underlying infrastructure.

Note: It is quite common for containers to be provisioned within a highly optimized VM to provide enhanced isolation and security. The next topic describes isolation modes, which provide an option to provision container runtime isolation.

Overview of container isolation modes

Windows containers can run in one of two distinct isolation modes. Both modes support identical processes for creating, managing, and running containers. However, there is a difference between the degree of isolation and security between the container, other containers, and the host operating system.

Windows containers support the following isolation modes:

- **Process Isolation.** Considered the traditional isolation mode for Windows containers, process isolation allows multiple container instances to run concurrently on a host. When running in this mode, containers share the same kernel with each other and the host operating system. Each provisioned container features its own user mode to allow Windows and app processes to run isolated from other containers. When you configure Windows containers to use process isolation, containers

can run multiple apps in isolated states on the same computer, but they do not provide security-enhanced isolation.

- **Hyper-V Isolation.** With Hyper-V isolation, each container runs inside a highly optimized virtual machine (VM). The advantage of this mode is that each container effectively gets its own kernel, providing an enhanced level of stability and security. The VM provides an additional layer of hardware-level isolation between each container and to the host computer. When deployed, a container using Hyper-V isolation mode starts in seconds, which is much faster than a VM with a full Windows operating system.

Note: Windows containers running on Windows server default to using process isolation. Windows containers running on Windows 10 Pro and Enterprise default to running with Hyper-V isolation mode.

When you create a container using Docker, you can specify the isolation mode by using the `-isolation` parameter. The following examples illustrate commands used to create a container using each of the isolation modes:

Process isolation mode:

```
docker run -it --isolation=process mcr.microsoft.com/windows/server-core:ltsc2019 cmd
```

Hyper-V isolation mode:

```
docker run -it --isolation=hyperv mcr.microsoft.com/windows/server-core:ltsc2019 cmd
```

Note: Additional information on using Docker to create and manage containers is provided later in this module.

Manage containers using Docker

Overview

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, called a Docker container, is software wrapped in a complete file system that includes everything it needs to run: 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, called the *Docker Engine*. The Docker Engine is a lightweight runtime environment that runs on Linux, macOS, or Windows-based operating systems.

Another component called the *Docker client* can be used as a command line interface (CLI) to integrate with the engine and run commands for building and managing the Docker containers provisioned on the host computer.

Docker containers are based upon open standards that allow containers to run on all major Linux distributions and Microsoft operating systems 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.

Running Docker on Windows Server

To install Docker on Windows Server, you can use a OneGet provider PowerShell module published by Microsoft called the *DockerMicrosoftProvider*. This provider enables the Containers feature in Windows and installs the Docker engine and client.

Note: If you plan to use Hyper-V isolation mode for your containers, you will also need to install the Hyper-V server role on the host server. If the host server is itself a virtual machine (VM), nested virtualization will need to be enabled before installing the Hyper-V role.

To install Docker on Windows Server, perform the following tasks:

1. Open an elevated PowerShell session and install the Docker-Microsoft PackageManagement Provider from the PowerShell Gallery:

```
Install-Module -Name DockerMsftProvider -Repository PSGallery -Force
```

2. Use the PackageManagement PowerShell module to install the latest version of Docker:

```
Install-Package -Name docker -ProviderName DockerMsftProvider
```

3. When installation of the Docker engine is complete, restart the computer.

Tip: You can verify the version of Docker that is installed by running the `docker version` command. If you need to install a specific version of Docker, add the `-RequiredVersion` switch when running the `Install-Package` command.

Running Docker on Windows 10

Windows 10 Professional and Enterprise edition supports the running and management of Windows containers using Docker. You can download and install the Docker Desktop, which provides the Docker Engine and other tools used to test and implement a containerized app.

Note: Windows containers use Hyper-V isolation mode on Windows 10 by default. To support this default mode, the Hyper-V feature must be installed on the Windows 10 computer.

Overview of Docker Desktop

The *Docker Desktop* is available for both Mac and Windows 10 desktop environments. This toolset enables you to build and distribute containerized applications and services. Docker Desktop includes the following:

- **Docker Engine.** This is a lightweight runtime environment for building and running Docker containers.
- **Docker Compose.** This tool 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 CLI client.** This tool includes a command shell that is preconfigured as a Docker command-line environment used to interface with the Docker Engine.
- **Kubernetes.** This tool is used to manage and orchestrate containerized applications across multiple hosts. It helps to deploy, maintain, and scale applications.
- **Credential Helper.** Used to help safely store Docker login credentials.

The Docker Hub

The *Docker Hub* is a web-based online library service in which you can:

- Register, store, and manage your own Docker images in an online repository and share them with others.
- Access over 100,000 container images from software vendors, open-source projects, and other community members.
- Download latest versions of the Docker Desktop.

Download container base images

After you install the Docker engine, the next step is to pull a base image, which is used to provide a foundational layer of operating system services for your container. You can then create and run a container, which is based upon the base image.

A container base image includes:

- The user mode operating system files needed to support the provisioned application.
- Any runtime files or dependencies required by the application.
- Any other miscellaneous configuration files needed by the app to provision and run properly.

Microsoft provides the following base images as a starting point to build your own container image:

- **Windows Server Core.** An image that contains a subset of the Windows Server application programming interfaces (APIs) such as the full .NET framework. It also includes most server roles.
- **Nano Server.** The smallest Windows Server image, with support for the .NET Core APIs and some server roles.
- **Windows.** Contains the full set of Windows APIs and system services; however, does not contain server roles.
- **Windows Internet of Things (IoT) Core.** A version of Windows used by hardware manufacturers for small IoT devices that run ARM or x86/x64 processors.

Note: The Windows host operating system version must match the container operating system version. To run a container based on a newer Windows build, you need to ensure that an equivalent operating system version is installed on the host. If your host server contains a newer operating system version, you can use Hyper-V isolation mode to run an older version of Windows containers. To determine the version of Windows installed, run the **ver** command from the command prompt.

The Windows container base images are discoverable through the **Docker Hub** and are downloaded from the **Microsoft Container Registry (MCR)**. You can use the *Docker pull* command to download a specific base image. When you enter the pull command, you specify the version that matches the version of the host machine.

For example, if you wanted to pull a Nano Server image based upon version 1903, you would use the following command:

```
docker pull mcr.microsoft.com/windows/nanoserver:1903
```

If you wanted to pull a 2019 LTSC Server core image, you would use the following command:

```
docker pull mcr.microsoft.com/windows/servercore:ltsc2019
```

After you download the base images needed for your containers, you can verify the locally available images and display metadata information by entering the following command:

```
docker image ls
```

Run a Windows container

With Docker, you can create, remove, and manage containers. You can also browse the Docker Hub to access and download prebuilt images. In most organizations, the most common management tasks that use Docker include:

- Automating the process of creating container images by using Dockerfile on a Windows OS.
- Managing containers by using Docker.

Automating container image creation by using Dockerfile on Windows

The Docker Engine includes tools for automating the process of creating container images. While you can create container images manually, adopting an automated image creation process provides many benefits, including:

- The ability to store container images as code.
- The rapid and precise re-creation of container images for maintenance and upgrade purposes.
- Continuous integration between container images and the development cycle.

The Docker components that drive this automation are the Dockerfile text file and the `docker build` command:

- **Dockerfile.** This text file contains the instructions needed to create a new container image. These instructions include the identification of an existing image to use as a base, commands to run during the image creation process, and a command that runs when new instances of the container image deploy.
- **docker build.** This Docker Engine command consumes a Dockerfile and then triggers the image creation process.

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 available images on your container host. As you might recall, you use container images as a base for new containers.
- **docker run.** This creates a container by using a container image. For example, the following command creates a container using the default process isolation mode, named IIS and based on the Windows Server Core container image:

```
docker run --name IIS -it windowsservercore
```

- **docker commit.** This commits the changes you made to a container to a new container image. The commit operation doesn't include data contained in volumes mounted within the container. Note that by default, the container will be paused as the new container image is created.
- **docker stop.** This stops a running container.
- **docker rm.** This removes an existing container.

Manage containers using Windows Admin Center

Windows Admin Center is a browser-based, graphical user interface (GUI) used to manage Windows servers, clusters, hyperconverged infrastructure (HCI), and Windows 10 PCs. It is used to provide a single administrative tool that can perform many of the tasks that were commonly performed using a variety of consoles, tools, or processes. Another powerful aspect of the Windows Admin Center is that it is extensible and supports third-party hardware manufacturers for configuration or monitoring capabilities related to their specific hardware.

After you install the Windows Admin Center, you might need to add additional extensions to allow you to manage the services you intend to use with the tool. You can add extensions by selecting the **Settings** button and then selecting **Extensions**. Windows Admin Center pulls the latest extension list from its default feed to display the available extensions that can be installed.

Overview of the Containers extension

The Windows Admin Center Containers extension is used to perform remote management and monitoring of the containers running on the targeted host machine. The Containers extension displays the following information and management tabs:

- **Summary.** Provides a summary of the status of the containers, including how many are running, which images are available, the networks that have been configured, and the volumes used.
- **Containers.** Provides an extensive amount of status information in both text and graphical format. You can also obtain details, logs, and events from specific containers and perform tasks such as End or Delete containers.
- **Images.** Displays and provides details on the images that are available in your local repository. You can also use this tab to delete images from the repository.
- **Networks.** Displays the networks that have been configured for the containers. Includes networking information such as subnet and gateway addresses used for the network connection.
- **Volumes.** Displays volume information used by containers on the host machine.

Demonstration: Deploy containers by using Docker

This demonstration illustrates how to:

- Install Docker on Windows Server.
- Download and run a Windows container.
- Use Windows Admin Center to manage containers.

Demonstration Steps

Install Docker on Windows Server

1. On SEA-ADM1, open Windows Admin Center, and then use Remote PowerShell to connect to SEA-SVR1.

2. At the PowerShell command prompt, enter the following command, and then select Enter:

```
Install-Module -Name DockerMsftProvider -Repository PSGallery -Force
```
3. At the PowerShell command prompt, enter the following command, and then select Enter:

```
Install-Package -Name docker -ProviderName DockerMsftProvider
```
4. After the installation is complete, restart the computer by using the following command:

```
Restart-Computer -Force
```

Download and run a Windows container

1. After SEA-SVR1 restarts, verify the installed version of Docker by using the following command:

```
Get-Package -Name Docker -ProviderName DockerMsftProvider
```
2. To display whether any Docker images are currently pulled, use the following command:

```
Docker images
```
3. To review docker base images from the online Microsoft repository, use the following command:

```
Docker search Microsoft
```
4. To download and run a sample hello-world test Nano Server image, run the following command:

```
docker container run hello-world:nanoserver
```
5. To download a Nano server image that matches the host operating system version of 1809, run the following command:

```
docker pull mcr.microsoft.com/windows/nanoserver:1809
```
6. Switch to SEA-SVR1.
7. To display the Docker images that are currently pulled, use the following command:

```
Docker images
```
8. To run the container in interactive mode, enter the following command:

```
Docker run -it --name NanoImage mcr.microsoft.com/windows/nanoserver:1809
```
9. To verify that you are focused on the container in interactive mode, enter the following command:

```
hostname
```
10. Switch back to SEA-ADM1, and then ensure that SEA-SVR1 is still connected to **Windows Admin Center**. In the Tools list select **Processes**. At the prompt, select **Continue**.
11. In the **Processes** list, take note of the **CExecSvc.exe** process. This is the *Container Execution service* which uses a named pipe to communicate with Docker and the VMcompute services on the host. This service indicates that the container is running in process isolation mode, which still uses the kernel from the host machine.
12. Switch to SEA-SVR1.
13. To close the interactive session with the container, use the following command:

```
Exit
```
14. To verify that the container has stopped, use the following command:

```
docker ps
```

15. To run the container in Hyper-V isolation mode, enter the following command:

```
Docker run -it --name NanoHVImage --isolation=hyperv mcr.microsoft.com/windows/nanoserver:1809
```

16. Switch back to SEA-ADM1. In the **Processes** list, notice that the **CExecSvc.exe** process is not running. This indicates that the container is running in Hyper-V isolation mode, which does not share processes with the host machine.

17. In the **Tools** list, select **PowerShell**. Provide the **Contoso\Administrator** credentials, and then select **Enter**.

18. In the remote PowerShell session, enter the following command:

```
docker ps
```

19. In the remote PowerShell session, enter the following command:

```
docker stop <ContainerID>
```

20. Rerun the `docker ps` command to confirm that the container has stopped.

Use Windows Admin Center to manage containers

1. On SEA-ADM1, ensure that SEA-SVR1 is targeted in the Windows Admin Center, and then select the **Containers** tool.
2. Browse through each of the **Summary**, **Containers**, **Images**, **Networks**, and **Volumes** tabs.

Test Your Knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

Describe the primary difference between a container and a virtual machine.

Question 2

Which container management provider is supported with Windows?

Question 3

Which container base image is used primarily to support .NET core APIs and is good to use if you want to have a very small base image starting point?

Question 4

What can you use to help automate container image creation and management?

Overview of Kubernetes

Lesson overview

Building and managing modern applications using containers requires methods and processes that allow efficient deployment of containers to both on-premises and cloud-based resources such as Microsoft Azure. Kubernetes is open-source orchestration software used to efficiently deploy, manage, and scale containers in a hosted environment.

In this lesson, you are introduced to the concept of Kubernetes and its benefits for managing container technology.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe container orchestration.
- Explain Kubernetes.
- List high-level steps for deploying Kubernetes resources.

What is Windows container orchestration?

Containers provide many benefits for managing applications and supporting agile delivery environments and microservice-based architectures. However, application components often grow to span multiple containers spread across multiple host servers. If you don't have automation processes in place, trying to manage and operate these as distributed architecture in an efficient and scalable manner can be difficult.

Automating the processes within a containerized environment is the job of an *orchestrator*. The orchestrator is used to automate and manage large numbers of containers and control how the containers interact with one another.

A typical orchestrator performs the following tasks:

- **Scheduling:** Finds a suitable machine on which to run the container when given a container image and a resource request.
- **Affinity/Anti-affinity:** Specifies whether a set of containers should run near each other for performance or far apart for availability.
- **Health monitoring:** Watches for container failures and automatically reschedules them.
- **Failover:** Keeps track of what's running on each machine and reschedules containers from failed machines to healthy nodes.
- **Scaling:** Manually or automatically adds or removes container instances to match demand.
- **Networking:** Provides an overlay network that coordinates containers to communicate across multiple host machines.
- **Service discovery:** Enables containers to locate each other automatically even as they move between host machines and change IP addresses.
- **Coordinated application upgrades:** Manages container upgrades to avoid application downtime and enables rollback if something goes wrong.

Types of container orchestration tools

Several container orchestration tools are available and are used depending upon the specific needs of the architecture to be managed. Common orchestration tools include:

- **Kubernetes:** Considered the main standard for container orchestration, Kubernetes is an open-source platform used for deploying and managing containers at scale. Note that Kubernetes is often abbreviated to *K8s* ('8' represents the eight characters between the K and the s of the word Kubernetes).
- **Docker Swarm:** Docker's own fully integrated container orchestration tool. Considered less extensible and complex than Kubernetes, it is a good choice for Docker-specific enthusiasts. Docker bundles both Swarm and Kubernetes with the Docker Desktop.
- **Apache Mesos:** Open source software that can provide management of a container cluster. Requires additional add-on frameworks to support full orchestration tasks.

Kubernetes is typically the clear standard for container orchestration. Most cloud providers offer Kubernetes-as-a-service to help manage the deployment and management of containerized applications for you. For example, **Azure Kubernetes Service (AKS)** integrates with Azure Container Registry (ACR) and provides its own provisioning portal where you can secure your container clusters with Azure's Active Directory and deploy apps across Azure's datacenter offerings. By using AKS, you can take advantage of the enterprise-grade features of Azure while still maintaining application portability through Kubernetes and the Docker image format.

Overview of Kubernetes on Windows

Kubernetes architecture

Kubernetes is based upon common cluster technology in which a set of resources works together and is managed as a single system performing similar kinds of tasks.

A *cluster* consists of centralized software, also known as the *master* or *control plane*, that is responsible for scheduling and controlling tasks within the cluster. The components that run the tasks within a cluster are called *nodes*. A cluster typically contains multiple nodes, which are managed by the control plane.

A Kubernetes cluster contains a Master/Control plane and a combination of either Windows or Linux-based worker node instances, described as follows:

- **At least one Master/Control plane:** Runs a number of services used to manage the orchestration in Kubernetes. Currently only the Linux operating system is supported as the host operating system for the Kubernetes master. Components that make up the master include *kube-api-server*, *controller*, and *scheduler* services.
- **One or more Linux-based node instances:** Worker nodes that are based upon the Linux operating system. Includes components called the Kubelet, Kube-proxy, and the Container runtime services.
- **One or more Windows-based node instances:** Worker nodes that are based upon the Windows Server 2019. Includes components called the Kubelet, Kube-proxy, and the Container runtime services.

Kubernetes pods

A Kubernetes workload is typically made up of several Docker-based containers that are often disbursed throughout multiple worker nodes within the cluster. A Kubernetes object called a *pod* is used to group one or more containers to represent a single instance of an application.

Note: A single pod can hold one or more containers; however, a pod usually does not contain multiple versions of the same application.

A pod includes information about the shared storage and network configuration, and a specification on how to run its packaged containers. You use pod templates to define the information about the pods that run in your cluster.

Deploy Kubernetes resources

As of Kubernetes 1.14, support for Windows-based nodes are supported for both worker nodes and scheduling Windows containers. Windows-based worker nodes must be on Windows Server 2019 and Windows Server version 1809 or later.

The process for creating a Kubernetes orchestration solution includes the following general steps:

1. **Create a Kubernetes master.** The Kubernetes master is configured using the Linux operating system. You use a tool called the *kubeadm* tool which initializes the master and provides all the associated tools used to administer cluster nodes. You will also need to install Docker to allow for container support.
2. **Configure a network solution.** The network solution typically used to create routable cluster subnets is a Linux-based Container Network Interface (CNI) plugin called *Flannel*. Other possible solutions include configuring a smart top-of-rack (ToR) switch or using Open vSwitch (OvS) or Open Virtual Network (OVN) technology.
3. **Join worker nodes.** After creating the Kubernetes master and configuring the network solution, you can then join Windows Server and Linux-based worker nodes to the cluster. Joining the worker nodes consists of configuring operating system-specific binaries and connecting to the network solution.
4. **Manage Kubernetes resources.** You use the *Kubectl* command to deploy and manage Kubernetes pods containing the containers which make up the application.

Tip: Cloud services such as **Azure Kubernetes Service (AKS)** reduce many of the challenges of manually configuring Kubernetes clusters by providing a hosted Kubernetes environment. These services also simplify the deployment and management of containerized applications in Azure. With AKS, you get the benefits of open-source Kubernetes without the complexity or operational overhead of running your own custom Kubernetes cluster.

Test Your Knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

Describe three tasks that a typical container orchestrator performs.

Question 2

Describe the primary components of a Kubernetes cluster.

Question 3

Which Microsoft cloud-based service can be used to provide a hosted Kubernetes environment?

Module 05 lab and review

Lab: Implementing and configuring virtualization in Windows Server

Scenario

Contoso is a global engineering and manufacturing company with its head office in Seattle, USA. An IT office and data center are in Seattle to support the Seattle location and other locations. Contoso recently deployed a Windows Server 2019 server and client infrastructure.

Because of many physical servers being currently underutilized, the company plans to expand virtualization to optimize the environment. Because of this, you decide to perform a proof of concept to validate how Hyper-V can be used to manage a virtual machine environment. Also, the Contoso DevOps team wants to explore container technology to determine whether they can help reduce deployment times for new applications and to simplify moving applications to the cloud. You plan to work with the team to evaluate Windows Server containers and to consider providing Internet Information Services (Web services) in a container.

Objectives

After completing this lab, you'll be able to:

- Create and configure VMs.
- Install and configure containers.

Estimated Time: 60 minutes

Module review

Use the following to check what you've learned in this module.

Question 1

Which of the following are requirements for installing the Hyper-V server role in Windows Server? Choose two.

- A 32-bit processor
- Minimum 32 GB of memory
- A 64-bit processor
- BitLocker enabled
- Intel VT or AMD-V enabled

Question 2

You plan to enable nested virtualization on a Hyper-V host. What do you need to do to ensure that network traffic of nested VMs can reach an external network?

- Enable BitLocker
- Enable MAC address spoofing
- Enable Device Guard
- Configure a switch with the Internal Network type
- Configure a switch with the Private Network type

Question 3

Which of the following are true for considerations when implementing a Host Guardian service? Choose two.

- A new Active Directory forest is created dedicated to the Host Guardian service.
- The Host Guardian service must be installed on a server containing the Linux operating system.
- The Host Guardian service must be installed in a virtual machine.
- The Host Guardian service uses certificates for signing and encryption tasks.
- The Host Guardian service must be installed in the same domain as the Hyper-V guarded hosts.

Question 4

Which of the following are requirements for creating a shielded template disk? Choose two.

- A generation 2 virtual machine
- A basic disk
- A generation 1 virtual machine
- A dynamic disk
- Must be generalized

Question 5

You download a container base image. When you attempt to create and run a container using the base image, you get an error message that relates to incompatibility with the host machine. What should you do?

- Download a new container base image that matches the version of the operating system installed on the host machine.
- Run the container using the `--isolation=process` switch.
- Update the version of Docker installed on the host machine.
- Install a self-signed authentication certificate on the host machine.
- Use BitLocker to encrypt the Operating system drive of the host machine.

Question 6

Which of the following can be used as worker nodes in a Kubernetes cluster? Choose two.

- Nano Server
- Windows Server 2019
- MacOS
- Linux

Answers

Question 1

What is the correct term for the virtualization layer that is inserted into the boot process of the host machine that controls access to the physical hardware?

*A software layer known as the **hypervisor** is inserted into the boot process. The hypervisor is responsible for controlling access to the physical hardware.*

Question 2

Name four methods for managing Hyper-V virtual machines.

Four methods include Hyper-V Manager, Windows PowerShell, PowerShell Direct, and Windows Admin Center.

Question 3

What is the PowerShell command for enabling nested virtualization?

```
Set-VMProcessor -VMName <VMName> -ExposeVirtualizationExtensions $true
```

Question 1

You need to create a virtual machine (VM) that supports Secure boot. Which generation would you choose when you create the VM?

You need to select generation 2, which is needed to support Secure boot.

Question 2

Which virtual hard disk (VHD) type only uses the amount of space that needs to be allocated and grows in size as more space is necessary?

Dynamically expanding VHD.

Question 3

Which Hyper-V virtual switch allows communication between the VMs on a host computer and also between the VMs and the host itself only?

The Internal network switch

Question 4

You need to preserve the state and configuration of a VM at a set time period. What can you do?

You can create a checkpoint to preserve the state and configuration of a VM at a set time period.

Question 1

Describe three main benefits of running protected virtual machines (VMs) in a guarded fabric.

Benefits include securing an authorized Hyper-V host, verification that a host is in a healthy state, and providing a secure method to release keys to healthy hosts to allow for unlocking and starting a protected VM.

Question 2

Which component in a guarded fabric is used to enforce security and manage the keys to start protected VMs?

The Host Guardian Service.

Question 3

Describe three types of VMs that can be run in a guarded fabric.

A shielded VM, an encryption-supported VM, and a normal VM.

Question 4

Which tool is used to prepare and encrypt a VM template disk?

The Shielded Template Disk Creation Wizard, which is part of the Shielded VM Tools available from the Remote Administration Tools feature.

Question 1

Describe the primary difference between a container and a virtual machine.

A container shares the kernel with the host operating system and other containers. A virtual machine is totally isolated and has its own kernel and user mode.

Question 2

Which container management provider is supported with Windows?

Docker containers are fully supported by the latest releases of the Windows operating system.

Question 3

Which container base image is used primarily to support .NET core APIs and is good to use if you want to have a very small base image starting point?

The Nano Server container base image is the smallest images and has support for the .NET Core APIs.

Question 4

What can you use to help automate container image creation and management?

A Dockerfile is used to automate tasks, which contains instructions on how to create a new container.

Question 1

Describe three tasks that a typical container orchestrator performs.

Tasks may include scheduling, affinity/anti-affinity, health monitoring, failover, scaling, networking, service discovery, and coordinated application upgrades.

Question 2

Describe the primary components of a Kubernetes cluster.

A Kubernetes cluster contains at least one Master/Control plane and one or more Linux or Windows-based worker nodes.

Question 3

Which Microsoft cloud-based service can be used to provide a hosted Kubernetes environment?

The Azure Kubernetes Service (AKS).

Question 1

Which of the following are requirements for installing the Hyper-V server role in Windows Server? Choose two.

- A 32-bit processor
- Minimum 32 GB of memory
- A 64-bit processor
- BitLocker enabled
- Intel VT or AMD-V enabled

Explanation

To install the Hyper-V server role, you need a 64-bit processor with second-level address translation (SLAT). You also need to enable Intel VT or AMD-V. You also must have a processor with VM Monitor Mode extensions and must enable Hardware-enforced Data Execution Prevention (DEP).

Question 2

You plan to enable nested virtualization on a Hyper-V host. What do you need to do to ensure that network traffic of nested VMs can reach an external network?

- Enable BitLocker
- Enable MAC address spoofing
- Enable Device Guard
- Configure a switch with the Internal Network type
- Configure a switch with the Private Network type

Explanation

To enable network packets to be routed through two virtual switches, you must enable MAC address spoofing on the physical Hyper-V host.

Question 3

Which of the following are true for considerations when implementing a Host Guardian service? Choose two.

- A new Active Directory forest is created dedicated to the Host Guardian service.
- The Host Guardian service must be installed on a server containing the Linux operating system.
- The Host Guardian service must be installed in a virtual machine.
- The Host Guardian service uses certificates for signing and encryption tasks.
- The Host Guardian service must be installed in the same domain as the Hyper-V guarded hosts.

Explanation

The Host Guardian Service (HGS) can be run on physical or virtual machines. The HGS can run on Windows Server 2019 or Windows Server 2016 Standard or Datacenter editions. The HGS will set up the server in a new AD DS forest dedicated so that HGS ensures sensitive key information is as secure as possible. The Hyper-V guarded hosts are installed in the standard AD DS environment.

Question 4

Which of the following are requirements for creating a shielded template disk? Choose two.

- A generation 2 virtual machine
- A basic disk
- A generation 1 virtual machine
- A dynamic disk
- Must be generalized

Explanation

When creating a shielded template disk, the disk must be Basic and cannot be dynamic because BitLocker does not support dynamic disks. The operating system also needs to be generalized, which can be done using `sysprep.exe`.

Question 5

You download a container base image. When you attempt to create and run a container using the base image, you get an error message that relates to incompatibility with the host machine. What should you do?

- Download a new container base image that matches the version of the operating system installed on the host machine.
- Run the container using the `--isolation=process` switch.
- Update the version of Docker installed on the host machine.
- Install a self-signed authentication certificate on the host machine.
- Use BitLocker to encrypt the Operating system drive of the host machine.

Explanation

The Windows host operating system version needs to match the container operating system version. To run a container based on a newer Windows build, you need to ensure that an equivalent operating system version is installed on the host. Note that if your host server contains a newer operating system version, you can use Hyper-V isolation mode to run an older version of Windows containers.

Question 6

Which of the following can be used as worker nodes in a Kubernetes cluster? Choose two.

- Nano Server
- Windows Server 2019
- MacOS
- Linux

Explanation

Windows Server 2019 and Linux are both supported as worker nodes in a Kubernetes cluster.



Module 6 High availability in Windows Server

Planning for failover clustering implementation

Lesson overview

Failover clusters in Windows Server provide a high-availability solution for many server roles and applications. By implementing failover clusters, you can maintain application or service availability in the event that one or more computers in a cluster fails. Before implementing failover clustering, you should be familiar with general high-availability concepts. You must be familiar with clustering terminology and understand how failover clusters work, and you must be familiar with the new clustering features in Windows Server 2019.

Lesson objectives

After completing this lesson, you'll be able to:

- Describe failover clustering.
- Explain high availability with failover clustering.
- Describe clustering terminology.
- Describe failover clustering components.
- Explain cluster quorum in Windows Server.
- Discuss considerations for planning failover clustering.

What is failover clustering?

A *failover cluster* is a group of independent computers that work together to increase the availability and scalability of clustered roles, formerly called *clustered applications and services*. Clustered servers, called *nodes*, are connected by physical cables and by software. If one or more cluster nodes fail, other nodes begin to provide service in a process known as *failover*. The clustered roles are proactively monitored to verify that they're working properly. If they're not working, they're restarted or moved to another node.

Additionally, failover clusters provide Cluster Shared Volume (CSV) functionality that provides a consistent, distributed namespace that clustered roles use to access shared storage from all nodes. By using failover clustering, users experience a minimum of disruptions in service.

Failover clustering has many practical applications, including:

- Highly available or continuously available file share storage for applications such as Microsoft SQL Server and Hyper-V virtual machines (VMs).
- Highly available clustered roles that run on physical servers or on VMs that are installed on servers running Hyper-V.

Failover clustering improvements in Windows Server 2019

Microsoft has made several investments for failover clustering in Windows Server 2019. Some of the introductions include failover clustering integration into the Windows Admin Center, cluster performance history, system insights, presentient history, cluster sets, the in-place upgrade of the cluster nodes, and cross-domain cluster migration.

There have been many improvements over the Windows Server 2016 release with two-node hyperconverged infrastructures (HCI), Microsoft Distributed Coordinator, Scale-Out File Servers, Cluster Aware Updating Enhancements, and Enhancements to the naming of Cluster Networks.

Windows Admin Center

The Windows Admin Center, while not just for failover clustering, offers a browser-based management tool that allows you to manage your Windows Servers with no Azure or cloud dependencies. Although you can manage failover cluster nodes as individual servers, you can use the Windows Admin Center to add the failover clusters to a view and manage your cluster, storage, network, nodes, roles, VMs, and virtual switch resources. To add a failover cluster to the Windows Admin Center, you choose to add a Failover Connection through the UI. Additionally, you can also manage hyperconverged clusters by adding the cluster as a hyperconverged cluster connection.

The biggest advantage to the new Windows Admin Center is that it is one user interface for everything. When you use the Failover Cluster Manager in Windows Server 2019, you are prompted to try managing your clusters with the Windows Admin Center. The Windows Admin Center tool gives you the ability to examine your Cluster Performance History on how the cluster/nodes are performing; it offers System Insights using machine learning (ML) and predictive analytics (AI), and *persistent memory* (PMEM) which is the latest in storage and memory technology.

Cluster Performance History

Cluster Performance History is a new feature in Windows Server 2019 that gives Storage Spaces Direct administrators easy access to historical compute, memory, network, and storage measurements across your organization's servers, drives, volumes, VMs, and many other resources. The performance history is automatically collected and stored on the cluster for up to a year. You can examine the metrics—which are aggregated for all the servers in the cluster—by using Windows PowerShell's **Get-ClusterPerf** alias, which calls the **Get-ClusterPerformanceHistory** cmdlet, or by using the Windows Admin Center.

Windows Server System Insights

New to Windows Server 2019 is the release of a feature called *Windows Server System Insights* (System Insights). This new feature is a machine learning (ML) and predictive analytics (AI) tool that analyzes data

to give insight into the functioning of your server deployments. System Insights collects, persists, and analyzes your server data locally on the Windows Server machine; this data might be forwarded to Azure Log Analytics (OMS) if you want to have a unified view of your entire environment. By default, System Insights offers CPU capacity forecasting, total storage consumption forecasting, networking capacity forecasting, and volume consumption forecasting. You can manage System Insights by using Windows PowerShell or Windows Admin Center. The Windows PowerShell cmdlet for System Insights is **Get-InsightsCapability**. After running this cmdlet, you will have the option to either use PowerShell to enable or disable the features you wish.

Persistent memory

Windows Server 2019 also introduced PMEM, which uses a new type of memory technology that delivers a combination of capacity and persistence. It is essentially super-fast storage on a memory stick. PMEM is deployed using Storage Spaces Direct.

Cluster sets

Cluster sets are new in Windows Server 2019 and offer cloud scale-out technology. Cluster sets dramatically increase cluster node count in a single Software Defined Data Center (SDDC). A cluster set is a loosely coupled grouping of multiple failover clusters that enables VM fluidity across member clusters within a cluster set and a unified storage namespace across the set. Cluster sets preserve existing failover Cluster management experiences on member clusters and offer key use cases around forecasting lifecycle management. You can implement and maintain cluster sets using Windows PowerShell. In addition to offering VM fluidity, cluster sets can place new VMs on an optimal node, allow movement of VMs between clusters without changing their storage paths, and take advantage of the new **Infrastructure Scale Out File Server** role and Server Message Block (SMB) Loopback, a new SMB feature.

In-place upgrade of cluster nodes

Windows Server 2019 is the first version where Microsoft truly supports upgrading the operating system of a cluster node in-place. In previous versions, administrators were required to drain the node, evict it, install everything fresh, and then add the node back to the cluster. The cluster rolling upgrade process that was introduced in Windows Server 2016 remains the same. You can upgrade Windows Server up to two versions, so if you are currently on Windows Server 2012 R2, you can upgrade to Windows Server 2019. From a cluster perspective in a rolling-up fashion, you can only upgrade one version at a time.

Improvements to Microsoft Distributed Coordinator

With Windows Server 2019, Microsoft has added support for MSDTC on Cluster Shared Volumes (CSV). MSDTCs can run directly on Storage Spaces Direct clusters; they no longer need to run on shared disks. This is a tremendous feature for organizations looking to store SQL Server Failover Cluster Instances (FCI) on Storage Space Direct clusters. Also, Azure now supports FCI.

Improvements to two-node HCI

Windows Server 2019 now offers two-node HCI. HCI originally launched with Windows Server 2016, but 2019 has enhancements such as resilience to a loss of a node plus one additional drive. The HCI can be set up as a nested two-way mirror which uses one data partition and three copies, where all writes go to all copies and reads can come from any copy. Another HCI setup option is the nested mirror accelerate parity model that combines nested mirroring with nested parity. This mode will rotate data between

mirror and parity in real time. Each of these modes comes at a cost when you examine that a normal two-way mirror costs you about 50% of your space, and nested two-way mirrors will use about 75% of your space. The nested mirror-accelerated parity offers a nice balance, consuming 60% of your space.

In situations where internet or domain controller access cannot be guaranteed, Windows Server 2019 now supports using a USB memory device as a witness. This means you no longer need a Cluster Name Object (CNO), Kerberos, a domain controller, certificates, or even an account on the nodes. You now can plug a USB into your network router and have a very simple cluster witness. This can all be set up by using the **Set-ClusterQuorum** cmdlet.

Failover Cluster no longer uses NTLM authentication

Windows Server 2019 offers new security enhancements in relation to failover clustering, supporting clusters without NTLM dependencies. Microsoft has moved away from NTLM in favor of certificate-based intra-cluster SMB authentication.

Improvements to Scale-Out File Servers

Windows Server 2019 has several improvements to the Scale-Out File Servers (SOFS) infrastructure and to the method in which SMB works with the SMB local loopback. In Windows Server 2019, identity tunneling on infrastructure shares is applied when using SOFS. This means when you access infrastructure shares from the same cluster or cluster set, the application token is serialized and tunneled to the server, and then the VM access is completed using the provided token.

Azure-aware Clusters

With Azure-aware clusters, you can detect issues during creation of a cluster if it is on Azure IaaS. Windows Server 2019 allows you to optimize configurations for Azure IaaS (new property). There is no longer the need for the internal load balancer IP address to connect to the cluster.

Note: Any roles that are required will still need the internal load balancer.

Azure host maintenance

In Windows Server 2019, the experience of Azure host maintenance is improved. It includes monitoring for Azure Scheduled Events, which makes troubleshooting device disconnects easier to detect.

Cross-domain Cluster Migration

With Windows Server 2019, you can migrate clusters from one domain to another using a series of Windows PowerShell scripts and without destroying your original cluster. The scripts allow the process to dynamically change the NetNames AD integration and whether or not the server is joined to a domain or a workgroup. This new process keeps the cluster intact.

Cluster Aware updating enhancements

Cluster aware updating is enhanced in Windows Server 2019 in the following ways:

- Windows Server 2019 now checks to be sure that Storage Spaces Direct storage jobs are making progress and waits before skipping them.

Note: You should test, force, drain, and reboot, even if no updates are available for troubleshooting and testing purposes.

- The **Save-CauDebugTrace** cmdlet collects the Windows Update Logs.
- The wait time has been extended for machines needing longer boot times.
- The server no longer drains unless an update requires a reboot.

Enhancements to the naming of Cluster Networks

With the introduction of Windows Server 2019, you now have the option to use a distributed network name for the cluster much like you can for SOFS. A distributed network name uses the IP addresses of the member servers instead of requiring a dedicated IP address for the cluster.

By default, Windows uses a distributed network name if it detects that you are creating the cluster in Azure; this removes the need to create an internal load balancer for the cluster. Use the **New-Cluster** cmdlet to take advantage of the naming enhancement.

High availability with failover clustering

Availability refers to a level of service that applications, services, or systems provide. It's 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 organization's needs and budget. For example, a system that's unavailable for 8.75 hours per year would have a 99.9 percent availability rating.

To improve availability, you should implement mechanisms that mitigate the effect of failing components and dependencies. You can achieve fault tolerance by implementing redundancy to single points of failure.

The costs of high availability don't just affect an organization financially—high availability also has more central processing unit (CPU), memory, and storage requirements. This might also equate to infrastructure and continued yearly support costs. The customer might have incorrect service-level expectations that could result in poor business decisions, such as unsuitable investment levels, and customer dissatisfaction. As an administrator working in the IT department, you must disclose the availability requirements to the customer to avoid possible misunderstandings.

Another factor that might significantly affect the definition of high availability is the availability measurement period. For example, a requirement for 99.9 percent availability over a one-year period allows for 8.75 hours of downtime for the entire year, while a requirement for 99.9 percent availability over a rolling four-week period allows for only 40 minutes of downtime every four weeks.

You must also identify and negotiate planned outages, maintenance activities, and software updates. For the purpose of defining service-level agreements, these are scheduled outages and usually aren't included as downtime when calculating a system's availability, because you typically calculate availability based only on unplanned outages. However, you must negotiate exactly which planned outages you consider to be downtime.

Clustering terminology

The following three tables contain clustering terms and definitions with which you should be familiar.

Infrastructure terminology

Table 1: Clustering terminology - infrastructure

Term	Term definition
Active node	An active node has a cluster currently running on it. A resource or resource group can only be active on one node at a time.
Cluster resource	A cluster resource is a hardware or software component in the cluster such as a disk, virtual name, or IP address.
Cluster sets	Use Cluster sets to scale out your topology by using the cloud. Cluster sets enable virtual machine (VM) fluidity across member clusters within a cluster set and a unified storage namespace across the set.
Node	A node is an individual server in a cluster.
Passive node	A passive node doesn't have a cluster currently running on it.
Public network or private network	Each node needs two network adapters: one for the public network and one for the private network. The public network is connected to a local area network (LAN) or a wide area network (WAN). The private network exists between nodes and is used for internal network communication, which is called the heartbeat.
Resource group	A resource group is a single unit within a cluster that contains cluster resources. A resource group is also called an application and service group.
Virtual server	A virtual server consists of the network name and IP address to which clients are connected. A client can connect to a virtual server, which is hosted in the cluster environment, without knowing the details of the server nodes.

Failover terminology

Table 2: Clustering terminology - failover

Term	Term definition
Azure Cloud Witness	In Windows Server, you can use a Microsoft Azure Cloud Witness share to create a quorum. Previously with Windows Server 2012 R2, when creating a stretch cluster, a third offsite quorum was recommended. With Windows Server, you can create an Azure Cloud Witness instead.
Cluster quorum	The cluster quorum maintains the definitive cluster configuration data and the current state of each node. It also maintains each service, application group, and resource network in the cluster.

Term	Term definition
Cluster Shared Volumes (CSVs)	CSVs in Windows Server provide support for a read cache, which can significantly improve performance in certain scenarios. Additionally, a CSV File System can perform chkdsk without affecting applications with open handles on the file system.
Heartbeat	The heartbeat is a health check mechanism of the cluster, where a single User Datagram Protocol (UDP) packet is sent to all nodes in the cluster through a private network to check whether all nodes in the cluster are online. One heartbeat is sent every second. By default, the cluster service will wait for five seconds before the cluster node is considered unreachable.
Private storage	Local disks are referred to as private storage.
Shared disk	Each server must be attached to external storage. In a clustered environment, data is stored in a shared disk that's accessible by only the nodes in the system.
Storage Replica	Storage Replica provides disaster recovery by enabling block-level, storage-agnostic, synchronous replication between servers. You can use Storage Replica in a wide range of architectures, including stretch clusters, cluster-to-cluster, and server-to-server.
witness disk or file share	The cluster witness disk or the witness file share are used to store the cluster configuration information. They help to determine the state of a cluster when some or all of the cluster nodes can't be contacted.

Tools and features terminology

Table 3: Clustering terminology - tools and features

Term	Term definition
Cluster Performance History	Cluster Performance History is a new feature in Windows Server 2019 that gives Storage Spaces Direct administrators easy access to historical compute, memory, network, and storage measurements across an organization's servers, drives, volumes, VMs, and many other resources. The performance history is automatically collected and stored on a cluster for up to a year. The metrics are aggregated for all the servers in the cluster, and they can be examined by using the Windows PowerShell Get-ClusterPerf alias, which calls the Get-ClusterPerformanceHistory cmdlet, or by using Windows Admin Center.
Cross-domain cluster migration	With Windows Server 2019 cross-domain cluster migration, you can migrate clusters from one domain to another by using a series of Windows PowerShell scripts without destroying your original cluster. The scripts allow you to dynamically change the NetNames Active Directory integration and change to and from domain joined to a workgroup and vice versa.
Persistent memory	Windows Server 2019 introduced persistent memory, or PMEM, which offers a new type of memory technology that delivers a combination of capacity and persistence. Essentially, PMEM is super-fast storage on a USB flash drive. PMEM deploys by using Storage Spaces Direct.
System Insights	The System Insights feature of Windows Server provides machine learning and predictive analytics to analyze data on your servers.
Windows Admin Center	Windows Admin Center offers a browser-based management tool that you can use to manage Windows Server computers with no Azure or cloud dependencies. You can use Windows Admin Center to add failover clusters to a view and to manage your cluster, storage, network, nodes, roles, VMs, and virtual switch resources.

Additional reading: For more information on persistent memory, refer to **Understand and deploy persistent memory**¹.

Additional reading: For more information on CSVs, refer to **Use Cluster Shared Volumes in a failover cluster**².

¹ <https://aka.ms/deploy-pmem>

² <https://aka.ms/failover-cluster-csvs>

Failover clustering components

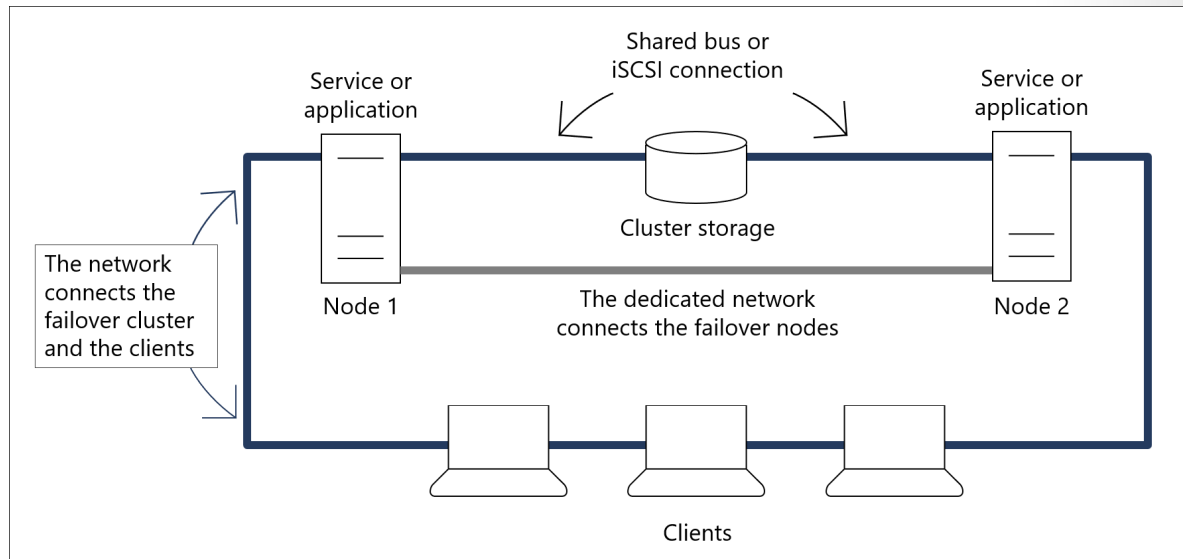


Figure 1: Two-node failover cluster

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, a process known as *failover*. Using failover ensures that users experience a minimum amount of service disruptions.

Components of a failover clustering solution

A failover clustering solution consists of several components, including the following:

- **Nodes.** *Nodes* are computers that are members of a failover cluster. These computers run the Cluster service and any resources and applications associated with the cluster.
- **A network.** Cluster nodes can communicate with one another and with clients across a *network*.
- **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 Microsoft Exchange Server, shared storage isn't required.
- **Clients.** *Clients* are computers or users that use the Cluster service.
- **A service or application.** A *service or application* is a software entity that Microsoft presents to clients and that clients use.
- **A witness.** In Windows Server, a *witness* can be a file share, disk, Azure Cloud witness, or in the case of Storage Spaces Direct, a USB device can be a witness. You use a witness to maintain a quorum. Ideally, a witness is on a network that's both logically and physically separate from those that the failover cluster uses. However, a witness must remain accessible by all cluster node members. Later lessons discuss the concepts of quorum and witnesses in more detail.

Failover cluster nodes

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 SCSI (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.

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 a node detects 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 a failure, client requests automatically redirect to an alternative node in the cluster as quickly as possible.

Cluster storage and networks

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. 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 is an isolated network that allows cluster node members to communicate directly with one another. If directly connected shared storage isn't in use, a third network segment (for iSCSI or Fibre Channel) can exist between the cluster nodes and a data storage network.

Cluster quorum in Windows Server

What is a quorum?

A *quorum* is the majority of voting nodes in an active cluster membership plus a witness vote. In effect, each cluster node is an element that can cast one vote to determine whether the cluster continues to run. If an even number of nodes exists, another element, referred to as a *witness*, is assigned to the cluster. The witness element can be a disk, a file share, or a Microsoft Azure Cloud Witness. Each voting element has a copy of the cluster configuration, and the Cluster service works to always keep all the copies synced.

The cluster stops providing failover protection if more than half 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.

A quorum prevents two or more nodes from concurrently operating as a failover cluster resource. If you don't achieve a clear majority among the node members, the vote of the witness becomes crucial to maintain the validity of the cluster. Concurrent operations 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 data-

base became available on a 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 required votes for the cluster to continue providing failover protection. If the number of votes drops below a majority, the cluster will stop running. That is, it won't 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 won't function as a cluster until a majority consensus occurs or they 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 other resources that might fail over to it.

Achieving quorum

A cluster must complete several phases to achieve a quorum. After a node starts running, 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. The nodes determine if this collection of members has a quorum or enough members to create enough votes such that a split scenario can't exist. A *split scenario* means that another set of nodes in this cluster run on a part of the network that's inaccessible to these nodes.

Therefore, more than one node might be actively trying to provide access to the same clustered resource. If enough votes don't exist to achieve a quorum, the voters (the currently recognized members of the cluster) wait for more members. After reaching at least the minimum vote total, the Cluster service begins to bring cluster resources and applications into service. After reaching a quorum, the cluster becomes fully functional.

Quorum modes in Windows Server failover clustering

In Windows Server, the process and recommendations for configuring a quorum have changed. As before, votes determine whether a cluster achieves a quorum. Nodes can vote and, where appropriate, so can a disk in cluster storage (known as a *witness disk*), a file share (known as a *file share witness*), or an Azure Cloud Witness. The following are the available quorum mode configurations:

- **Node majority.** Each available and communicating node can vote. The cluster functions only with a majority of the votes. This model is preferred when a cluster consists of an odd number of server nodes; no witness is necessary to maintain or achieve a quorum.
- **Node and disk majority.** Each node plus a designated disk in the cluster storage (the witness disk) can vote when they're available and in communication. The cluster functions only with a vote majority. This model is based on an even number of server nodes being able to communicate with one another in the cluster and with the witness disk.
- **Node and file share majority.** Each node plus a designated file share that an administrator created, the file share witness, can vote when they're available and in communication. The cluster functions only with a vote majority. This model is based on an even number of server nodes being able to communicate with one another in the cluster and with the file share witness.

- **No majority:** In a disk-only scenario, the cluster has a 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

The dynamic quorum mode dynamically adjusts the quorum votes based on the number of servers that are online. For example, assume that you have a five-node cluster, you place two of the nodes in a paused state, and then one of the remaining nodes fails. In any of the earlier configurations, the cluster would fail to achieve a quorum and would go offline. However, a 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. A cluster with a dynamic quorum stays online.

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 doesn't 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 adjust voting to always have an odd number of votes in a 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.

You can choose whether to use a witness disk, file share witness, or Azure Cloud Witness:

- **Witness disk.** A witness disk is still the primary witness in most scenarios, especially for local cluster 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 witness disk.
- **File share witness.** A file share witness is ideal when shared storage isn't available or when the cluster spans geographical locations. This option doesn't store a copy of the cluster database.
- **Azure Cloud Witness.** The Azure Cloud Witness is the ideal option when you run internet-connected stretch clusters. This removes the need to set up a file share witness at a third datacenter location or a virtual machine in the cloud. Instead, this option is built into a failover cluster. This doesn't store a copy of the cluster database.
- **In Windows Server 2019,** you can now create a file share witness that doesn't utilize the cluster name object, but in fact, simply uses a local user account on the server to which the file share witness is connected. This means that you no longer need Kerberos authentication, a domain controller, a certificate, or even a cluster name object. You also no longer need an account on the nodes.

You should also consider the capacity of the nodes in a 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 witness disk 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.

Considerations for planning failover clustering

Planning considerations

Before implementing failover clustering technology, you must pick the services and applications that you want to make highly available. You can't 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 virtual machines (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 doesn't 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 enough capacity to service the highly available services or applications that you allocate to it when another node fails. This capacity should provide enough of a buffer to avoid nodes that run at near capacity after a failure event. Failing to adequately plan resource utilization can result in decreased 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 operational 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's important that the standby servers have enough capacity to run the load from more than one node failure.

You should also examine all cluster configuration components to identify single points of failure. You can remedy many single points of failure with simple solutions, such as adding storage controllers to separate and stripe disks, 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 to use multiple power supplies to provide power redundancy and options to create 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:

- The hardware should be certified for Windows Server.
- The same or similar hardware must be installed 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 serial attached SCSI (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 should also use the same firmware version.
- If you use Internet SCSI (iSCSI) storage connections, each clustered server should have one or more network adapters or host bus adapters that are dedicated to the cluster storage. You shouldn't use the network that you use for iSCSI storage connections for non-storage network communication. In all clustered servers, the network adapters that you use to connect to the iSCSI storage target should be identical, and you should use Ethernet that's 10-gigabit or more.
- After configuring the servers with hardware, all the tests in the **Validate a Configuration Wizard** must pass before the cluster is considered a configuration that Microsoft supports.

- Each node must run the same processor architecture. This means that each node must have the same processor family.

Network component requirements for failover cluster implementation

In addition to hardware requirements, the components that you use must meet certain requirements and pass the tests in the **Validate a Configuration Wizard**. The suggested network requirements include the following:

- The network adapters in each node should be identical and have the same IP version, speed, duplex, and flow control capabilities.
- The networks and network equipment with which you connect the nodes should be redundant so that if a failure occurs, the nodes can continue communicating with one another. You can team network devices to provide single network redundancy. You should use multiple networks to provide multiple paths among nodes for internode communication; otherwise, a warning message will appear during the validation process.
- The network adapters in a cluster network should have the same IP address assignment method, such as static IP addresses or by using Dynamic Host Configuration Protocol (DHCP).
- Use unique subnets. If you have private networks that aren't 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, don't specify 10.0.0.0/24 for both networks even if you give each adapter a unique IP address. Using unique subnets helps avoid routing loops and other network communication problems. For example, if the segments were accidentally configured into the same collision domain because of incorrect virtual local area network (VLAN) assignments.

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 shouldn't have single points of failure.

AD DS and infrastructure requirements for failover clusters

Failover clusters depend on infrastructure services. Windows Server supports multiple-domain clusters and workgroup clusters.

Note: Though you can deploy multiple-domain clusters and workgroup clusters, you shouldn't use this configuration for Hyper-V or file server clusters.

You should install the same Windows Server features and roles on each node. Inconsistent configurations on cluster nodes might cause instability and performance issues. Additionally, you shouldn't install the Active Directory Domain Services (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 Domain Name System (DNS) for name resolution. DNS dynamic update is a supported configuration.
- A domain role. In Windows Server failover clusters, the cluster nodes don't 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 doesn't have to be a **Domain Admins** account. It can be a **Domain Users** account that's in the **Administrators** group on each clustered server. Additionally, if the account isn't 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 isn't required when you create detached clusters in AD DS.

In Windows Server, you don't need to have a cluster service account. Instead, the cluster service automatically runs in a special context that provides the specific permissions and credentials that are necessary for the service (like 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. Additionally, 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

In Windows Server 2019, Microsoft supports in-place upgrades of a cluster node's operating system. To upgrade cluster nodes in previous versions, it was necessary to drain the node, evict it, install everything fresh, and then add it back to the cluster. The rolling upgrade process for clusters that Windows Server 2016 introduced stays the same. You can upgrade Windows Server up to two versions at once. This removes the need to island hop from one version to the next. For example, if you're currently running Windows Server 2012 R2, you can upgrade to Windows Server 2019 directly without having to first upgrade to Windows Server 2016.

As a best practice, each cluster node should run the same edition of Windows Server 2019. The edition can be either Windows Server 2019 Standard or Windows Server 2019 Datacenter. The nodes should also have the same software updates. Depending on the role that's clustered, a Server Core installation of Windows Server 2019 might also meet software requirements.

The same versions of any operating system updates should exist on all the nodes that are part of a cluster.

Test your knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

What component provides block-level replication for any type of data in complete volumes?

- Storage Replica
- Cluster Shared Volume (CSV) Replica
- Cluster set
- Quorum

Question 2

Which term is defined as the majority of voting nodes in an active cluster membership plus a witness vote?

- Failover voting
- CSV
- Cluster set
- Quorum

Question 3

What quorum configuration is a best practice for Windows Server 2019 failover clusters?

Creating and configuring failover clusters

Lesson overview

Failover clusters that you create in Windows Server 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 often 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 2019 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'll be able to:

- Describe the **Validate a Configuration Wizard** and cluster support policy requirements.
- Describe the process to create a failover cluster.
- Create a failover cluster.
- Configure storage.
- Configure networking.
- Configure quorum options.
- Configure a quorum.
- Configure roles.
- Manage failover clusters.
- Configure cluster properties.
- Configure failover and failback.

The Validate a Configuration Wizard and cluster support policy requirements

The Validate a Configuration Wizard

Whether you're configuring a brand new Windows failover cluster or are maintaining an existing one, the **Validate a Configuration Wizard** is a tool for verifying a storage configuration. Use the **Validate a Configuration Wizard** to perform a variety of tests to help ensure that cluster components are accurately configured and supported in a clustered environment.

The wizard includes various tests, such as listing the system configuration or performing storage and network tests. These tests can run on a new, proposed member of a cluster, or you can run them to establish a baseline for an existing cluster. The wizard can also troubleshoot a broken cluster by isolating the network, storage, or system component that's failing a particular test.

Support policy requirements

Before you create a new failover cluster, Microsoft strongly recommends that you validate the configuration to make sure that the hardware and hardware settings are compatible with failover clustering. Run the failover cluster validation tests on a fully configured failover cluster before you install the Failover Clustering feature.

Cluster validation is intended to:

- Find hardware or configuration issues before a failover cluster goes into production.
- Help ensure that the clustering solution that you deploy is dependable.
- Provide a way to validate changes to the hardware of an existing cluster.
- Perform diagnostic tests on an existing cluster.

Note: Microsoft supports a cluster solution only if the complete configuration passes all validation tests and if all hardware is certified for the version of Windows Server that the cluster nodes are running.

Indicators and their meanings

The possible indicators that the wizard will present include:

- A green check mark (passed). This indicates that the failover cluster is valid.
- A yellow yield sign (warning). The yellow yield sign indicates that the aspect of the proposed failover cluster that's being tested isn't in alignment with Microsoft best practices. Investigate this aspect to make sure that the configuration of the cluster is acceptable for the cluster's environment, the requirements of the cluster, and the roles that the cluster hosts.
- A red circle with a single bar (canceled). If a failover cluster receives a red "X" (fail) in one of the tests, you can't use the part of the failover cluster that failed in a Windows Server failover cluster. Additionally, if a test fails, all other tests don't run, and you must resolve the issue before you install the failover cluster.

Validate after changes

Run validation tests when a major component of the cluster is changed or updated. For example, run validation tests when you make any of the following configuration changes to a failover cluster:

- Add a node to the cluster.
- Upgrade or replace the storage hardware.
- Upgrade the firmware or the driver for host bus adapters.
- Update the multipathing software or the version of the device-specific module.
- Change or update a network adapter.

Microsoft Support might also ask you to run validation tests against a production cluster. When you do this, failover cluster validation tests perform a hardware and software inventory, test the network, validate the system configuration, and perform other relevant tests. In some scenarios, you can run only a subset of the tests. For example, when troubleshooting a problem with networking, Microsoft Support might ask you to run only the hardware and software inventory and the networking test against the production cluster.

When an underlying storage configuration change or problem causes a cluster storage failure, Microsoft Support might also ask that you run validation tests on production clusters. The relevant disk resources

and the resources on which the disks depend are taken offline during the test. Therefore, run validation tests when the production environment isn't in use.

Create a failover cluster

Before creating a failover cluster, verify the following prerequisites:

- Make sure that all the servers that will function as nodes are running the same version of Windows Server.
- Ensure that you meet all hardware and software requirements.
- To add clustered storage during the creation process, make sure that all servers can access the storage.

Adding a failover cluster by using the Create Cluster Wizard

Follow the instructions in the **Create Cluster Wizard** to specify:

- The servers to include in the cluster.
- The name of the cluster.
- Any IP address information that your Dynamic Host Configuration Protocol (DHCP) settings don't automatically supply.

After the wizard runs, a **Summary** page appears. Select the **View Report** option to access a report on the tasks that the wizard performed. After you close the wizard, you can find the report at `<SystemRoot>\Cluster\Reports`, where *SystemRoot* is the location of the operating system; for example, `C:\Windows`.

Note: If you're using Windows Server 2019, you can use a distributed network name for the cluster. A distributed network name uses the IP addresses of the member servers instead of requiring a dedicated IP address for the cluster. By default, Windows uses a distributed network name if it detects that you're creating a cluster in Microsoft Azure, which means that you don't have to create an internal load balancer for the cluster. Windows will use a normal static or IP address if you're running on-premises.

Adding a failover cluster in Windows Admin Center

Windows Admin Center is a browser-based management tool that allows you to manage Windows Server computers with no Azure or cloud dependencies. You can manage failover cluster nodes as individual servers by adding them as server connections in Windows Admin Center. You can also add them as failover clusters to a view and manage cluster resources, storage, network, nodes, roles, virtual machines, and virtual switches.

Windows Admin Center provides one user interface (UI) in which you can:

- Examine cluster performance history to assess how clusters and nodes are performing.
- Examine the system insights feature of Windows Server, which uses machine learning and predictive analytics.
- Utilize persistent memory.

When you use **Failover Cluster Manager** in Windows Server 2019, you'll receive a prompt to try managing your clusters with Windows Admin Center. To add a failover cluster to Windows Admin Center, add a failover connection through the UI. Additionally, you can manage hyper-converged clusters by adding a cluster as a hyper-converged cluster connection.

To create a failover cluster by using Windows Admin Center, follow these steps:

1. Under **All Connections**, select **Add**.
2. Select **Failover Connection**.
3. Enter the name of the cluster, and if prompted, enter the credentials to use.
4. Add the cluster nodes as individual server connections.
5. Select **Submit** to finish.

After creating a cluster, you can use the **Failover Cluster Management** console to monitor its status and manage the available options.

Additional reading: For more information about failover clustering requirements and storage, refer to **Failover clustering hardware requirements and storage options**³.

Demonstration: Create a failover cluster

In this demonstration, you will learn how to:

- Validate a cluster configuration.
- Create a failover cluster.

Demonstration steps

Validate and create a failover cluster

1. On **SEA-SVR2**, sign in as **Contoso\Administrator** with password **Pa55w.rd**.
2. Select **Start**, and then select **Windows PowerShell**.
3. Use the **Test-Cluster SEA-SVR2, SEA-SVR3** cmdlet to start cluster validation.
4. Review the validation report. You can expect a few warning messages to display, but there should be no errors.
5. Use the **New-Cluster -Name WFC2019 -Node sea-svr2 -StaticAddress 172.16.10.125** cmdlet to create a new cluster.
6. Use the **Add-ClusterNode -Name SEA-SVR3** cmdlet to add **SEA-SVR3** as a cluster node.

Configure storage

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). Shared SAS provides the lowest cost option; however, it isn't very flexible for deployment, because cluster nodes must be physically close together. Additionally, the shared storage devices that support shared SAS have a limited number of connections for cluster nodes.

³ <https://aka.ms/clustering-requirements>

- iSCSI. Internet SCSI (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 gigabits 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, 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, you can use a shared virtual hard disk (VHD) as storage for virtual machine (VM) guest clustering. A shared VHD should be located on a Cluster Shared Volume (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.

In addition to using storage as a cluster component, you can also 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 protect your environment from risks such as:

- Data access failures.
- Volume unavailability.
- Server node failures.

You must use Storage Spaces Direct or shared storage that's compatible with Windows Server. You can use shared storage that's attached, and you can also use Server Message Block (SMB) 3.0 file shares as shared storage for servers that are running Hyper-V that are configured in a failover cluster.

Storage Spaces Direct

Storage Spaces Direct uses servers with locally attached drives to create highly available, highly scalable software-defined storage at less cost than that of traditional SAN or network-attached storage arrays. Its converged or hyper-converged architecture simplifies procurement and deployment, while features such as caching, storage tiers, and erasure coding, together with hardware improvements such as Remote Direct Memory Access networking and Non-Volatile Memory Host Controller Interface Specification-Enhanced drives, deliver efficiency and performance.

Additional reading: For an overview of Storage Spaces Direct, refer to **Storage Spaces Direct overview**⁴.

Additional reading: For more information about deploying Storage Spaces Direct, refer to **Deploy Storage Spaces Direct**⁵.

Storage requirements

In most cases, attached storage should have multiple separate disks (logical unit numbers, or LUNs) that are configured at the hardware level. For some clusters, one disk functions as the witness disk, which is described at the end of this topic. Other disks have the required files for the clustered roles, formerly called *clustered services and applications*.

Storage requirements include the following:

- Use basic disks, not dynamic disks, to use the native disk support that Failover Clustering includes.

⁴ <https://aka.ms/storage-spaces-direct-overview>

⁵ <https://aka.ms/deploy-storage-spaces-direct>

- If you use CSV to format the partitions, each partition must be NTFS or Resilient File System (ReFS). We recommend that you format the partitions with NTFS.

Note: If you have a witness disk for your quorum configuration, you can format the disk with NTFS or ReFS.

For the partition style of the disk, you can use master boot record (MBR) or GUID partition table (GPT). A *witness disk* is a disk in the cluster storage that's designated to hold a copy of the cluster configuration database. A failover cluster has a witness disk only if this is specified as part of the quorum configuration.

Deploying SANs with failover clusters

When deploying a SAN with a failover cluster, follow these guidelines:

- Confirm compatibility of the storage. Confirm with manufacturers and vendors that the storage, including drivers, firmware, and software used for the storage are compatible with the failover clusters in the version of Windows Server that you're running.
- Isolate storage devices, one cluster per device. Servers from different clusters must not be able to access the same storage devices. In most cases, using a LUN for one set of cluster servers should isolate them from all other servers through LUN masking or zoning.
- Consider using multipath input/output (I/O) software or teamed network adapters. In a highly available storage fabric, you can deploy failover clusters with multiple host bus adapters (HBAs) by using multipath I/O software or network adapter teaming (also called load balancing and failover). This provides the highest level of redundancy and availability.

Device controllers or appropriate adapters for storage

SAS or Fibre Channel

If you're using SAS or Fibre Channel, all components of the storage stack should be identical in all clustered servers. It's required that the Microsoft Multipath I/O (MPIO) software and device-specific module software components be identical. It's recommended that the mass-storage device controllers that are attached to the cluster storage, such as the HBA, HBA drivers, and HBA firmware, be identical. If you use dissimilar HBAs, you should verify with the storage vendor that you're following their supported or recommended configurations.

iSCSI

If you're using iSCSI, each clustered server must have one or more network adapters or HBAs that are dedicated to the iSCSI storage. The network that you use for iSCSI can't be used for network communication. In all clustered servers, the network adapters that you use to connect to an iSCSI storage target should be identical, and we recommend that you use Gigabit Ethernet or faster.

Other storage requirements

You must use shared storage that's certified for Windows Server.

For example, when configuring a two-node failover cluster, the storage should have at least two separate volumes (LUNs) if using a witness disk for quorum. The witness disk is a disk in the cluster storage that's designated to hold a copy of the cluster configuration database. For this two-node cluster example, the quorum configuration will be node and disk majority. Node and disk majority means that the nodes and the witness disk each have copies of the cluster configuration, and the cluster has a quorum as long as a

majority (two out of three) of these copies are available. The other volume (LUN) will contain the files that are being shared with users.

Other storage requirements include the following:

- To use the native disk support in failover clustering, use basic disks, not dynamic disks. Microsoft recommends that you format the partitions with NTFS for the witness disk—the partition must be NTFS. For the partition style of the disk, you can use MBR or GPT.
- The storage must respond correctly to specific SCSI commands, and the storage must follow the standard called SCSI Primary Commands-3 (SPC-3). 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.

Configure networking

For a failover cluster in Windows Server to be considered an officially supported solution by Microsoft, all hardware and software components must meet the qualifications for Windows Server. The fully configured solution (servers, network, and storage) must pass all tests in the **Validate a Configuration Wizard**, which is part of the failover cluster snap-in.

A failover cluster requires the following:

- Servers:
 - We recommend using matching computers with the same or similar components.
 - The servers for a two-node failover cluster must run the same version of Windows Server. They should also have the same software updates (patches).
- Network adapters and cables:
 - The network hardware, like other components in the failover cluster solution, must be compatible with Windows Server.
 - If you use Internet SCSI (iSCSI), the network adapters must be dedicated either to network communication or iSCSI, not both.
 - In the network infrastructure that connects your cluster nodes, avoid having single points of failure. You can do this in multiple ways, including:
 - Connecting your cluster nodes by multiple distinct networks.
 - Connecting your cluster nodes with one network that's constructed with teamed network adapters, redundant switches, redundant routers, or similar hardware that removes single points of failure.

Network and domain account requirements

To install a failover cluster, you'll need the following network infrastructure:

- Network settings and IP addresses. When you use identical network adapters for a network, also use identical communication settings on those adapters; for example, speed, duplex mode, flow control, and media type. Additionally, compare the settings between the network adapter and the switch to which it connects, and make sure that no settings are in conflict.
- Subnets. If you have private networks that aren't routed to the rest of your network infrastructure, ensure that each of these private networks uses a unique subnet. This is necessary even if you give

each network adapter a unique IP address. For example, if you have a cluster node in a central office that uses one physical network, and you have another node in a branch office that uses a separate physical network, don't specify 10.0.0.0/24 for both networks, even if you give each adapter a unique IP address.

- DNS. The servers in the cluster must be using Domain Name System (DNS) for name resolution. The DNS dynamic update protocol can be used.
- Domain role. All servers in a cluster must be in the same Active Directory domain. As a best practice, all clustered servers should have the same domain role, either a member server or domain controller. The recommended role is member server.
- Domain controller. Microsoft recommends that clustered servers be member servers. If they are, you need another server that acts as the domain controller in the domain that contains your failover cluster.
- Clients. As needed for testing, you can connect one or more networked clients to the failover cluster that you create, and you can observe the effect on a client when you move or fail over the clustered file server from one cluster node to the other.

You also will need an administrative 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 the servers in that cluster.
- The account doesn't need to be a **Domain Admins** account—it can be a **Domain Users** account that's in the **Administrators** group on each clustered server.
- Additionally, if the account isn't a **Domain Admins** account, the account, or the group in which the account is a member, must be given the Create Computer Objects and Read All Properties permissions in the domain organizational unit that it will reside in.

Network improvements

As it relates to networks, the Windows Server 2019 release has some notable improvements, including:

- Failover Cluster no longer uses NT LAN Manager (NTLM) authentication.
- Cross-cluster domain migration was introduced.
- Cluster network naming has been enhanced.

Failover Cluster no longer uses NTLM authentication

Windows Server 2019 offers improved security enhancements in relation to failover clustering. Windows Server 2019 now supports clusters without NTLM dependencies, because Microsoft has moved away from NTLM in favor of certificate-based intra-cluster Server Message Block (SMB) authentication.

Cross-domain cluster migration

Windows Server 2019 offers the ability to migrate clusters from one domain to another without destroying the original cluster. You can now move clusters from one domain to another by using a series of Windows PowerShell scripts. The scripts allow the process to dynamically change the NetNames Active Directory integration and then dynamically change to and from domain joined to workgroup and vice versa. The new process keeps your cluster intact.

Enhancements to cluster network naming

With the introduction of Windows Server 2019, you can now use a distributed network name for a cluster much like you can for Scale-Out File Server clusters. A distributed network name uses the IP addresses of member servers instead of requiring a dedicated IP address for the cluster. By default, Windows Server uses a distributed network name if it detects that you're creating the cluster in Microsoft Azure, which removes the need to create an internal load balancer for the cluster. To take advantage of the naming enhancement, use the **New-Cluster** cmdlet.

Failover cluster networks

Networks and network adapters are important parts of each cluster implementation. You can't 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 security-enhanced physical 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.

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, don't use a storage network for node or client communication.

Though not a best practice, you can use 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 using a separate Ethernet network for iSCSI, which is primarily 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.

Cluster networking features

The networking features in clusters that are based on Windows Server 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 that are 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's 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's 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 the nodes in a 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.

Configure quorum options

A quorum for a cluster is determined by the number of voting elements that must be part of active cluster membership for that cluster to start properly or continue running.

The quorum model in Windows Server is flexible. If you must modify the quorum configuration for your cluster, you can use the **Configure Cluster Quorum Wizard** or the failover clusters Windows PowerShell cmdlets.

Modify quorum configurations

If you must modify the quorum configuration for a cluster, you can use the **Configure Cluster Quorum Wizard** or the failover clusters Windows PowerShell cmdlets.

The following three quorum configuration options are available in the **Configure Cluster Quorum Wizard**.

Use typical settings

The cluster automatically assigns a vote to each node and dynamically manages the node votes. If it's suitable for your cluster and cluster shared storage is available, the cluster selects a witness disk. This option is recommended in most cases because the cluster software automatically chooses a quorum and a witness configuration that provides the highest availability for your cluster.

Add or change the quorum witness

You can add, change, or remove a witness resource. You can configure a witness file share or witness disk. The cluster automatically assigns a vote to each node and dynamically manages the node votes.

Advanced quorum configuration and witness selection

You should select this option only when you have application-specific or site-specific requirements for configuring a quorum. You can modify the quorum witness, add or remove node votes, and choose whether the cluster dynamically manages node votes. By default, votes are assigned to all nodes, and the node votes are dynamically managed.

Quorum modes

Depending on the quorum configuration option that you choose and your specific settings, the cluster will be configured in one of the following quorum modes.

Table 1: Quorum modes

Quorum mode	Description
Node majority	Only nodes have votes, and no quorum witness is configured. The cluster quorum is the majority of voting nodes in the active cluster membership.
Node majority with witness	Nodes have votes, and a quorum witness has a vote. The cluster quorum is the majority of voting nodes in the active cluster membership plus a witness vote. A quorum witness can be a designated witness disk or a designated file share witness.
No majority	No nodes have votes, and only a witness disk has a vote. The cluster quorum is determined by the state of the witness disk. Generally, this mode isn't recommended and shouldn't be selected, because it creates a single point of failure for the cluster.
Witness configuration	As a general rule, when you configure a quorum, the voting elements in the cluster should be an odd number. Therefore, if the cluster has an even number of voting nodes, you should configure a witness disk or a file share witness. The cluster will be able to sustain one additional node down. Additionally, adding a witness vote enables the cluster to continue running if half the cluster nodes simultaneously fail or are disconnected. A witness disk is usually recommended if all nodes can access the disk. A file share witness is recommended when you must consider multisite disaster recovery with replicated storage. Configuring a witness disk with replicated storage is possible only if the storage vendor supports read/write access from all sites to the replicated storage. A witness disk isn't supported with Storage Spaces Direct.

Quorum mode	Description
Node vote assignment	As an advanced quorum configuration option, you can choose to assign or remove quorum votes on a per-node basis. By default, all nodes are assigned votes. Regardless of vote assignment, all nodes continue to function in the cluster, receive cluster database updates, and can host applications. You might want to remove votes from nodes in certain disaster recovery configurations. For example, in a multisite cluster, you can remove votes from the nodes in a backup site so that those nodes don't affect quorum calculations. This configuration is recommended only for manual failover across sites. The configured vote of a node can be verified by getting the <code>NodeWeight</code> common property of the cluster node by using the <code>Get-ClusterNode</code> Windows PowerShell cmdlet. A value of 0 indicates that the node doesn't have a quorum vote configured. A value of 1 indicates that the quorum vote of the node is assigned and that the cluster is managing it. The vote assignment for all cluster nodes can be verified by using the <code>Validate Cluster Quorum</code> wizard.

General recommendations for quorum configuration

Failover clustering in Windows Server automatically configures the quorum for a new cluster based on the number of nodes configured and the availability of shared storage. This is usually the most appropriate quorum configuration for that cluster. However, it's a good idea to review the quorum configuration after the cluster is created, before placing the cluster into production. To access the detailed cluster quorum configuration, you can use the **Validate a Configuration Wizard** or the **Test-Cluster** Windows PowerShell cmdlet to run the validate quorum configuration test. In **Failover Cluster Manager**, the basic quorum configuration is part of the summary information for the selected cluster, or you can review the information about quorum resources that returns when you run the **Get-ClusterQuorum** Windows PowerShell cmdlet.

You can run the **Configure Cluster Quorum Wizard** at any time to validate that the quorum configuration is optimal for a cluster. The test output indicates if a change to the quorum configuration is recommended and the settings that are optimal. If a change is recommended, you can use the **Configure Cluster Quorum Wizard** to apply the recommended settings.

After a cluster is in production, don't change the quorum configuration unless you decide that a change is appropriate for the cluster. You might want to consider changing the quorum configuration when:

- You need to add or evict nodes.
- You need to add or remove storage.
- A long-term node or witness failure occurs.
- You need to recover a cluster in a multisite disaster recovery scenario.

Demonstration: Configure a quorum

In this demonstration, you will learn how to change the quorum configuration.

Demonstration steps

Configure a quorum

1. On **SEA-ADM1**, open **Failover Cluster Manager**, and then select `WFC2019.Contoso.com`.
2. Select **Configure Cluster Quorum Settings**.
3. On the **Select Quorum Configuration Option** page, select **Use default quorum configuration**. Explain the other available options for quorum settings.
4. Browse to the **Disks** node, and then point out that one of the cluster disks is assigned as **witness disk in Quorum**.

Configure roles

Create clustered roles

After creating a failover cluster, you can create clustered roles to host cluster workloads.

The following table lists the clustered roles that you can configure in the **High Availability Wizard** and the associated server role or feature that you must install as a prerequisite.

Table 1: Clustered roles and server roles or features

Clustered role	Role or feature prerequisite
Namespace Server	Namespaces (part of the File Server role)
DFS Namespace Server	DHCP Server role
Distributed Transaction Coordinator (DTC)	None
File Server	File Server role
Generic Application	Not applicable
Generic Script	Not applicable
Generic Service	Not applicable
Hyper-V Replica Broker	Hyper-V role
iSCSI Target Server	iSCSI Target Server (part of the File Server role)
iSNS Server	iSNS Server Service feature
Message Queuing	Message Queuing Services feature
Other Server	None
Virtual Machine	Hyper-V role
WINS Server	WINS Server feature

To create a clustered role, follow these steps:

1. Use **Server Manager** or Windows PowerShell to install the role or feature that's required for a clustered role on each failover cluster node. For example, if you want to create a clustered file server, install the File Server role on all cluster nodes.

2. In **Failover Cluster Manager**, expand the cluster name, right-click or access the context menu for **Roles**, and then select **Configure Role**.
3. Follow the steps in the **High Availability Wizard** to create the clustered role.
4. To verify that the clustered role was created, in the **Roles** pane, make sure that the role has a status of **Running**.

The **Roles** pane indicates the owner node. You use the owner node to specify the nodes that will be first to take over in case of a failure. To test failover:

1. In the **Roles** pane, in the **Owner Node** column, right-click or access the context menu for the role, select **Move**, and then select **Select Node**.
2. In the **Move Clustered Role** dialog box, select the desired cluster node, and then select **OK**.
3. In the **Owner Node** column, verify that the owner node changed.

Manage failover clusters

After your cluster infrastructure is running, you should set up monitoring procedures to prevent failures. Additionally, you should have backup and restore procedures for the cluster configuration. With Cluster-Aware Updating in Windows Server, you can update cluster nodes without downtime. In this lesson, you'll learn about monitoring, backing up and restoring, and updating cluster nodes.

Manage cluster nodes

After you create a cluster and put it into production, you might have to perform occasional 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 new node.
- Pausing a node. You can pause a node to prevent resources from failing over or moving to the node. You typically pause a node when it's undergoing maintenance or troubleshooting.
- Evicting a node. You can evict a node from a cluster. After you evict the node, you must add it back to the cluster. You evict a node when it's 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**.

You can manage a cluster in the **Actions** pane of the **Failover Cluster Management** console.

Note: If you plan to change a failover cluster configuration, you'll need to run the **Validate a Configuration Wizard** again for it to continue being supported.

Tools to help manage 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 can also use the **tracert.exe** tool to export data for analysis. Additionally, you can use the Multipurpose Internet Mail Extension Hypertext Markup Language (MHTML)-formatted cluster configuration reports and the **Validate a Configuration Wizard** to troubleshoot problems with cluster configuration and hardware changes.

Event Viewer

If problems arise in a cluster, use **Event Viewer** to examine events with a Critical, Error, or Warning severity level. Additionally, you can access informational-level events in the Failover Clustering Operations log, which you can access in **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.

Windows Server doesn't replicate event logs among nodes. However, the **Failover Cluster Management** snap-in has a Cluster Events option that you can use to access 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 more logs, such as the Debug and Analytic logs, in **Event Viewer**. To display these logs, change the view on the menu by selecting the **Show Analytic** and **Debug Logs** options.

Event tracing for Windows Server

Event tracing for Windows Server is a kernel component that's available soon after startup and late into shutdown. It's designed to allow the fast tracing and delivery of events to both trace files and consumers. Because it's designed to be fast, it allows only basic, in-process filtering of events based on event attributes.

The event trace log has a comprehensive accounting of failover cluster actions. Use **tracertp.exe** to access the information in the event trace log. **Tracertp.exe** parses event trace logs only on the node on which it runs. All the individual logs are collected in a central location. To transform the XML file into a text file or an HTML file that you can open in Microsoft Edge or Internet Explorer, you can parse the XML-based file by using the Microsoft Extensible Stylesheet Language (XSL) parsing command-line tool, **msxsl.exe**, and an XSL style sheet.

Performance and Reliability Monitor snap-in

The **Performance and Reliability Monitor** snap-in is another option to help monitor failover clusters, and you can use it to:

- Monitor how application performance is trending on each node. To determine how an application is performing, you can access specific information on the system resources that are used on each node and assess how that information is trending.
- Monitor how application failures and stability on each node are trending. You can pinpoint when application failures occur and match them with other events on the node.
- Modify trace log settings. You can start, stop, and adjust trace logs, including their size and location.

Configure cluster properties

Each failover cluster object has a set of properties that define its identity and behavior in the cluster. Applications and resource DLLs manipulate the properties of failover cluster objects to maintain awareness of the failover cluster and to change the failover cluster to suit their needs. The following table defines properties of failover cluster objects.

Table 1: Failover cluster object properties

Property	Property definition
Cluster Common Properties	Cluster common properties are stored in the cluster database and apply to the cluster as a whole.
Groupset Common Properties	Common properties for groupsets are data values stored in the cluster database that describe the identity and behavior of each groupset in a cluster.
Group Common Properties	Common properties for groups are data values stored in the cluster database that describe the identity and behavior of each group in a cluster.
Network Common Properties	Common properties for networks are data values stored in the cluster database that describe the identity and behavior of each network in a cluster.
Network Interface Common Properties	Common properties for network interfaces are data values stored in the cluster database that describe the identity and behavior of each network interface in a cluster.
Node Common Properties	Common properties for nodes are data values stored in the cluster database that describe the identity and behavior of each node in a cluster.
Resource Common Properties	Common properties for resources are data values stored in the cluster database that describe the identity and behavior of each resource in a cluster.
Resource Type Common Properties	Common properties for resource types are data values stored in the cluster database that describe the identity and behavior of each resource type in a cluster.
Virtual Machine Common Properties	Common properties for virtual machine (VM) resource types are data values stored in the cluster database that describe the identity and behavior of each VM resource type in a cluster.

Configure failover and failback

You can adjust failover settings, including preferred owners and failback settings, to control how a cluster responds when roles or services fail. You can configure these settings on the property sheet (on either the **General** or **Failover** tab) for the clustered service or application.

You can individually set preferred owners in the properties of each role. You can select multiple preferred owners and place them in any order. Selecting preferred owners provides more control over what node a particular role fails overs to and actively runs on.

Each role for failover and failback has settings that you can change. Failover settings can control how many times a cluster can try restarting a role in a particular amount of time. In Windows Server, the default is to allow only one failure every six hours. You can set the failback setting to **Prevent Failback**, which is the default, or **Allow Failback**. When allowing failback, you can set the role to immediately use failback or use failback during a certain number of hours.

Examples of failover settings

The following table provides examples that illustrate how these settings work.

Table 1: Failover settings examples

Example	Settings	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, the service or application fails back to Node1 when Node1 is available again.
Example 2	Failover tab, Maximum failures in the specified period: 2; Failover tab, Period (hours): 6	If the application or service fails no more than two times in a six-hour period, it restarts or fails 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 won't move among nodes indefinitely.

Test your knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

Does Windows Server 2019 require all nodes to be in the same domain?

- Yes
- No

Question 2

Can a node that runs Windows Server 2016 and one that runs Windows Server 2019 both run in the same cluster?

- Yes
- No

Question 3

You must install what feature on every server that you want to add as a failover cluster node?

- Cluster set
- Failback Clustering
- Hyper-V
- Failover Clustering

Question 4

When running the Validate a Configuration Wizard, what does the yellow yield symbol indicate?

- The failover cluster needs to fail back to the original node.
- The wizard is waiting for a file to download.
- The failover cluster creation is in progress.
- The failover cluster that's being tested isn't in alignment with Microsoft best practices.

Overview of stretch clusters

Lesson overview

In some scenarios, you must deploy cluster nodes on different sites. Usually, you do this when building disaster recovery solutions. In this lesson, you'll learn about deploying stretch clusters.

Lesson objectives

After completing this lesson, you'll be able to:

- Describe what a stretch cluster is.
- Explain the Storage Replica feature.
- Describe the prerequisites for implementing a stretch cluster.
- Explain synchronous and asynchronous replication.
- Select a quorum mode for a stretch cluster.
- Configure a stretch cluster.

What is a stretch cluster?

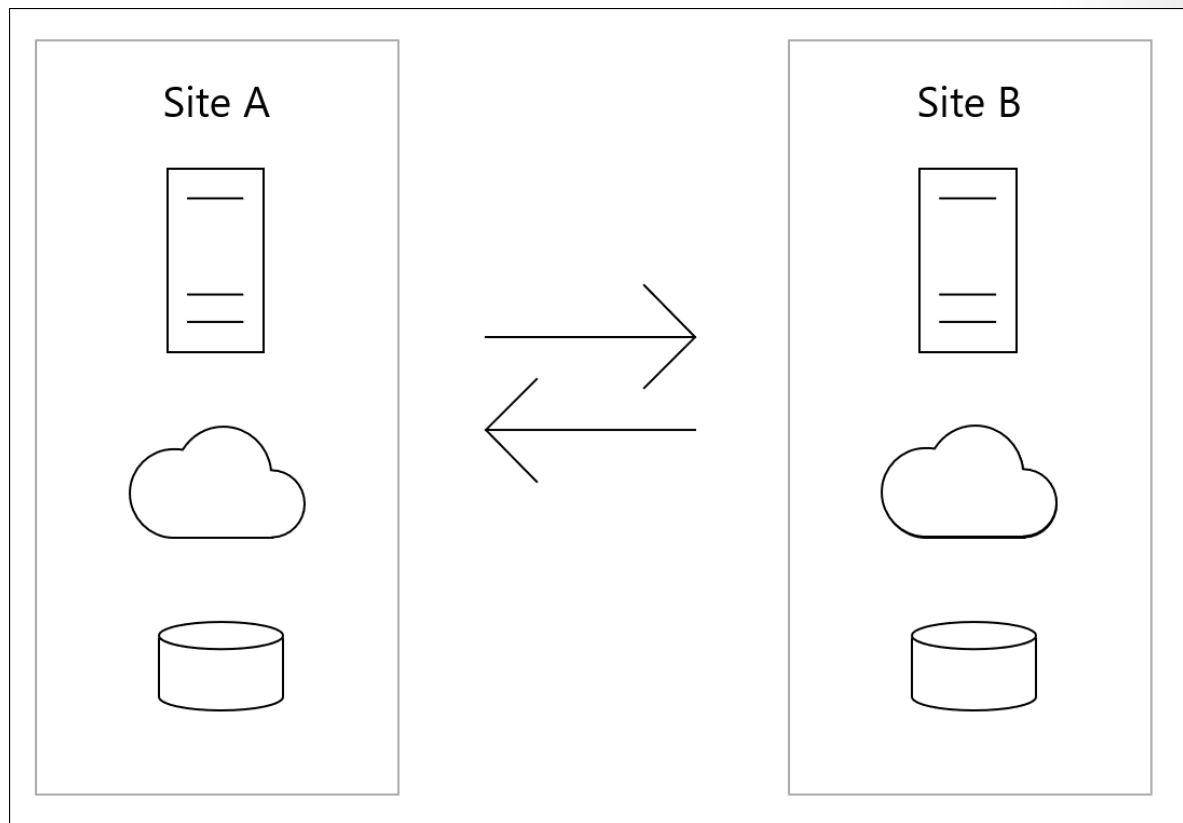


Figure 1: Stretch cluster

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 storage in the local site; however, with separate storage systems, you can't share a disk among sites.

A stretch cluster in a failover site has three main advantages over a remote server:

- 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 replicate changes manually.
- 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 organization. When you're considering whether to deploy a stretch cluster, you should evaluate the importance of the applications to the organization, 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 enough availability with only a modest increase in cost and complexity.

The complexity of a stretch cluster requires architectural and hardware planning that's more detailed than is necessary for a single-site cluster. It also requires you to develop organizational processes to test cluster functionality routinely.

Overview of Storage Replica

Storage Replica is Windows Server technology that enables replication of volumes between servers or clusters for disaster recovery. You can use Storage Replica to create stretch failover clusters that span two sites, with all nodes staying in sync.

Storage Replica supports synchronous and asynchronous replication:

- Synchronous replication mirrors data within a low-latency network site with crash-consistent volumes to ensure zero data loss at the file system-level during a failure.
- Asynchronous replication mirrors data across sites beyond metropolitan ranges over network links with higher latencies, but without a guarantee that both sites have identical copies of the data at the time of a failure.

Note: Although Storage Replica can technically replicate volume data for all applications, you should use the best practices for those applications that provide replication functionality, such as Microsoft Exchange Server, Hyper-V, or Microsoft SQL Server.

Storage Replica features

The main features of Storage Replica include:

- Zero data loss, block-level replication. With synchronous replication, there's no possibility of data loss. With block-level replication, there's no possibility of file locking.
- Simplified. Storage Replica has a design mandate for ease of use for straightforward deployment and management. Creation of a replication partnership between two servers can utilize Windows Admin Center. Deployment of stretch clusters uses a wizard in the familiar **Failover Cluster Manager** console.

- Guest and host. All Storage Replica capabilities are exposed in both virtualized guest-based and host-based deployments. This means that guests can replicate their data volumes even if running on non-Windows virtualization platforms or in public clouds if the guest is using Windows Server.
- It's Server Message Block (SMB) 3.0–based. Storage Replica uses the proven and mature technology of SMB 3.0, which was first released in Windows Server 2012. This means that all of SMB's advanced characteristics — such as multichannel and SMB direct support on RDMA over Converged Ethernet (RoCE), and iWARP — are available to Storage Replica.
- Security. Unlike many vendors' products, Storage Replica has industry-leading security technology built in. This includes packet signing, AES-128-GCM full data encryption, support for third-party encryption acceleration, and pre-authentication integrity man-in-the-middle attack prevention. Storage Replica uses Kerberos AES256 for all authentication between nodes.
- High performance initial sync. Storage Replica supports seeded initial sync, where a subset of data already exists on a target from older copies, backups, or shipped drives. Initial replication only copies the differing blocks, potentially shortening initial sync time and preventing data from using up limited bandwidth. Storage Replica block checksum calculations and aggregation means that initial sync performance is limited only by the speed of the storage and network.
- Consistency groups. Write ordering helps ensure that applications such as SQL Server can write to multiple replicated volumes and know that the data is written on the destination server sequentially.
- User delegation. Users can be delegated permissions to manage replication without being a member of the built-in **Administrators** group on replicated nodes, thereby limiting their access to unrelated areas.
- Network constraint. Storage Replica can be limited to individual networks by server and by replicated volumes to provide application, backup, and management software bandwidth.
- Thin provisioning. Thin provisioning in Storage Spaces Direct and storage area network (SAN) devices is supported, which provides near-instantaneous initial replication times under many circumstances.

Prerequisites for implementing a stretch cluster

Prerequisites

The prerequisites for implementing a stretch cluster include:

- An Active Directory Domain Services (AD DS) forest. Note that it doesn't need to run Windows Server 2019.
- Between 2 and 64 servers that are running Datacenter editions of Windows Server 2019 or Windows Server 2016. If they're running Windows Server 2019, you can use the Standard edition if replicating only a single volume of up to 2 terabytes (TB) in size is adequate.
- Two sets of shared storage using serial attached SCSI (SAS) just a bunch of disks (JBODs), such as with Storage Spaces Direct, Fibre Channel storage area networks (SANs), Shared Windows Virtual Hard Drive files (VHDX), or an Internet SCSI (iSCSI) target. The storage should have a mix of hard disk drives and solid-state drive media and must support persistent reservation. You'll make each storage set available to two of the servers only (asymmetric).
- Each set of storage must allow the creation of at least two virtual disks, one for replicated data and one for logs. The physical storage must have the same sector sizes on all the data disks. The physical storage must have the same sector sizes on all the log disks.

- At least one Gigabit Ethernet connection on each server for synchronous replication, but preferably Remote Direct Memory Access (RDMA).
- At least 2 gigabytes (GB) of random access memory (RAM) and two cores per server. You'll need more memory and cores for more virtual machines (VMs).
- Appropriate firewall and router rules to allow Internet Control Message Protocol (ICMP), Server Message Block (SMB) (port 445, plus 5445 for SMB Direct), and Web Services-Management (WS-MAN) (port 5985) bi-directional traffic between all nodes.
- A network between servers with enough bandwidth to contain your input/output (I/O) write workload and an average of =5 millisecond (ms) round-trip latency for synchronous replication. Asynchronous replication doesn't have a latency recommendation.
- The replicated storage can't be on the drive with the Windows operating system folder.
- Many of these requirements can be determined by using the **Test-SRTopology** cmdlet. You get access to this tool if you install the Storage Replica feature or the Storage Replica Management Tools feature on at least one server. There's no need to configure Storage Replica to use this tool, only to install the cmdlet.

Considerations for deploying a stretch cluster

Stretch clusters aren't appropriate for every application or organization. When you design a stretch cluster solution with a hardware vendor, clearly identify the organizational requirements and expectations. Not every scenario that involves more than one location is appropriate for a stretch cluster.

Stretch clustering is a high-availability strategy that primarily focuses on hardware platform availability. However, specific stretch cluster configurations and deployments have availability ramifications, including users' ability to connect to the application and 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 require more overhead than local clusters. Instead of a local cluster in which each cluster node attaches to a mass-storage device, each site of a stretch cluster must have comparable storage. Additionally, you must consider setting up replication among the cluster sites. Storage Replica in Windows Server provides storage-agnostic replication. However, you must also consider paying for more network bandwidth among the sites and developing the management resources with your organization to efficiently administer a 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 Microsoft SQL Server, Hyper-V, Microsoft Exchange Server, and AD DS should use their individual application stretch configurations (Hyper-V Replica, database availability groups, and so on).

Considerations for stretch cluster failover and failback

When you establish a stretch clustering structure, it's important that you define a procedure for the tasks that you should perform in case of a site disaster. Additionally, you should define a procedure for the tasks that you should perform for failback.

In most cases, failover of critical services to another site doesn't occur automatically, but instead consists of a manual or partially manual procedure. When defining your failover process, consider the following factors:

- Failover time. You must decide how long to 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, Domain Name System (DNS), and Dynamic Host Configuration Protocol (DHCP), that should fail over to another site. It isn't enough to have a cluster that's designed to fail over to another site. Failover clustering requires that you have Active Directory services running in a second site. You can't 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 more domain controllers and DNS servers or VMs in a second site.
- Quorum maintenance. It's important to design the quorum model in a way that each site has enough votes for maintaining cluster functionality. If that isn't possible, you can use options such as forcing a quorum or dynamic quorum to create a quorum in case of a disaster.
- Published services and name resolution. If you have services that are 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's the case, you should have a procedure for changing DNS records in the internal or public DNS. To reduce downtime, we recommended that you reduce the Time to Live (TTL) for critical DNS records.
- Client connectivity. A failover plan must include a design for client connectivity in 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. The failback process is rarely automated, and it usually occurs 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, consider the much larger picture of failover as part of a disaster recovery strategy. Windows Server has several technologies that can help with failover and failback, but you should also consider the other technologies in your infrastructure. Additionally, each failover and failback procedure depends greatly on the services that are implemented in a cluster.

Synchronous and asynchronous replication

Storage Replica uses synchronous or asynchronous replication. The following list describes the differences between synchronous and asynchronous replication:

- When using 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 doesn't write successfully to both storage systems, the application must try to write to the disk again. With synchronous replication, the data on both storage systems is identical.
- When using 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 doesn't support stretch clustering replication. This is because DFS Replication is designed to replicate smaller documents that aren't continuously kept open. As a result, it wasn't designed for high-speed, open-file replication.

When to use synchronous or asynchronous replication

Use synchronous replication when it's 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 increased distance might cause higher latency. If the latency is high, the performance and even the stability of the application might be affected.

Asynchronous replication overcomes 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.

Select a quorum mode for a stretch cluster

Stretch cluster

When creating a stretch cluster across geographically dispersed nodes, you should use a Microsoft Azure Cloud Witness whenever possible. However, in some cases, it might be more practical to use a file share witness.

Regardless of the selected witness, you should use dynamic witness mode, which is the default. A stretch cluster spreads across multiple datacenters, and it's 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's important to use dynamic quorum and dynamic witness to avoid a shutdown of the cluster.

Note: In a stretch cluster, shared storage isn't accessible to the nodes in different locations. A witness disk isn't a suggested witness selection for this scenario.

File share witness

The major issue with a file share witness for most scenarios is the minimum recommendation of three datacenters to create the witness. However, if you're working in an environment where three or more datacenters are already in operation, creating a file share witness on a share in one of the locations might be the quickest and easiest witness option. A file share witness requires a file share that all the nodes in the cluster can access by using the Server Message Block (SMB) protocol. A file share witness doesn't keep a copy of the cluster database.

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 it doesn't keep a copy of the cluster database. However, rather than requiring a share and writing over SMB, Azure Cloud Witness uses Azure Blob storage and the REST-based API for Azure Storage.

No witness

You can also configure a cluster to not use any witness. Although you should avoid this solution, it's supported to prevent split-brain syndrome. You perform this configuration in Windows Server 2019 by using site-aware clustering. You can also configure no witness for manual failovers; for example, 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.

Configure a stretch cluster

You can configure a stretch cluster by using **Failover Cluster Manager** or Windows PowerShell. You can perform all the following steps directly on the cluster nodes or from a remote management computer that has Windows Server Remote Server Administration Tools (RSAT). The following steps are required to configure a stretch cluster by using **Failover Cluster Manager**.

The Failover Cluster Manager method

1. Don't add all the disks—just add a single disk. At this point, half the disks will register as offline because this is asymmetric storage. If replicating a physical disk resource workload like File Server for general use, you already have a role-attached disk ready to go.
2. Right-click or access the context menu for the Cluster Shared Volume (CSV) disk or role-attached disk, select **Replication**, and then select **Enable**.
3. Select the appropriate destination data volume, and then select **Next**. The displayed destination disks will have a volume that's the same size as the selected source disk. When moving between these wizard dialogs, the available storage will automatically move and come online in the background as needed.
4. Select the appropriate source log disk, and then select **Next**. The source log volume should be on a disk that uses solid-state drives or similarly fast media, not spinning disks.
5. Select the appropriate destination log volume, and then select **Next**. The displayed destination log disks will have a volume that's the same size as the selected source log disk volume.
6. Leave the Overwrite Volume value at Overwrite destination volume if the destination volume doesn't have a previous copy of the data from the source server. If the destination does have similar data from a recent backup or previous replication, select **Seeded destination disk**, and then select **Next**.
7. Leave the Replication Mode value at Synchronous Replication if you plan to use zero recovery point objective (RPO) replication. Change it to **Asynchronous Replication** if you plan to stretch your cluster over higher latency networks or need lower input/output (I/O) latency on the primary site nodes.
8. Leave the Consistency Group value at Highest Performance if you don't plan to use write ordering later with more disk pairs in the replication group. If you plan to add further disks to this replication group and you require guaranteed write ordering, select **Enable Write Ordering**, and then select **Next**.
9. Select **Next** to configure replication and the stretch cluster formation.
10. On the **Summary** screen, you can open the report in a web browser.

11. At this point, you have configured a Storage Replica partnership between the two halves of the cluster, but replication is ongoing. You can obtain the state of replication with a graphical tool in several ways:
- Use the **Replication Role** column and the **Replication** tab. When the initial sync is complete, the source and destination disks will have a **Replication Status** of **Continuously Replicating**.
 - Start **eventvwr.exe**.
 - On the source server, browse to **Applications and Services\Microsoft\Windows\StorageRepl-ica\Admin**, and then examine events 5015, 5002, 5004, 1237, 5001, and 2200.
 - On the destination server, browse to **Applications and Services\Microsoft\Windows\Storag-eReplica\Operational**, and then wait for event 1215. This event states the number of copied bytes and the time taken.

Test your knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

Which type of witness uses a basic format and doesn't keep a copy of the cluster database?

- USB witness
- Failback witness
- File share witness
- Microsoft Azure Cloud Witness

Question 2

What technology enables replication of volumes between servers or clusters for disaster recovery?

- File share witness
- Cluster set
- Cluster Shared Volume (CSV)
- Storage Replica

Question 3

What added features does enabling site-aware clustering in a stretch cluster provide?

High availability and disaster recovery solutions with Hyper-V VMs

Lesson overview

Moving virtual machines (VMs) from one server to another is a common procedure in the administration of Hyper-V environments. Most of the techniques for moving VMs in previous versions of Windows Server required downtime. With Windows Server 2019, VM migration has no downtime. In this lesson, you'll learn about VM migration and the available migration options.

Lesson objectives

After completing this lesson, you'll be able to:

- Describe high availability options for Hyper-V VMs.
- Explain **Live Migration**.
- Describe **Live Migration** requirements.
- Provide high availability with storage migration.

High-availability options for Hyper-V VMs

Most organizations have some business-critical applications that must be highly available. To make an application or service highly available, you must deploy it in an environment that provides redundancy for all the components that the application requires. To provide high availability for virtual machines (VMs) and the services hosted within VMs, you can choose to:

- Implement VMs as a clustered role (host clustering).
- Implement clustering inside VMs (guest clustering).
- Use Network Load Balancing (NLB) inside VMs.

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 VM 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 VM don't have to be cluster-aware. However, the VM is still highly available.

Some examples of cluster-unaware applications are a print server or a proprietary, network-based application such as an accounting application. Should the host node that controls the VM unexpectedly become unavailable, the secondary host node takes control and restarts or resumes the VM as quickly as possible. Additionally, you could move the VM from one node in the cluster to another in a controlled manner. For example, you could move the VM from one node to another while updating the host management Windows Server 2016 operating system.

The applications or services that are running on a VM don't have to be compatible with failover clustering, and they don't have to be aware that the VM is clustered. Because failover is at the VM-level, there are no dependencies on software that's installed on the VM.

Guest clustering

You configure guest failover clustering similarly to physical-server failover clustering, except that the cluster nodes are VMs. In this scenario, you create two or more VMs 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 VMs. Because you implement failover clustering within the guest operating system of each VM node, you can put the VMs 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 an application or service more robustly if you deploy the VMs on separate failover clustering-enabled Hyper-V host computers. With failover clustering implemented at both the host and VM levels, you can restart the resource regardless of whether the node that fails is a VM or a host. Such high-availability configurations for VMs that are running mission-critical applications in a production environment are considered optimal.

You should consider several factors when implementing guest clustering:

- The application or service must be failover cluster-aware. This includes any of the Windows Server services that are cluster-aware, in addition to any applications such as clustered Microsoft SQL Server and Microsoft Exchange Server.
- Hyper-V VMs in Windows Server can use Fibre Channel-based connections to shared storage, or you can implement Internet SCSI (iSCSI) connections from the VMs to the shared storage. You can also use the shared virtual hard disk feature to provide shared storage for VMs.

You should deploy multiple network adapters on the host computers and the VMs. Ideally, you should dedicate a network connection to the iSCSI connection (if you're 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 VMs 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's running on a host in the NLB cluster. NLB transparently distributes client requests among the hosts, and it enables clients to access the cluster by using a virtual host name or a virtual IP address. From a 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.

For these reasons, NLB is an appropriate solution for resources that don't have to accommodate exclusive read or write requests. Examples of NLB-appropriate applications are web-based front-end VMs to database applications, or Exchange Server Client Access servers.

When you configure an NLB cluster, you must install and configure the application on all the VMs that will participate in the NLB cluster. After you configure the application, you install the NLB feature in Windows Server within each VM's guest operating system (not on the Hyper-V hosts), and then configure an NLB cluster for the application.

Overview of live migration

Live Migration is a Hyper-V feature in Windows Server, and with it, you can transparently move running virtual machines (VMs) from one Hyper-V host to another without perceived downtime. The primary benefit of live migration is flexibility—running VMs aren't tied to a single host machine. This allows actions like draining a specific host of VMs before decommissioning or upgrading it. When paired with

Windows Failover Clustering, live migration allows the creation of highly available and fault-tolerant systems.

How Live Migration works

By using the **Live Migration** feature in Hyper-V failover clustering, you can move running VMs from one failover cluster node to another node in the same cluster. With **Live Migration**, users who are connected to the VM should experience almost no server outages.

Note: Shared-nothing live migrations are able to migrate without any shared storage or cluster. However, this relies on your network to perform reads and writes to both locations simultaneously. Because of this, a clustered solution is better for production environments.

You can control a 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), but you can also use Kerberos authentication. When choosing which authentication protocol to use, consider the following:

- CredSSP is the default configuration, and it's easy to configure; however, it's less secure than Kerberos authentication. To live-migrate VMs, CredSSP requires signing in to the source server, remote desktop, or remote Windows PowerShell session.
- Kerberos authentication is the more secure of the two options. It requires manual selection and constrained delegation configuration for that host. However, this doesn't require signing in to the Hyper-V host server for live migrations. You can't use Hyper-V Manager to perform a live migration of a clustered (highly available) VM.

You can start the live migration process by using:

- The **Failover Cluster Management** console.
- The **Virtual Machine Manager Administrator** console, if you use Virtual Machine Manager to manage your physical hosts.
- Hyper-V Manager.
- A Windows Management Instrumentation or Windows PowerShell script.

Note: Use **Live Migration** to significantly reduce a VM's downtime during a planned failover. During a planned failover, you start the failover manually. **Live Migration** doesn't apply during an unplanned failover, such as when the node hosting a VM fails. You can use Hyper-V Manager for moving VMs that aren't running as a failover cluster role; VMs that are running as a failover cluster role require the use of **Failover Cluster Manager**, and you can move them only to another node in the same failover cluster.

The live migration process

The live migration process consists of the following steps:

1. Migration setup. When an administrator starts a VM failover, the source node creates a TCP connection with the target physical host. This connection is used to transfer VM configuration data to the target physical host. **Live Migration** creates a temporary VM on the target physical host and allocates memory to the destination VM. The migration preparation also checks to determine whether a VM can be migrated.
2. Guest memory transfer. Guest memory transfers iteratively to the target host while the VM 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 VM continues to run. Hyper-V iterates the memory copy

process several times, and a smaller number of modified pages copy 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. State transfer. To migrate the VM to the target host, Hyper-V stops the source partition, transfers the state of the VM, including the remaining dirty memory pages, to the target host, and then begins running the VM on the target host.
4. Cleanup. The cleanup stage finishes the migration by tearing down the VM on the source host, terminating the worker threads, and signaling the completion of the migration.

Note: In Windows Server, you can perform live migration of VMs by using Server Message Block (SMB) 3.0 as a transport. This means that you can utilize key SMB features, such as SMB Direct and SMB Multi-channel, which provide high-speed migration with low central processing unit (CPU) utilization.

Live migration requirements

Live migration of virtual machines (VMs) in Windows Server delivers improved performance and flexibility. In Windows Server 2019, it's available inside and outside of clustered environments, both with and without shared storage. In this section, you'll learn about the common requirements for all live migrations, requirements for VMs in a cluster, VMs that use shared storage, and VMs that aren't clustered.

Common requirements

Common requirements for live migration include the following, with two or more servers running Hyper-V that:

- Support hardware virtualization.
- Use processors from the same manufacturer.
- Belong either to the same Active Directory domain or to domains that trust each other.
- Configure VMs to use virtual hard disks (VHDs) or virtual Fibre Channel disks (no physical disks).
- Use an isolated network, physically or through another networking technology such as virtual local area networks (VLANs), which is recommended for live migration network traffic.

Requirements for live migration in a cluster

Requirements for live migration in a cluster include ensuring that:

- Windows Failover Clustering is enabled and configured.
- Cluster Shared Volume (CSV) storage in the cluster is enabled.

Pass-through disks requirements

Physical disks that are directly attached to a VM, known as *pass-through disks*, are supported when all the following conditions are met:

- Requirements for live migration using shared storage:
 - All files that make up a VM (for example, VHDs, snapshots, and configuration) are stored on a Server Message Block (SMB) share.

- Permissions on the SMB share have been configured to grant access to the computer accounts of all servers that are running Hyper-V.
- Requirements for live migration with no shared infrastructure:
 - No extra requirements exist.

Requirements for non-clustered hosts

To set up non-clustered hosts for live migration, you'll need:

- A user account with permissions to perform the various steps. Membership in the local **Hyper-V Administrators** group or the **Administrators** group on both the source and destination computers meets this requirement, unless you're configuring constrained delegation. Membership in the **Domain Administrators** group is required to configure constrained delegation.
- Source and destination computers that either belong to the same Active Directory domain or belong to domains that trust each other.
- The Hyper-V management tools that are installed on a Windows Server or Windows 10 computer, unless the tools are installed on the source or destination server and you'll run the tools from the server.

Provide high availability with storage migration

With Hyper-V in Windows Server, you can use **Live Migration** to move storage with no downtime by moving it while a virtual machine (VM) is still running. You can perform this task by using the **Live Migration Wizard** in Hyper-V Manager or by using new Hyper-V cmdlets for Windows PowerShell.

You can add storage either to a standalone computer or to a Hyper-V cluster, and you can then move the VMs to the new storage while the VMs continue to run.

The most common reason for moving a VM's storage is to update the physical storage that's available to Hyper-V. You can also move VM storage between physical storage devices, at run time, to respond to reduced performance that results from bottlenecks in the storage throughput.

How storage migration works

Moving a VM to another host is a common procedure. As an administrator, you'll need to move VM files to another location for many reasons. For example, if the disk where a VM hard disk resides runs out of space, you must move the VM to another drive or volume.

Export operations can be time-consuming, depending on the size of the VM hard disks. VM and storage migration make it possible for you to move a VM to another location on the same host, or to another host computer, without turning off the VM.

To copy a virtual hard disk (VHD), begin live storage migration by using the Hyper-V console or Windows PowerShell, and then either complete the **Live Migration Wizard** or specify parameters in Windows PowerShell. Doing this creates a new VHD on the destination location, and the copy process starts. During the copy process, the VM is fully functional. However, all changes that occur during copying write 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 VMs to run on the destination VHD. Additionally, if the VM is moving to another host, the computer configuration is copied and the VM is associated with another host. If a failure were to occur on the destination side, a failback option is always

available to run on the source directory. After the VM is successfully migrated and associated with a new location, the process deletes the source VHDs.

The time that's necessary to move a VM depends on the source and destination location, the speed of the hard drives or storage, and the size of the VHDs. The moving process accelerates if source and destination locations are on storage that supports Windows Offloaded Data Transfers.

When you move a VM's VHDs to another location, the **Move Wizard** presents three available options in Hyper-V Manager:

- Move all the VM's data to a single location. You specify one single destination location, such as disk file, configuration, checkpoint, or smart paging.
- Move the VM's data to a different location. You specify individual locations for each VM item.
- Move only the VM's VHD. You move only the VHD file.

The **Move Wizard** and these options are only available if the Hyper-V VM isn't part of a failover cluster. All three of the options are achievable in **Failover Cluster Manager** by using the **Move Virtual Machine Storage** options.

Test your knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

Which feature would you use to configure a failover cluster when you use Hyper-V host servers?

- Site-aware clustering
- Client clustering
- Live clustering
- Host clustering

Question 2

Which feature can you use to transparently move running VMs from one Hyper-V host to another without perceived downtime?

- Site-aware cluster
- Storage migration
- Cluster set
- Live Migration

Module 06 lab and review

Lab: Implementing failover clustering

Scenario

As the business of Contoso, Ltd. grows, it's becoming increasingly important that many of the applications and services on its network are always available. Contoso 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 can't be made highly available by using Network Load Balancing (NLB). Therefore, you should use a different technology to make these applications highly available.

As one of the senior network administrators at Contoso, you're responsible for implementing failover clustering on the servers that are running Windows Server 2019 to provide high availability for network services and applications. You're also responsible for planning the failover cluster configuration and deploying applications and services on the failover cluster.

Objectives

After completing this lab, you'll 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.

Estimated time: 60 minutes

Module review

Use the following questions to check what you've learned in this module.

Question 1

What term describes a loosely coupled grouping of multiple failover clusters?

- Cluster set
- Failback Clustering
- Hyper-V
- Failover Clustering

Question 2

When running the Validate a Configuration Wizard, what does the red "X" indicator mean?

- The failover cluster needs to fail back to the original node.
- You can't use the part of the failover cluster that failed.
- Failover cluster creation is in progress.
- The failover cluster that's being tested isn't in alignment with Microsoft best practices.

Question 3

What component provides a consistent, distributed namespace that clustered roles can use to access shared storage from all nodes?

- Storage Replica
- Cluster Shared Volume (CSV)
- Cluster set
- Quorum

Question 4

Which type of witness is ideal when shared storage isn't available or when the cluster spans geographical locations?

- USB witness
- Failback witness
- File share witness
- Microsoft Azure Cloud Witness

Question 5

What technology provides high availability where each site has a separate storage system with replication among the sites?

- Stretch cluster
- Cluster set
- CSV
- Storage Replica

Question 6

Which feature provides high availability for applications or services running on the VM that don't have to be compatible with failover clustering?

- Site-aware clustering
- Client clustering
- Live clustering
- Host clustering

Question 7

Which feature distributes IP traffic to multiple instances of a TCP/IP service?

- Site-aware cluster
- Storage migration
- Network Load Balancing (NLB)
- Live Migration

Answers

Question 1

What component provides block-level replication for any type of data in complete volumes?

- Storage Replica
- Cluster Shared Volume (CSV) Replica
- Cluster set
- Quorum

Explanation

Storage Replica is the correct answer. Storage Replica provides block-level replication for any type of data in complete volumes. This allows disaster recovery in stretch cluster, cluster-to-cluster, or server-to-server situations.

Question 2

Which term is defined as the majority of voting nodes in an active cluster membership plus a witness vote?

- Failover voting
- CSV
- Cluster set
- Quorum

Explanation

Quorum is the correct answer. A quorum is the majority of voting nodes in an active cluster membership plus a witness vote. In effect, each cluster node is an element that can cast one vote to determine whether the cluster continues to run. In case an even number of nodes exists, another element, referred to as a "witness," is assigned to the cluster. The witness element can be a disk, a file share, or a Microsoft Azure Cloud Witness. Each voting element contains a copy of the cluster configuration, and the Cluster service works to always keep all the copies synced.

Question 3

What quorum configuration is a best practice for Windows Server 2019 failover clusters?

Dynamic quorum mode and dynamic witness provide the highest level of scalability for a cluster in most standard configurations.

Question 1

Does Windows Server 2019 require all nodes to be in the same domain?

- Yes
- No

Explanation

No is the correct answer. Windows Server 2019 doesn't require all nodes to be in the same domain; however, we recommend having all nodes in the same domain.

Question 2

Can a node that runs Windows Server 2016 and one that runs Windows Server 2019 both run in the same cluster?

- Yes
- No

Explanation

Yes is the correct answer. A node that runs Windows Server 2016 and one that runs Windows Server 2019 both can run in the same cluster. This is part of the Cluster Operating System Rolling Upgrade feature that's new in Windows Server 2016. It's a best practice to move toward having the cluster run the same operating system and not run in mixed mode for an extended period.

Question 3

You must install what feature on every server that you want to add as a failover cluster node?

- Cluster set
- Failback Clustering
- Hyper-V
- Failover Clustering

Explanation

Failover Clustering is the correct answer. You must install the Failover Clustering feature on every server that you want to add as a failover cluster node.

Question 4

When running the Validate a Configuration Wizard, what does the yellow yield symbol indicate?

- The failover cluster needs to fail back to the original node.
- The wizard is waiting for a file to download.
- The failover cluster creation is in progress.
- The failover cluster that's being tested isn't in alignment with Microsoft best practices.

Explanation

When running the Validate a Configuration Wizard, the yellow yield symbol indicates that the aspect of the proposed failover cluster that's being tested isn't in alignment with Microsoft best practices. Investigate this aspect to make sure that the configuration of the cluster is acceptable for the environment of the cluster, for the requirements of the cluster, and for the roles that the cluster hosts.

Question 1

Which type of witness uses a basic format and doesn't keep a copy of the cluster database?

- USB witness
- Failback witness
- File share witness
- Microsoft Azure Cloud Witness

Explanation

Microsoft Azure Cloud Witness is the correct answer. 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 it doesn't keep a copy of the cluster database.

Question 2

What technology enables replication of volumes between servers or clusters for disaster recovery?

- File share witness
- Cluster set
- Cluster Shared Volume (CSV)
- Storage Replica

Explanation

Storage Replica is the correct answer. Storage Replica is Windows Server technology that enables replication of volumes between servers or clusters for disaster recovery. With it, you can also create stretch failover clusters that span two sites, with all nodes staying in sync.

Question 3

What added features does enabling site-aware clustering in a stretch cluster provide?

Your answers might vary, but they might include:

Question 1

Which feature would you use to configure a failover cluster when you use Hyper-V host servers?

- Site-aware clustering
- Client clustering
- Live clustering
- Host clustering

Explanation

Host clustering is the correct answer. 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 machines (VMs) 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 VM don't have to be cluster-aware. However, the VM is still highly available.

Question 2

Which feature can you use to transparently move running VMs from one Hyper-V host to another without perceived downtime?

- Site-aware cluster
- Storage migration
- Cluster set
- Live Migration

Explanation

Live Migration is the correct answer. You can use Live Migration, a Hyper-V feature in Windows Server, to transparently move running VMs from one Hyper-V host to another without perceived downtime.

Question 1

What term describes a loosely coupled grouping of multiple failover clusters?

- Cluster set
- Failback Clustering
- Hyper-V
- Failover Clustering

Explanation

Cluster set is the correct answer. A cluster set is a loosely coupled grouping of multiple failover clusters; it enables virtual machine (VM) fluidity across member clusters within the set and a unified storage namespace across the set.

Question 2

When running the Validate a Configuration Wizard, what does the red "X" indicator mean?

- The failover cluster needs to fail back to the original node.
- You can't use the part of the failover cluster that failed.
- Failover cluster creation is in progress.
- The failover cluster that's being tested isn't in alignment with Microsoft best practices.

Explanation

When running the Validate a Configuration Wizard, when a failover cluster receives a red "X" (fail) in one of the tests, it means that you can't use the part of the failover cluster that failed in a Windows Server failover cluster. Additionally, when a test fails, all other tests don't run, and you must resolve the issue before you install the failover cluster.

Question 3

What component provides a consistent, distributed namespace that clustered roles can use to access shared storage from all nodes?

- Storage Replica
- Cluster Shared Volume (CSV)
- Cluster set
- Quorum

Explanation

Cluster Shared Volume (CSV) is the correct answer. Failover clusters provide CSV functionality that provides a consistent, distributed namespace that clustered roles can use to access shared storage from all nodes. With the Failover Clustering feature, users experience a minimum of disruptions in service.

Question 4

Which type of witness is ideal when shared storage isn't available or when the cluster spans geographical locations?

- USB witness
- Failback witness
- File share witness
- Microsoft Azure Cloud Witness

Explanation

File share witness is the correct answer. A file share witness is ideal when shared storage isn't available or when the cluster spans geographical locations. This option doesn't store a copy of the cluster database.

Question 5

What technology provides high availability where each site has a separate storage system with replication among the sites?

- Stretch cluster
- Cluster set
- CSV
- Storage Replica

Explanation

Stretch cluster is the correct answer. A stretch cluster provides high availability where each site has a separate storage system with replication among the sites.

Question 6

Which feature provides high availability for applications or services running on the VM that don't have to be compatible with failover clustering?

- Site-aware clustering
- Client clustering
- Live clustering
- Host clustering

Explanation

Host clustering is the correct answer. Host clustering provides high availability for applications or services running in the VM that don't have to be compatible with failover clustering; additionally, they don't have to be aware that the VM is clustered. Because the failover is at the VM-level, there are no dependencies on the software that's installed in the VM.

Question 7

Which feature distributes IP traffic to multiple instances of a TCP/IP service?

- Site-aware cluster
- Storage migration
- Network Load Balancing (NLB)
- Live Migration

Explanation

Network Load Balancing (NLB) is the correct answer. NLB distributes IP traffic to multiple instances of a TCP/IP service.



Module 7 Disaster recovery in Windows Server

Hyper-V Replica

Lesson overview

Hyper-V Replica is a disaster recovery feature that's built into Hyper-V. You can use it to replicate a running virtual machine (VM) to a secondary location, a tertiary location, or to Microsoft Azure. While the live VM is running, its **Hyper-V Replica** VM is offline. When you update a Hyper-V host or when necessary, you can failover from the replica VM. Failovers are performed manually and can be either test failovers, planned failovers, or unplanned failovers. Planned failovers are without data loss, while unplanned failovers can cause the loss of recent changes—up to five minutes by default.

Lesson objectives

After completing this lesson, you'll be able to:

- Describe **Hyper-V Replica**.
- Plan for **Hyper-V Replica** configuration.
- Configure and implement **Hyper-V Replica**.
- Describe Azure Site Recovery.

Overview of Hyper-V Replica

Hyper-V failover clusters are used to make virtual machines (VMs) highly available, but they're often limited to a single location. Multisite clusters usually depend on specialized hardware and are expensive to implement, even with Windows Server Storage Replica. In case of a natural disaster such as an earthquake or a flood, all server infrastructure at the affected location can be lost.

Hyper-V Replica can protect against data loss from natural disasters, and it can be used to implement an affordable business continuity and disaster recovery (BCDR) solution for a virtual environment. Use **Hyper-V Replica** to replicate VMs to a Hyper-V host in a secondary location across a wide area network (WAN) link and even to a third location. If you have a single location, you can still use **Hyper-V Replica** to replicate VMs to a partner organization in another location, to a hosting provider, or to Microsoft Azure.

Hyper-V hosts that participate in replication don't have to be in the same Active Directory forest or have the same configuration. You can also encrypt network traffic between them.

Hyper-V Replica can have two instances of a single VM residing on different Hyper-V hosts. One of the instances will be the primary, running VM, and the other instance will be a replica—an offline copy. If necessary, you can even extend replication of the offline copy to a third location. Hyper-V syncs these instances, and you can perform manual failover at any time. If a failure occurs at a primary site, you can use **Hyper-V Replica** to perform a failover of the VMs to Replica servers at a secondary location with minimal downtime.

Prerequisites for Hyper-V Replica implementation

Before implementing **Hyper-V Replica**, ensure that the virtualization infrastructure has the following prerequisites:

- Windows Server 2012 or a newer version of Windows Server with the Hyper-V role installed at both locations. Server hardware should have enough capacity to run all VMs: the local VMs and the replicated VMs. Replicated VMs are in a turned-off state and start only if you perform a failover.
- Sufficient storage on both the primary and replica Hyper-V hosts to store both local and replicated VM data.
- Network connectivity between the locations that are hosting the primary and the replica Hyper-V hosts. Connectivity can be through a WAN or a local area network link.
- Firewall rules to allow replication between the primary and replica sites. When you install the Hyper-V role, Hyper-V Replica HTTP Listener (TCP-In) and Hyper-V Replica HTTPS Listener (TCP-In) rules are added to Windows Defender Firewall. Before you can use **Hyper-V Replica**, you must enable one or both rules on the replica Hyper-V host.
- An X.509v3 certificate from a trusted certification authority to support mutual authentication at both Hyper-V hosts if you plan to use certificate-based authentication. When you use certificate-based authentication, Hyper-V hosts can be in different Active Directory forests.
- Both Hyper-V hosts joined to the same Active Directory forest if you intend to use Kerberos authentication.

Hyper-V Replica high-level architecture

When you configure a VM for replication, it performs an initial replication and creates a copy of the VM on the second Hyper-V host at the recovery site. The replicated VM stays turned off until you initiate a failover, while the primary VM keeps running. Changes in the primary VM are written in a log file that's periodically replicated and applied to the replica. **Hyper-V Replica** has several components:

- Replication engine. This component manages the initial replication, replication configuration details, replication of delta changes, and failover and test failover operations. It also tracks VM and storage mobility events and takes appropriate actions when necessary. For example, it pauses replication and **Hyper-V Replica** configurations when the source or the replica Hyper-V hosts are part of a Hyper-V failover cluster.
- Change tracking module. This component tracks changes that occur to the VM on a source Hyper-V host. The change tracking module tracks write operations to the virtual hard disks (VHDs) regardless of where the VHDs are stored locally—on a storage area network, on a Server Message Block version 3 or newer share, or on a Cluster Shared Volume.

- Network module. This component provides a secure and efficient way to transfer VM data between Hyper-V hosts. By default, the network module minimizes traffic by compressing data. It can also encrypt data when HTTPS and certification-based authentication are used.
- **Hyper-V Replica Broker.** This component is used only when a Hyper-V host is a node in a failover cluster. **Hyper-V Replica Broker** enables you to use **Hyper-V Replica** with highly available VMs that can move between cluster nodes. The **Hyper-V Replica Broker** role queries the cluster database. It then redirects all requests to the cluster node where the VM is currently running.
- Management tools. With tools such as Hyper-V Manager and Windows PowerShell, you can configure and manage **Hyper-V Replica**. Use **Failover Cluster Manager** for all VM management and **Hyper-V Replica** configurations when the source or the replica Hyper-V hosts are part of a Hyper-V failover cluster.

Hyper-V Replica security considerations

You can set up **Hyper-V Replica** with a Hyper-V host regardless of its location and domain membership if you have network connectivity with it. Hyper-V hosts don't have to be part of the same Active Directory domain. You can implement **Hyper-V Replica** when Hyper-V hosts are members of untrusted domains by configuring certificate-based authentication. **Hyper-V Replica** implements security at the following levels:

- On each server, Hyper-V creates a local security group named **Hyper-V Administrators**. Members of this group, in addition to local administrators, can configure and manage **Hyper-V Replica**.
- You can configure a Replica server to allow replication from any authenticated server or to limit replication to specific servers. In the first case, you must specify a fully qualified domain name for the primary server (for example, `lon-svr1.adatum.com`), or use a wildcard character with a domain suffix (for example, `*.adatum.com`). Using IP addresses isn't allowed. If the Replica server is in a failover cluster, replication is allowed at the cluster level. When you limit replication to specific servers, you also must specify a trust group, which is used to identify the servers within which a VM can move. For example, if you provide disaster recovery services to partner organizations, the trust group prevents one organization from gaining access to another organization's replica machines.
- The replica Hyper-V host can authenticate a primary Hyper-V host by using Kerberos authentication or certificate-based authentication. Kerberos authentication requires both Hyper-V hosts to be in the same Active Directory forest, while you can use certificate-based authentication in any environment. Kerberos authentication is used with HTTP traffic, which isn't encrypted, while certificate-based authentication is used with HTTPS traffic, which is encrypted.
- You can establish **Hyper-V Replica** only if network connectivity exists between the Hyper-V hosts. Configure Windows Defender Firewall to allow HTTP or HTTPS **Hyper-V Replica** traffic as needed.

Plan for Hyper-V Replica

When planning for **Hyper-V Replica** deployment, you must define several parameters used in replica configuration. Careful planning is important before setting up replication between Hyper-V hosts.

Hyper-V Replica configurations

You can set up **Hyper-V Replica** between Hyper-V hosts irrespective of whether they're nodes in a failover cluster. You can also set up **Hyper-V Replica** irrespective of whether the Hyper-V hosts are

members of the same Active Directory forest or are in different Active Directory forests without any trust between them. You can use **Hyper-V Replica** in four different configurations:

- Both Hyper-V hosts are standalone servers. This configuration isn't recommended, because it includes only disaster recovery and not high availability.
- The Hyper-V host at the primary location is a node in a failover cluster, and the Hyper-V host at the secondary location is on a standalone server. Many environments use this type of implementation. A failover cluster provides high availability for running virtual machines (VMs) at the primary location. If a disaster occurs at the primary location, a replica of the VMs is still available at the secondary location.
- Each Hyper-V host is a node in a different failover cluster. This enables you to perform a manual failover and continue operations from a secondary location if a disaster occurs at the primary location.
- The Hyper-V host at the primary location is a standalone server, and the Hyper-V host at the secondary location is a node in a failover cluster. Although technically possible, this configuration is rare. You typically want VMs at the primary location to be highly available, while their replicas at the secondary location are turned off and aren't used until a disaster occurs at the primary location.

Note: You can configure **Hyper-V Replica** regardless of whether the Hyper-V host is a node in a failover cluster.

Replication settings

Because you must configure replication for each VM individually, you also must plan resources for each VM on replication hosts. Besides resources, you also must plan on how to configure the following replication settings:

- **Replica Server.** Specify the computer name or the fully qualified domain name (FQDN) of the Replica server—an IP address isn't allowed. If the Hyper-V host that you specify isn't yet configured to allow replication traffic, you can configure it here. If the Replica server is a node in a failover cluster, you should enter the name or FQDN of the connection point for the **Hyper-V Replica Broker**.
- **Connection Parameters.** If the Replica server is accessible, the **Enable Replication Wizard** populates the authentication type and replication port fields automatically with appropriate values. If the Replica server is inaccessible, you can configure these fields manually. However, you should be aware that you won't be able to enable replication if you can't create a connection to the Replica server. On the **Connection Parameters** page, you can also configure Hyper-V to compress replication data before transmitting it over a network.
- **Replication virtual hard disks.** By default, all virtual hard disks (VHDs) are replicated. If some of the VHDs aren't required at the replica Hyper-V host, exclude them from replication; for example, a VHD that's dedicated to storing page files. Excluding VHDs that include operating systems or applications can result in that particular VM being unusable at the Replica server.
- **Replication Frequency.** You can set replication frequency to 30 seconds, 5 minutes, or 15 minutes based on the network link to the Replica server and the acceptable state delay between primary and replica VMs. Replication frequency controls how often data replicates to the Hyper-V host at the recovery site. If a disaster occurs at the primary site, a shorter replication frequency means less loss as fewer changes aren't replicated to the recovery site.
- **Additional recovery points.** You can configure the number and types of recovery points to send to a Replica server. By default, the option to maintain only the latest point for recovery is selected, which means that only the parent VHD replicates. All changes merge into that VHD. However, you can choose to create more hourly recovery points and then set the number of additional recovery points

(up to 24). You can configure the Volume Shadow Copy Service snapshot frequency to save application-consistent replicas for the VM and not just the changes in the primary VM.

- **Initial replication method and schedule.** VMs have large virtual disks, and initial replication can take a long time and cause a lot of network traffic. While the default option is to immediately send the initial copy over the network, if you don't want immediate replication, you can schedule it to start at a specific time. If you want an initial replication but want to avoid network traffic, you can opt to send the initial copy to external media or use an existing VM on the Replica server. Use the last option if you restored a copy of the VM at the Replica server and you want to use it as the initial copy.
- **Extended replication.** With Windows Server 2012 R2 and later Windows Server operating systems, you can replicate a single VM to a third server. Thus, you can replicate a running VM to two independent servers. However, the replication doesn't happen from one server to the two other servers. The server that's running an active copy of the VM replicates to the Replica server, and the Replica server then replicates to the extended Replica server. You create a second replica by running the **Extend Replication Wizard** on a passive copy. In this wizard, you can set the same options that you chose when you configured the first replica.

Note:Hyper-V Replica now allows administrators to use a Microsoft Azure instance as a replica repository. This enables administrators to take advantage of Azure rather than having to build out a disaster recovery site or manage backup tapes offsite. To use Azure for this purpose, you must have a valid subscription. Note that this service might not be available in all world regions.

Configure and implement Hyper-V Replica

Hyper-V Replica implements as part of the Hyper-V role. You can use it on standalone Hyper-V servers or on servers that are part of a failover cluster, in which case you should configure **Hyper-V Replica Broker**. Unlike failover clustering, the Hyper-V role doesn't depend on Active Directory. You can use the Hyper-V role with standalone Hyper-V servers or servers that are members of different Active Directory domains, except when servers that participate in **Hyper-V Replica** are part of the same failover cluster. To enable **Hyper-V Replica** technology, complete following steps:

- In the **Replication Configuration** group of options, enable the Hyper-V server as a Replica server.
- Configure Hyper-V server settings. Select the authentication and port options, and then configure the authorization options. You can choose to enable replication from any server that successfully authenticates. This is convenient in scenarios where all servers are part of the same domain, or you can enter the fully qualified domain names (FQDNs) of servers that you accept as Replica servers. Additionally, you must configure the location for replica files. You should configure these settings on each server that serves as a Replica server.
- Specify both the Replica server name and the connection options.
- Select the virtual hard disks (VHDs) to replicate in cases where a virtual machine (VM) has more than one VHD. You can also configure the recovery history and the initial replication method. Configure the replication interval for 30 seconds, 5 minutes (the default value), or 15 minutes.
- After configuring these options, you can start replication. After you make the first replica, you can also make an extended replica to a third physical or cloud-based instance that's running Hyper-V. The extended replica site is built from the first replica site, not from the primary VM. It's possible to configure different replication intervals for the replica and extended replica instances of a VM.

After you establish the replication relationship, the **Status** column in Hyper-V Manager displays the replication progress as a percentage of the total replication for the configured VM. The VM replica is in a turned-off state and will start only when you perform a failover.

After the initial replication is done, the replica updates regularly with changes from the primary VM. One of the configuration steps is configuring the replication frequency setting. This setting controls the longest time interval until changes from the primary VM are applied to the replica. In a real-world environment, however, there can be many reasons why changes from a primary VM aren't applied to the replica for extended periods; for example, because network connectivity is lost or because you pause the replication. This will be reflected in replication health, but when replication is established again, all changes will be applied to the replica.

When you enable replication, VM network adapters receive more settings that were previously unavailable. These new settings pages are **Failover TCP/IP** and **Test Failover**. Failover TCP/IP is available only for network adapters and not for legacy network adapters. The settings on this page are useful when a VM has a static IP address assigned and the replica site is using IP settings different from the primary site. You can configure the TCP/IP settings that a network adapter will use after a failover is performed. If you use static IP addresses to configure VMs, you should configure failover TCP/IP settings on the primary and replica VMs. VMs must also have integration services installed to be able to apply failover TCP/IP settings.

Replication health monitoring

When you enable replication for a VM, changes in the primary VM write to a log file, which periodically transfers to the replica Hyper-V host and is then applied to a VHD of a replica VM. Replication health monitors the replication process and displays important events in addition to the replication and sync state of the Hyper-V host.

Replication health includes the following data:

- **Replication State.** This indicates whether replication is enabled for a VM.
- **Replication Type.** This indicates whether you're monitoring replication health on a primary VM or replica VM.
- **Primary and Replica server names.** This indicates which Hyper-V host the primary VM is running on and which Hyper-V host is the replica.
- **Replication Health.** This indicates replication status. Replication health can have one of three values: Normal, Warning, or Critical.
- **Replication statistics.** This displays replication statistics since the time that the VM replication started or since you reset the statistics. Statistics include data such as maximum and average sizes of a replication, average replication latency, number of errors encountered, and the number of successful replication cycles.
- **Pending replication.** This displays information about the size of data that still needs to replicate and when the replica was last synced with the primary VM.

Failover options

Three types of failovers are possible with **Hyper-V Replica**: test failover, planned failover, and failover:

- **Test failover.** A test failover is a nondisruptive task that enables you to test a VM on a Replica server while the primary VM is running without interrupting the replication. You can perform it after you configure **Hyper-V Replica** and after the VMs start replicating. Initiating a test failover on a replicated VM creates a new checkpoint, and you can use this checkpoint to select a recovery point from which to create a new test VM. The test VM has the same name as the replica, but with "- Test" appended to the end. The test VM stays disconnected by default to avoid potential conflicts with the running primary VM. After you finish testing, to stop the test VM and delete it from the replica Hyper-V host,

stop the test failover. This option is available only if a test failover is running. If you run a test failover on a failover cluster, you'll have to manually remove the Test-Failover role from the failover cluster.

- **Planned failover.** You can start a planned failover to move the primary VM to a replica site, for example, before site maintenance or before an expected disaster. Because this is a planned event, no data loss will occur, but the VM will be unavailable for some time during its startup. A planned failover confirms that the primary VM is turned off before the failover runs. During the failover, the primary VM sends all the data that it hasn't yet replicated to the Replica server. The planned failover process then fails over the VM to the Replica server and starts the VM on the Replica server. After the planned failover, the VM will run on the Replica server, and it doesn't replicate its changes. If you want to set up replication again, you should reverse the replication. You'll have to configure settings similar to when you enabled replication, and it will use the existing VM as an initial copy.
- **Failover.** If a disruption occurs at the primary site, you can perform a failover. You start a failover at the replicated VM only if the primary VM is either unavailable or is turned off. A failover is an unplanned event that can result in data loss because changes at the primary VM might not have replicated before the disaster happened. The replication frequency setting controls how often changes replicate. During a failover, the VM runs on a Replica server. If you start the failover from a different recovery point and discard all the changes, you can cancel the failover. After you recover the primary site, you can reverse the replication direction to reestablish replication. This also removes the option to cancel failover.

Configuration options for replication

Besides performing various types of failovers, you can configure several other options for replication. Other Hyper-V replication-related actions include:

- **Pause Replication.** This action pauses replication for the selected VM.
- **Resume Replication.** This action resumes replication for the selected VM. It's available only if replication for the VM is paused.
- **View Replication Health.** This action provides data about the replication events for a VM.
- **Extend Replication.** This action is available on the replica VMs, and it extends VM replication from a Replica server to a third server (the extended Replica server).
- **Remove Recovery Points.** This action is available only during a failover. If you select it, all recovery points (checkpoints) for a replica VM are deleted, and their differing VHDs are merged.
- **Remove Replication.** This action stops replication for the VM.

Demonstration: Implement Hyper-V Replica

In this demonstration, you'll learn how to implement the **Hyper-V Replica** feature on a Windows Server computer that's running Server Core.

Demonstration steps

1. On **SEA-ADM1**, open Windows PowerShell as an administrator.
2. In PowerShell, create a new remote PowerShell session to `sea-svr1.contoso.com`. Use **Contoso\Administrator** credentials to connect to the remote PowerShell on **SEA-SVR1**.
3. In the remote PowerShell session on `sea-svr1.contoso.com`, use the **Enable-Netfirewallrule** cmdlet to enable the firewall rule named Hyper-V Replica HTTP Listener (TCP-In).

4. Use the **Get-Netfirewallrule** cmdlet to verify that the Hyper-V Replica HTTP Listener (TCP-In) rule is enabled.
5. Use the following command to configure **SEA-SVR1** for **Hyper-V Replica**:

```
powershellSet-VMReplicationServer -ReplicationEnabled $true -AllowedAuthenticationType Kerberos -ReplicationAllowedFromAnyServer $true -DefaultStorageLocation c:\ReplicaStorage
```
6. Use the **Get-VM** cmdlet to verify that the **SEA-CORE1** virtual machine (VM) is present on **SEA-SVR1**.
7. Open a new remote PowerShell session for `sea-svr2.contoso.com` in a new PowerShell window. Repeat steps 2 through 5 to configure **SEA-SVR2** for **Hyper-V Replica**.
8. Switch to the PowerShell window where you have the remote PowerShell session opened for `sea-svr1.contoso.com`, enter the following command, and then select Enter:

```
powershellEnable-VMReplication SEA-CORE1 -ReplicaServerName SEA-SVR2.contoso.com -ReplicaServerPort 80 -AuthenticationType Kerberos -computername SEA-SVR1.contoso.com
```
9. Start replication with the following command:

```
powershellStart-VMInitialReplication SEA-CORE1
```

Overview of Azure Site Recovery

Organizations implement business continuity and disaster recovery (BCDR) processes to help ensure that their applications, workloads, and data are available during planned and unplanned downtime so that they can return to working conditions after a disaster as quickly as possible. Site Recovery is a BCDR solution that can replicate on-premises virtual machines (VMs) and physical servers to Microsoft Azure or to a second site. You can use it to monitor service availability and perform automatic orchestrated site recovery, which helps ensure better data management and business continuity in the event of an outage.

Site Recovery is a hybrid product that mixes onsite and cloud automation. While your primary site might be on-premises, the secondary site can be on-premises in the same datacenter, in a geographically separate datacenter, or in Azure. After a secondary site is in place, you can use it to test the failover and recovery of your primary site. Because Site Recovery is used to replicate on-premises server loads to Azure or a secondary datacenter, and because it can provide orchestrated failover, you can use it in different scenarios, such as:

- Disaster recovery. Servers replicate to a secondary location and can fail over if the primary location becomes unavailable, such as in a disaster. Site Recovery can control and orchestrate failover, ensuring that at the secondary location, the workload starts in a specific order and that additional preparation tasks are performed.
- Workload migration to Azure. Physical servers that run Windows or Linux operating systems can be virtualized and moved to Azure. Many servers are already virtualized, but if you're still running a physical server, use Site Recovery to virtualize it.
- Cloud bursting. Move applications to Azure temporarily when workload demands exceed on-premises capacity and then bring them back when demand stabilizes. You can use Azure scalability to eliminate long and costly procurement cycles for more hardware.

- DevTest. Replicate workloads to Azure for testing purposes so that you don't need to buy and maintain an onsite test environment. You can safely test with replicated live data without affecting users or production environments.
- Analytics and reporting. Replicate workloads to run reports, check the health of applications, and improve performance. You can analyze production workloads by running compute-intensive diagnostics without affecting users. You can understand where performance issues occur by removing infrastructure variables through cloud replication.

The following table details which types of machines or servers that Site Recovery can replicate and to which locations.

Table 1: Types of machines or servers that Site Recovery can replicate

Replicate	Replicate from	Replicate to
Hyper-V VM	Hyper-V host in a Microsoft System Center Virtual Machine Manager (VMM) cloud	Azure
Hyper-V VM	Hyper-V host in a VMM cloud	VMM cloud in a secondary site
Hyper-V VM	Hyper-V host in a VMM cloud with storage area network (SAN) storage	VMM cloud in a secondary site with SAN storage
Hyper-V VM	Hyper-V host without VMM	Azure
VMWare VM	VMWare server	Azure
VMWare VM	VMWare server	Secondary VMWare site
Physical Windows Server or Linux servers	Physical server	Azure
Workloads that are running on physical Windows Server or Linux servers	Physical server	Secondary datacenter

Note: Site Recovery is workload-agnostic. You can use it for protecting all kinds of workloads, including Microsoft SharePoint Server, Microsoft Exchange Server, Microsoft Dynamics CRM, Microsoft SQL Server, Active Directory Domain Services, and third-party applications.

Use Site Recovery to:

- Configure and monitor the asynchronous Hyper-V replication of VMs between locations by using the **Hyper-V Replica** feature.
- Manage and orchestrate VM failover in VMM clouds. In this case, Site Recovery can protect only the VMs in VMM clouds. VMs running on your virtualization infrastructure that are managed by VMM but aren't deployed to the VMM cloud can't be protected.
- Manage virtual network mapping between the primary and secondary locations so that you don't need to change the TCP/IP settings of VMs that fail over between the primary and secondary locations.
- Continually monitor service availability.
- Test a recovery configuration by performing a test failover without affecting the primary location.
- Perform an automatic orchestrated site recovery by using recovery plans. Note: A recovery plan is a plan that defines the required steps to recover applications and services in the correct sequence. Plans include triggers for invocation, manual actions, and custom scripts.

- Configure and control VM replication between two on-premises locations and orchestrate failover if a disaster occurs. In such configurations, data never replicates to Azure—Site Recovery only monitors replication between sites and controls failover.
- Act as a standalone solution for virtualizing physical servers, replicating VMs to Azure and performing failover.

Site Recovery uses the **Hyper-V Replica** feature to protect VMs on Hyper-V servers or in VMM clouds. Therefore, Site Recovery provides the same types of failover as **Hyper-V Replica**, but with Site Recovery, you can further automate and orchestrate the failover of multiple VMs between VMM clouds. You can also use Site Recovery at a third location when replicating VMs between two Hyper-V hosts. Site Recovery supports following failover types:

- Test failover
- Planned failover
- Unplanned failover

Implement Site Recovery

If you want to implement Site Recovery, you must meet several prerequisites. Prerequisites vary based on the scenario that you want to implement and whether you want to use Site Recovery to replicate workloads in Azure or between two on-premises datacenters. To use Site Recovery, you need an Azure subscription because Site Recovery is one of the services that Azure offers.

It's recommended to use VMM in your local environment to establish and manage replication to Site Recovery. You can use Site Recovery without having VMM deployed by choosing the option to establish replication between your on-premises Hyper-V and Azure.

To implement Site Recovery with VMM, perform the following high-level configuration steps:

1. Create an Azure Recovery Services vault or a Site Recovery vault. In a Recovery Services vault, you can specify the scenario in which you want to use Site Recovery, register a VMM (or VMMs) with Site Recovery, configure replication settings, and manage recovery plans. You can have multiple Recovery Services vaults in a single Azure subscription.
2. Register VMM with Site Recovery. Based on the scenario that you want to implement and the protection goals that you configure in the Recovery Services vault, you can download a Site Recovery provider and vault registration key. When you install the provider on a VMM management server, it registers the VMM with Site Recovery and sends configuration data about the cloud (or clouds) that are defined in the VMM.
3. Prepare the infrastructure for Site Recovery. In the infrastructure-planning step, you must specify the location of the machines with which you want to use Site Recovery. You must also specify where they should be replicated, if they're virtualized, and if you have completed deployment planning. As part of this preparation step, you must also register the VMM with Site Recovery. After registering the VMM, you can configure the cloud environments that you want to protect and the cloud to which the machines should be replicated, and then create and associate a replication policy. A *replication policy* defines the replication settings, such as the copy frequency, authentication type, and data transfer compression. As part of this step, you can also download and run the Site Recovery Capacity Planner to more accurately estimate network bandwidth, storage, and other requirements to meet your replication needs.
4. Replicate the application. In this step, you enable the replication of VMs from the protected cloud. You can enable replication of a single VM or multiple VMs based on the infrastructure that you configured in the previous step.

5. Manage the recovery plan. A recovery plan controls how a Site Recovery failover is performed. It specifies the order in which VMs should start at the secondary location and more actions that should be performed during the failover. You can also specify which VMs are included in the recovery plan, perform test failover, planned failover, unplanned failover, and reverse replication after failover. Site Recovery can orchestrate the failover of multiple VMs between the primary and secondary VMM cloud.

Test your knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

What's the difference between a planned failover and a failover?

Question 2

Can you use Hyper-V Replica to replicate only VMs that have integration services installed?

Backup and restore infrastructure in Windows Server

Lesson overview

Having a backup infrastructure is mandatory for most organizations. You can choose to use built-in tools to perform backups, or you can use third-party tools. Windows Server has built-in backup software called Windows Server Backup. You can use this software for simple backup and restore tasks. In this lesson, you'll learn about Windows Server Backup and Microsoft Azure Backup.

Lesson objectives

After completing this lesson, you'll be able to:

- Describe Windows Server Backup.
- Implement backup and restore with Windows Server Backup.
- Back up and restore Hyper-V virtual machines.
- Describe Azure Backup.
- Implement backup and restore with Azure Backup.

Overview of Windows Server Backup

Windows Server Backup is a feature in Windows Server that consists of a **Microsoft Management Console (MMC)** snap-in, the **wbadmin** command, and Windows PowerShell commands. With Windows Server Backup, you can perform backup and recovery in a Windows Server environment.

Use the wizards in Windows Server Backup to guide you through running backups and recoveries. You can also configure backup jobs by using Windows PowerShell.

You can use Windows Server Backup to back up:

- A full server (all volumes) or just selected volumes.
- Individual files and folders.
- System state.
- Individual virtual machines on a Hyper-V host.
- Microsoft Exchange Server databases.
- Cluster Shared Volumes.

Additionally, Windows Server Backup allows you to perform some specific backup and restore tasks, such as:

- Performing a bare-metal restore. A bare-metal backup contains at least all critical volumes, and it allows you to restore without first installing an operating system. You do this by using the product media on a DVD or USB key and the Windows Recovery Environment (Windows RE). You can use this backup type together with the Windows RE to recover from a hard disk failure or to recover a whole computer image to new hardware.
- Restoring system state. The backup contains all information to roll back a server to a specific time. However, you must install an operating system before you can recover the system state.

- Recovering individual files and folders or volumes. The Individual files and folders option enables you to selectively back up and restore specific files, folders, or volumes. You can add specific files, folders, or volumes to a backup even when you use an option such as critical volume or system state.
- Excluding selected files or file types. You can exclude unwanted files or file types, such as temporary files, from a backup.
- Storing backups in multiple storage locations. You can store backups on remote shares or non-dedicated volumes.
- Performing backups to Azure. Azure Online Backup is a cloud-based backup solution for Windows Server that enables you to back up and recover files and folders offsite by using cloud services. You can use Windows Server Backup with an appropriate agent to store backups in Azure.

Note: Windows Server Backup is a single-server backup solution. To back up multiple servers, you must install and configure Windows Server Backup on each server.

WBAdmin (WBAdmin.exe¹) is a command-line tool that's built into Windows Server. The command is used to perform backups and restores of operating systems, drive volumes, files, folders, and applications from a command-line interface.

By default, Windows Server Backup isn't installed. You can install it from **Server Manager** Add Roles and Features or with the Windows PowerShell **Add-WindowsFeature Windows-Server-Backup -Include-AllSubfeature** cmdlet. You can also use the **Windows Admin Center** to install it.

Volume Shadow Copy Service

Correctly performing backup and restore operations requires close coordination between the backup applications, the line-of-business applications that are being backed up, and the storage management hardware and software. Windows Server Backup uses Volume Shadow Copy Service (VSS) to perform backups. VSS facilitates the necessary conversations between these components to allow them to work better together. VSS coordinates the actions that are necessary to create a consistent shadow copy, also known as a *snapshot* or a *point-in-time copy*, of the data that's to be backed up.

VSS solutions have the following basic parts:

- VSS service. This is part of the Windows operating system, and it ensures that the other components can communicate with each other properly and work together.
- VSS requester. This software requests the actual creation of shadow copies or other high-level operations like importing or deleting them. Typically, this is the backup application, such as Windows Server Backup.
- VSS writer. This component guarantees that you have a consistent dataset to back up. This is typically part of a line-of-business application such as Microsoft SQL Server or Microsoft Exchange Server. The Windows operating system includes VSS writers for various Windows components such as the registry.
- VSS provider. This component creates and maintains the shadow copies. This can occur in the software or in the hardware. The Windows operating system includes a VSS provider that uses copy-on-write.

¹ <https://aka.ms/wbadmin>

Implement backup and restore

Depending on what you must back up, the procedures and options in Windows Server Backup might vary. In this topic, you'll discuss various backup options, which depend on the scenario and the resources being backed up.

Back up file servers and web servers

Consider a scenario where you want to provide a backup for a file or a web server. You should protect your server and its data automatically by scheduling daily backups. It's recommended that you have a daily backup plan because most organizations can't afford to lose data that's been created over several days.

When performing a backup of a web or file server with Windows Server Backup software, you should choose the specific folders that contain critical files or resources for the applications that you're running.

Back up AD DS

Backing up the Active Directory Domain Services (AD DS) role should be an important part of any backup and recovery process or strategy. An AD DS role backup can restore data in different data-loss scenarios, such as deleted data or a corrupted AD DS database.

You can perform three types of backups using Windows Server Backup to back up AD DS on a domain controller. You can also use all three backup types to restore AD DS. A full server backup contains all the volumes on a domain controller. To back up only the files that are required to recover AD DS, you can perform a system state backup or a critical-volumes backup.

A system state backup isn't incremental. Therefore, each system state backup requires a similar amount of space. A critical-volumes backup is incremental, which means that it includes only the difference between the current backup and the previous backup. However, because a critical-volumes backup can include other files in addition to the volumes that are required for system state, you can expect critical-volume backups to grow with unnecessary files over time.

When you back up AD DS, consider your backup schedule. Plan your AD DS backup schedule properly because you can't restore from a backup that's older than 180 days—the deleted object lifetime. When a user removes an object from AD DS, it keeps the information about that deletion for 180 days. If you have a backup that's newer than 180 days, you can successfully restore the deleted object. If the backup is older than 180 days, however, the restore procedure won't replicate the restored object to other domain controllers, which means that the state of AD DS data will be inconsistent.

Back up Exchange Server

The preferred architecture for Microsoft Exchange Server takes advantage of a concept known as Exchange Native Data Protection. Exchange Native Data Protection relies on native Exchange features to protect mailbox data without using traditional backups. But if you want to create backups, Exchange Server includes a plug-in for Windows Server Backup that enables you to create Exchange-aware Volume Shadow Copy Service (VSS)-based backups of Exchange data. To make Exchange-aware backups, you must have the Windows Server Backup feature installed.

The plug-in, `WSBExchange.exe`, runs as a service named Microsoft Exchange Server Extension for Windows Server Backup; the short name for this service is `WSBExchange`. This service is automatically installed and configured for manual startup on all Mailbox servers. The plug-in enables Windows Server Backup to create Exchange-aware VSS backups.

Advanced settings

When you schedule or modify a backup by using the **Backup Schedule Wizard**, you can modify the following settings:

- **Exclusions.** You can exclude file types within specific folders, and optionally, their subfolders. For example, if you back up a Hyper-V host with several virtual machines, you might not want to back up any .iso files that have been attached.
- **VSS backup.** With VSS backup options, you can select either a VSS full backup or VSS copy backup. The full backup updates the backup history and clears the log file. However, if you use other backup technologies that also use VSS, you might want to choose the VSS copy backup, which retains the VSS writer log files.

Back up and restore Hyper-V VMs

When planning a backup strategy for a virtualized server environment, you need to consider several factors. You must consider the types of backups that you can make, the state of virtual machines (VM), and the type of storage that the VMs use. This topic discusses the advantages, disadvantages, and considerations for these factors.

When creating a backup solution for VMs, you should consider the data that's being backed up. You can install Windows Server Backup on a host and perform a host-level backup, or you can install Windows Server Backup on a VM to perform an in-guest backup. In many cases, you might want to use both host and in-guest backups. It's recommended that you study the technical documentation and best practices on how to back up a specific application. For example, Microsoft SQL Server and Microsoft Exchange Server have different best practices for backups. Furthermore, some applications support only in-guest backups.

In Hyper-V environments, you can use both Windows Server Backup and Microsoft System Center Data Protection Manager to back up the VMs and the Hyper-V hosts. Third-party backup tools also exist.

You can use the following methods to back up VMs:

- **Back up the VM from the host.** Use a host-level backup when performing a full server backup where the data in the backup includes VMs' configurations, VMs' associated checkpoints, and VMs' virtual hard disks (VHDs). When restoring data from a backup, it's not necessary to recreate VMs or reinstall Windows Server roles. However, the backup doesn't include virtual network settings, which must be recreated and reattached to the VMs. For this purpose, you might create Windows PowerShell scripts that automate the process of creating and attaching the virtual switches. Most businesses today back up VMs from the host. Often, the host supports a way to flush application data to disk before performing a backup. On Hyper-V hosts, the Volume Shadow Copy Service (VSS) handles this. This backup method also makes it simpler to recover from backup if something goes wrong.
- **Back up the VM's VHDs.** Access the storage and back up the VHDs for all VMs. Because all important data from the VMs are stored in the VHDs, backing up the VHDs should be sufficient. But in a recovery scenario, this might prove ineffective because creating new VMs and attaching the VHDs would be necessary.
- **Back up inside the VM.** When you perform a backup within the guest operating system, the procedure is the same as performing a backup for a physical computer. When performing both a host-level backup and VM backup, you should complete the backup within the guest operating system before performing a full backup on the host computer. If you use backup software inside the VM, be sure that the backup software supports the workloads in the VM by making an application-consistent backup. The downside of this backup type is that the configuration of the VM on the host isn't backed up.

Note: If you create checkpoints in Hyper-V, be sure that you can return to the state that the VM was in at that point in time. Checkpoints aren't backups—if the storage fails, the VM must be recovered from backup.

In the unlikely scenario that the entire virtualization environment fails, it's necessary to have a backup of the individual virtualization hosts. If something happens to your virtualization cluster, it's also recommended that you back up the cluster nodes.

Understand online and offline backups in VMs

You can perform online backups that don't incur VM downtime if you meet the following conditions:

- The VM being backed up has integration services installed and enabled.
- Each disk that the VM uses is running NTFS file system basic disks.
- VSS is enabled on all volumes within the VM and snapshots for each volume are stored on the same volume. For example, volume **D** must store shadow copies on volume **D**.

Note: During the backup procedure, you'll receive a warning that reminds you not to mix virtual volume backups with physical disk backups.

Overview of Azure Backup

Microsoft Azure Backup is a subscription-based service that you can use to provide offsite protection against critical data loss caused by disasters. You back up files and folders and recover them from the cloud as necessary. Backup replaces your existing on-premises or offsite backup solution with a cloud-based solution that's reliable, secure, and cost-competitive. It also helps protect assets that run in the cloud.

Backup offers multiple components that you can download and deploy on a computer, server, or in the cloud. The component, or agent, that you deploy depends on what you want to protect. All Backup components—no matter whether you're protecting data on-premises or in the cloud—can be used to back up data to an Azure Recovery Services vault in Azure.

There are many reasons to consider using Backup, because it offers several features that aren't available in Windows Server Backup. Some of the most important Backup features include:

- Automatic storage management. No capital expenditure is necessary for on-premises storage devices. Backup automatically allocates and manages backup storage, and it uses a pay-as-you-use consumption model.
- Unlimited scaling. You can take advantage of high availability without the overhead of maintenance and monitoring. Backup uses the underlying power and scale of the Azure cloud, with its non-intrusive automatic scaling capabilities.
- Multiple storage options. Backup offers two types of replication to keep your storage and data highly available:
 - Locally redundant storage (LRS) replicates your data three times—it creates three copies of your data—in a storage scale unit in a datacenter. All copies of the data exist within the same region. LRS is a low-cost option for protecting your data from local hardware failures. It is ideal for price-conscious customers, and it helps protect against local hardware failures.
 - Geo-redundant storage (GRS) is the default and recommended replication option. GRS replicates your data to a secondary region that's far away from the primary location of the source data. It provides three copies in a paired datacenter. These extra copies help ensure that your backup data

is highly available even if a site-level failure occurs in Azure. GRS costs more than LRS, but it offers a higher level of durability for your data even if a regional outage occurs.

- Data encryption. Data encryption allows for highly secure transmission and storage of customer data in the public cloud. The encryption passphrase is stored at the source, and it's never transmitted or stored in Azure. The encryption key is required to restore any of the data, and only the customer has full access to the data in the service.
- Offloaded on-premises backup. Backup offers a simple solution for backing up your on-premises resources to the cloud. You can have short-term and long-term backups without needing to deploy complex on-premises backup solutions.
- Backup for Azure infrastructure as a service (IaaS) virtual machines (VMs). Backup provides independent and isolated backups to guard against accidental destruction of original data. Backups are stored in a Recovery Services vault with built-in management of recovery points. Configuration and scalability are simple, backups are optimized, and you can easily restore as needed.
- Retention of short-term and long-term data. You can use Recovery Services vaults for short-term and long-term data retention. Azure doesn't limit the length of time that data can remain in a Recovery Services vault—you can keep it for as long as you like. Backup has a limit of 9,999 recovery points per protected instance.

Implement backup and restore with Azure Backup

To use Microsoft Azure Backup, you must install its backup agent, the Azure Recovery Services agent, on your local servers and configure the Recovery Services vault. A Recovery Services vault is a storage entity in Azure that typically houses copies of data or configuration information for virtual machines (VMs), workloads, servers, or workstations. You can use Recovery Services vaults to hold backup data for various Azure services, such as infrastructure as a service VMs (Linux or Windows) and Azure SQL databases. Recovery Services vaults support Microsoft System Center Data Protection Manager, Windows Server, Windows 10, Azure Backup Server, and more. Recovery Services vaults help simplify organization of backup data while minimizing management overhead.

Creation and management of Recovery Services vaults in the Azure portal is simple because the Backup service integrates with the **Azure Settings** menu. This integration means that you can create or manage a Recovery Services vault in the context of the target service. For example, to observe the recovery points for a Contoso VM, select it, and then enter **Backup** in the **Settings** menu. The backup information specific to that VM appears. **ContosoVM-demovault** can be the name of the Recovery Services vault. You don't need to remember the name of the Recovery Services vault that stores the recovery points—you can access this information from the VM.

Note: If the same Recovery Services vault protects multiple servers, it might be more logical to refer to the Recovery Services vault. You can search for all Recovery Services vaults in your subscription and choose one from the list.

Within an Azure subscription, you can create up to 25 Recovery Services vaults per region. Backup for files and folders relies on the Recovery Services agent, which must be installed on the Windows client or server. The Recovery Services agent is a full-featured agent. Some of its key characteristics are:

- The ability to back up files and folders on physical or virtual Windows operating systems. VMs can be on-premises or in Azure.
- No requirement for a separate backup server.
- Not application-aware. File, folder, and volume-level restore only.

- The ability to back up and restore content.
- No support for Linux.

For more advanced scenarios, such as Hyper-V VMs, Microsoft SQL Server, Microsoft Exchange, Microsoft SharePoint, and system state and bare-metal recovery, you'll need Azure Backup Server. Azure Backup Server is similar to Data Protection Manager. It's a free subscription-based product. You deploy agents from Azure Backup Server to workloads, and the agents then back up to a Recovery Services vault.

Configuring Backup for files and folders involves the following steps:

1. Create the Recovery Services vault. Within your Azure subscription, you must create a Recovery Services vault for the backups.
2. Download the agent and credential file. The Recovery Services vault provides links to download the backup agent. There's also a credentials file that's required during the installation of the agent. You must have the latest version of the agent. Versions of the agent earlier than 2.0.9083.0 must be upgraded by uninstalling and reinstalling the agent.
3. Install and register the agent. The installer provides a wizard to configure the installation location, proxy server, and passphrase information. You use the downloaded credential file to register the agent.
4. Configure the backup. Use the agent to create a backup policy, including when to back up, what to back up, how long to keep items, and settings like network throttling.

Test your knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

Can you use Microsoft Azure Site Recovery to manage virtual machine (VM) replication between two Hyper-V hosts?

Question 2

Is Site Recovery used only as a disaster recovery solution?

Module 07 lab and review

Lab: Implementing Hyper-V Replica and Windows Server Backup

Scenario

You're working as an administrator at Contoso, Ltd. Contoso wants to assess and configure new disaster recovery and backup features and technologies. As the system administrator, you have been tasked with performing that assessment and implementation. You decided to evaluate **Hyper-V Replica** and Windows Server Backup.

Objectives

After completing this lab, you'll be able to:

- Configure and implement **Hyper-V Replica**.
- Configure and implement backup with Windows Server Backup.

Estimated time: 45 minutes

Module review

Use the following questions to check what you've learned in this lesson.

Question 1

How can you monitor virtual machine (VM) replication health by using Windows PowerShell?

Question 2

What's the difference between planned failover and failover?

Question 3

Is Azure Site Recovery used only as a disaster recovery solution?

Question 4

Can you use Azure Backup to back up VMs?

Answers

Question 1

What's the difference between a planned failover and a failover?

You can perform a planned failover when both Hyper-V hosts—at the primary site and at the recovery site—are available. A planned failover is performed without any data loss. When this isn't possible, for example if the primary site is no longer available because of a disaster, you can perform failover, which means unplanned failover. After failover, you'll be able to use a replicated virtual machine (VM), but changes that were performed at the primary site and weren't yet replicated will be lost.

Question 2

Can you use Hyper-V Replica to replicate only VMs that have integration services installed?

No. You can use Hyper-V Replica to replicate any VM regardless of whether it has integration services installed. However, some features such as Failover TCP/IP settings are applied to a replicated VM only if it has integration services installed.

Question 1

Can you use Microsoft Azure Site Recovery to manage virtual machine (VM) replication between two Hyper-V hosts?

No. You can't use Site Recovery to manage replication between two Hyper-V hosts. You can use Site Recovery to manage VM replication from a Hyper-V host to Azure or between two clouds that Microsoft System Center Virtual Machine Manager manages. If you want to manage VM replication between two Hyper-V hosts, you should use Hyper-V Manager.

Question 2

Is Site Recovery used only as a disaster recovery solution?

No. Although administrators often use Site Recovery as a disaster recovery solution, you can also use it in several other scenarios, such as migrating workloads to Azure, cloud bursting, DevTest, and analytics and reporting.

Question 1

How can you monitor virtual machine (VM) replication health by using Windows PowerShell?

*At a Windows PowerShell command prompt, you can run the **Get-VMReplication** and **Measure-VMReplication** cmdlets.*

Question 2

What's the difference between planned failover and failover?

You can perform planned failover when both the Hyper-V hosts at the primary site and the recovery site are available and planned failover is performed without any data loss. When this isn't possible—for example, if the primary site is no longer available because of a disaster—you can perform failover, which means unplanned failover. After failover, you'll be able to use a replicated VM, but changes at the primary site that weren't yet replicated will be lost.

Question 3

Is Azure Site Recovery used only as a disaster recovery solution?

No. You can use it to manage the failover of VMs and Microsoft System Center Virtual Machine Manager (VMM) clouds, to coordinate and monitor asynchronous replication, to continually monitor service availability, to test the recovery, and to manage virtual network mappings between sites.

Question 4

Can you use Azure Backup to back up VMs?

Yes. It's possible to back up both on-premises and Azure VMs by using Backup.



Module 8 Windows Server security

Credentials and privileged access protection in Windows Server

Lesson overview

The first step in securing Windows Server is ensuring that you've properly configured user accounts. This includes ensuring the accounts have only the privileges needed to perform necessary tasks. It also requires that you protect user account credentials from compromise.

Lesson objectives

After completing this lesson, you will be able to:

- Describe and configure user rights.
- Describe protected users and groups, authentication policies, and authentication-policy silos.
- Describe and configure Windows Defender Credential Guard.
- Describe NTLM blocking.
- Locate problematic accounts.

Configure user rights

When configuring user rights, it's important to follow the principle of least privilege. This means granting users only the rights and privileges they need to perform their tasks, and no more. As a result, if an unauthorized user compromises an account, they gain access only to the limited set of privileges assigned to that account. IT staff should also have separate accounts for day-to-day activities such as answering email, separate from the privileged accounts used to perform administrative tasks.

Additional reading: For more information about implementing the principle of least privilege, refer to **Implementing Least-Privilege Administrative Models**¹.

¹ <https://aka.ms/implementing-least-privilege-administrative-models>

You can assign user rights to account in Active Directory (AD DS) or by adding the account to a group to which you have assigned rights. In both cases, rights are assigned using Group Policy. You can review a list of user rights that you can use Group Policy to assign, in the following table.

Table 1: User rights

User rights assignment policy	Function
Access Credential Manager as a trusted caller	Used by Credential Manager during backup and restore. You should not assign this privilege to user accounts.
Access this computer from the network	Determines which users and groups can connect to the computer from the network. This right does not affect Remote Desktop Services.
Act as part of the operating system	Allows a process to impersonate a user without authentication. You typically would assign the LocalSystem account to processes that require this privilege.
Add workstations to a domain	Allows you to join workstations to the domain.
Adjust memory quotas for a process	Determines which security principals can adjust the maximum amount of memory assigned to a process.
Allow sign in locally	Determines which users can sign in locally to a computer. Alter this policy on Privileged Access Workstations to remove members of the Users group as a way of limiting which accounts can sign in to a computer. By default, any authenticated user can sign in to any workstation or server except for a Domain Controller, which is limited to members of certain groups.
Allow sign in through Remote Desktop Services	Determines which users and groups can sign in remotely by using a Remote Desktop Service connection.
Back up files and directories	Gives permission to back up files, directories, registry, and other objects to which the user normally would not have permission. Assigning this right gives indirect access to all data on a computer because the person with that right can back that data up and then recover it in an environment over which they have complete control.
Bypass traverse checking	Allows the user with this right to traverse directories on which they don't have permission. It does not allow the user to list the contents of that directory though.
Change the system time	Allows the user with this right to alter the system time, which is separate from the time zone.
Change the time zone	Allows the user with this right to alter the time zone, but not the system time.
Create a page file	Allows the user with this right to create and modify a page file.

User rights assignment policy	Function
Create a token object	Determines which user accounts that processes can use to create tokens that allow access to local resources. You should not assign this right to any user you don't want to have complete system control, because they can use it to leverage local Administrator privileges.
Create global objects	Determines which user accounts can create global objects that are available to all sessions. You should not assign this right to any user you don't want to give complete system control, because they can use it to leverage local Administrator privileges.
Create permanent shared objects	Determines which user accounts can create directory objects by using the object manager.
Create symbolic links	Determines which user accounts can create symbolic links from the computer they are signed in to. You should assign this right only to trusted users because symbolic links can expose security vulnerabilities in apps that aren't configured to support them.
Debug programs	Determines which user accounts can attach a debugger to processes within the operating system kernel. Only developers who are writing new system components require this ability. Developers who are writing applications do not.
Deny access to this computer from the network	Blocks specified users and groups from accessing the computer from the network. This setting overrides the policy that allows access from the network.
Deny sign in as a batch job	Blocks specified users and groups from signing in as a batch job. This overrides the sign in as a batch job policy.
Deny sign in as a service	Blocks service accounts from registering a process as a service. This policy overrides the sign in as a service policy. However, it doesn't apply to Local System, Local Service, or Network Service accounts.
Deny sign in locally	Blocks accounts from signing on locally. This policy overrides the allow sign in locally policy.
Deny sign in through Remote Desktop Services	Blocks accounts from signing in by using Remote Desktop Services. This policy overrides the Allow sign in through Remote Desktop Services policy.
Enable computer and user accounts to be trusted for delegation	Determines whether you can configure the Trusted for Delegation setting on a user or a computer object.
Force shutdown from a remote system	Users assigned this right can shut down computers from remote network locations.

User rights assignment policy	Function
Generate security audits	Determines which accounts processes can use to add items to the security log. Because this right allows interaction with the security log, it presents a security risk when you assign this to a user account.
Impersonate a client after authentication	Allows apps that are running on behalf of a user to impersonate a client. This right can be a security risk, and you should assign it only to trusted users.
Increase a process working set	Accounts assigned this right can increase or decrease the number of memory pages available for the process to use to the process in random access memory (RAM).
Increase scheduling priority	Accounts assigned this right can change the scheduling priority of a process.
Load and unload device drivers	Accounts assigned this right can dynamically load and unload device drivers into kernel mode. This right is separate from the right to load and unload plug and play drivers. Assigning this right is a security risk because it grants access to the kernel mode.
Lock pages in memory	Accounts assigned this right can use a process to keep data stored in physical memory, blocking that data from paging to virtual memory.
Sign in as a batch job	Users with accounts that have this permission can sign in to a computer through a batch-queue facility. This right is only relevant to older versions of the Windows operating system, and you should not use it with newer versions, such as Windows 10 and Windows Server 2016 or later.
Sign in as a service	Allows a security principal to sign in as a service. You need to assign this right when any service that you configure to use a user account, rather than one of the built-in service accounts.
Manage auditing and security log	Users assigned this right can configure object access auditing options for resources such as files and AD DS (Active Directory) objects. Users assigned this right can also review events in the security log and clear the security log. Because unauthorized users are likely to clear the security log as a way of hiding their tracks, you should not assign this right to user accounts to which you would not assign local Administrator permissions on a computer.
Modify an object label	Users with this permission can modify the integrity level of objects, including files, registry keys, or processes that other users own.

User rights assignment policy	Function
Modify firmware environment values	Determines which users can modify firmware environment variables. This policy is primarily for modifying the boot-configuration settings of non-x86-based computers
Perform volume maintenance tasks	Determines which user accounts can perform maintenance tasks on a volume. Assigning this right is a security risk, because users who have this permission might access data stored on the volume.
Profile single process	Determines which user accounts can leverage performance-monitoring tools to monitor nonsystem processes.
Profile system performance	Determines which user accounts can leverage performance-monitoring tools to monitor system processes.
Remove computer from docking station	When assigned, a user account can remove a portable computer from a docking station without signing in.
Replace a process-level token	When assigned, a user account can call the CreateProcessAsUser API so that one service can trigger another.
Restore files and directories	Allows users assigned this right to bypass permissions on files, directories, and the registry and overwrite these objects with restored data. This right is a security risk, as a user account with this right can overwrite registry settings and replace existing permissions.
Shut down the system	Assigns the ability for a locally signed-in user to shut down the operating system.
Synchronize directory service data	Assigns the ability to synchronize AD DS data.
Take ownership of files or other objects	When assigned, this user account can take ownership of any securable object, including AD DS objects, files, folders, registry keys, processes, and threads. This represents a security risk because it allows the user to take control of any securable object.

You can also configure additional account security options that limit how and when an account can be used, including:

- **Logon Hours.** Use this setting to configure when users can use an account.
- **Logon Workstations.** Use this setting to limit the computers an account can sign in to. By default, users can use an account to sign in to any computer in the domain.
- **Password Never Expires.** You should never configure this option for privileged accounts because it will exempt the account from the domain password policy.
- **Smart card is required for interactive logon.** In high-security environments, you can enable this option to ensure that only an authorized person that has both the smart card and the account credentials can use the privileged account.

- **Account is sensitive and cannot be delegated.** When you enable this option, you ensure that trusted applications cannot forward an account's credentials to other services or computers on the network. You should enable this setting for highly privileged accounts.
- **Use only Kerberos Data Encryption Standard (DES) encryption types for this account.** This option configures an account to use only DES encryption, which is a weaker form of encryption than Advanced Encryption Standard (AES). You should not configure this option on a secure network.
- **This account supports Kerberos AES 128-bit encryption.** When you enable this option, you are allowing Kerberos AES 128-bit encryption to occur.
- **This account supports Kerberos AES 256-bit encryption.** When possible, you should configure this option for privileged accounts and have them use this form of Kerberos encryption over the AES 128-bit encryption option.
- **Do not require Kerberos preauthentication.** Kerberos preauthentication reduces the risk of replay attacks. Therefore, you should not enable this option.
- **Account expires.** Allows you to configure an end date for an account so that it doesn't remain in AD DS after it is no longer used.

Protected Users group, authentication policies, and authentication-policy silos

The AD DS (Active Directory) security group Protected Users helps you protect highly privileged user accounts against compromise. The Protected Users group members have several security-related configuration settings applied that cannot be modified except by leaving the group.

Protected Users group prerequisites

To provide protection for members of the Protected Users group:

- The group must be replicated to all domain controllers.
- The user must sign in to a device running Windows 8.1 or Windows Server 2012 R2 or later.
- Domain controller protection requires that domains must be running at a Windows Server 2012 R2 or higher domain functional level. Lower functional levels still support protection on client devices.

Protected User group protections

When a user is a member of the Protected Users group, on their workstation or local device:

- User credentials are not cached locally.
- Credential delegation (CredSSP) will not cache user credentials
- Windows Digest will not cache user credentials.
- NTLM will not cache user credentials.
- Kerberos will not create DES (Data Encryption Standard) or RC4 keys, or cache credentials or long-term keys.
- The user can no longer sign-in offline.

On domain controllers running Windows Server 2012 R2 or later:

- NTLM authentication is not allowed.

- DES and RC4 encryption in Kerberos preauthentication cannot be used.
- Credentials cannot be delegated using constrained delegation.
- Cannot be delegated using unconstrained delegation.
- Ticket-granting tickets (TGTs) cannot renew past the initial lifetime.

Authentication policies

Authentication policies enable you to configure TGT lifetime and access-control conditions for a user, service, or computer account. For user accounts, you can configure the user's TGT lifetime, up to the maximum set by the Protected Users group's 600-minute maximum lifetime. You can also restrict which devices the user can sign in to, and the criteria that the devices need to meet.

Authentication policy silos

Authentication policy silos allow administrators to assign authentication policies to user, computer, and service accounts. Authentication policy silos work with the Protected User group to add configurable restrictions to the group's existing non-configurable restrictions. In addition, policy silos ensure that the accounts belong to only a single authentication policy silo.

When an account signs in, a user that is part of an Authentication policy silo is granted an Authentication Policy Silo claim. This silo claim controls access to claims-aware resources to verify whether the account is authorized to access that device. For example, you might associate accounts that can access particularly sensitive servers with a specific Authentication policy silo.

Additional reading: For more information about authentication policies and authentication policy silos, refer to **Authentication Policies and Authentication Policy Silos**².

What is Windows Defender Credential Guard?

Windows Defender Credential Guard

Windows Defender Credential Guard helps protect you against NTLM Pass-the-Hash attacks or Kerberos Pass-the-Ticket attacks. It protects you by restricting access to NTLM password hashes (*Pass-the-Hash*), Kerberos TGTs (*Pass-the-Ticket*), and application credentials stored as domain credentials to special processes and memory that manage and store that authorization and authentication-related data. Therefore, only specific, digitally authorized elements of the host operating system—which verifies those elements—can access the special processes and memory. This blocks unauthorized operations or unauthorized software from gaining access to the protected processes and memory, and subsequently limiting their access to authorization and authentication-related data. Windows Defender Credential Guard also provides hardware security by utilizing hardware security features such as secure boot and virtualization for NTLM, Kerberos, and Credential Manager.

How Windows Defender Credential Guard works

Windows Defender Credential Guard protects user credentials from compromise by isolating those credentials within a protected, virtualized container, separate from the rest of the operating system. Only privileged system software can access the credentials.

² <https://aka.ms/authentication-policies-and-policy-silos>

The virtualized container's operating system runs in parallel with, but independent from the host operating system. This operating system protects these processes from attempts by any external entity to read information that those processes store and use. This means that credentials are more protected, even if malware has penetrated the rest of your system.

Windows Defender Credential Guard requirements

Requirements

You can deploy Windows Defender Credential Guard only on devices that meet certain hardware requirements. Windows Defender Credential Guard should be used on any computer where IT staff use privileged credentials, especially workstations dedicated to privileged access.

Windows Defender Credential Guard requires the following:

- Windows 10 Enterprise or Windows Server 2016 or later
- 64-bit CPU
- CPU virtualization extensions plus extended page tables (Intel VT-x or AMD-V)
- Trusted Platform Module (TPM) 1.2 or 2.0
- Unified Extensible Firmware Interface (UEFI) firmware version 2.3.1.c or newer
- UEFI Secure boot
- UEFI secure firmware update

Windows Defender Credential Guard can protect secrets in a Microsoft Hyper-V virtual machine when:

- The Hyper-V host has an input/output memory management unit (IOMMU) and runs Windows Server 2016 or later, or runs Windows 10 Enterprise.
- The virtual machine must be Generation 2, have virtual TPM enabled, and run an operating system that supports Windows Defender Credential Guard.

Windows Defender Credential Guard does not support:

- Unconstrained Kerberos delegation
- NTLMv1
- MS-CHAPv2
- Digest authentication
- Credential (CredSSP) delegation
- Kerberos DES (Data Encryption Standard) encryption

Windows Defender Credential Guard is not supported on domain controllers. It also does not provide protections for the AD DS (Active Directory) database or Security Accounts Manager (SAM).

Additional credential security

Windows 10 certified devices available after mid-2016 support Windows Defender Credential Guard. There are also devices available that include additional hardware and firmware features that can provide additional protections for your administrative workstations.

Additional information: For more information, refer to **Security considerations**³.

Configure Windows Defender Credential Guard

Enabling Windows Defender Credential Guard

You can add support for and enable Windows Defender Credential Guard by:

- Using Group Policy.
- Updating the registry.
- Using the hypervisor-protected code integrity and Windows Defender Credential Guard hardware readiness tool (DG_Readiness_Tool.ps1)

To enable Windows Defender Credential Guard by using Group Policy:

1. Open the Group Policy Management Console and navigate to **Computer Configuration>Administrative Templates>System>Device Guard**.
2. Select or access the context menu for **Turn On Virtualization Based Security**, then select the spacebar and Enter, and then select **Enabled**.
3. In the **Select Platform Security Level** drop-down list, select either **Secure Boot** or **Secure Boot and DMA Protection**.
4. In the **Credential Guard Configuration** drop-down list, select either **Enabled with UEFI lock** or **Enabled without lock**.
5. In the **Secure Launch Configuration** drop-down list, select either **Not Configured**, **Enabled**, or **Disabled**.
6. Select **OK**, and then close the Group Policy Management Console.

The Secure Boot with Direct Memory Access (DMA) Protection option ensures that your system uses Windows Defender Credential Guard with direct memory access protection. If you set the **Enabled with UEFI lock**, Windows Defender Credential Guard cannot be disabled remotely. It can be disabled only if someone with local Administrator privileges signs in locally and disables Windows Defender Credential Guard configuration.

For more information about the **Secure Launch Configuration** options, refer to **System Guard Secure Launch and SMM protection**⁴.

The simplest option for enabling Windows Defender Credential Guard is by using the Hypervisor-Protected Code Integrity and Windows Defender Credential Guard hardware readiness tool.

To enable windows Defender Credential using the tool:

1. Download the tool into the script and extract the files to install the tool.
2. In Windows PowerShell, navigate to the **Installation** folder and run the following command:

```
DG_Readiness_Tool.ps1 -Enable -AutoReboot
```

You can also use the tool to verify that Windows Defender Credential guard is enabled by running the tool with the **-Ready** parameter.

³ <https://aka.ms/security-considerations>

⁴ <https://aka.ms/system-guard-secure-launch-and-smm-protection>

Disabling Windows Defender Credential Guard

If Credential Guard was enabled without UEFI Lock and you used Group Policy to enable Windows Defender Credential Guard, you can disable Windows Defender Credential Guard by disabling the Group Policy setting. Otherwise, you must complete following additional steps:

1. Delete the following registry settings:
 - **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\LSA\LsaCfgFlags**
 - **HKEY_LOCAL_MACHINE\Software\Policies\Microsoft\Windows\DeviceGuard\LsaCfgFlags**
2. If UEFI Lock was enabled, you must also delete the Windows Defender Credential guard EFI variables using **bcdedit**. From an elevated command prompt, enter the following commands:

```
mountvol X: /s
copy %WINDIR%\System32\SecConfig.efi X:\EFI\Microsoft\Boot\SecConfig.efi /Y
bcdedit /create {0cb3b571-2f2e-4343-a879-d86a476d7215} /d "DebugTool" /application osloader
bcdedit /set {0cb3b571-2f2e-4343-a879-d86a476d7215} path "\EFI\Microsoft\Boot\SecConfig.efi"
bcdedit /set {bootmgr} bootsequence {0cb3b571-2f2e-4343-a879-d86a476d7215}
bcdedit /set {0cb3b571-2f2e-4343-a879-d86a476d7215} loadoptions DISABLE-LSA-ISO
bcdedit /set {0cb3b571-2f2e-4343-a879-d86a476d7215} device partition=X:
mountvol X: /d
```

3. Restart the device.
4. Accept the prompt to disable Windows Defender Credential Guard.

Note: You must be at the physical machine to accept this prompt.

You can also disable Windows Defender Credential Guard by using the Hypervisor-Protected Code Integrity and Windows Defender Credential Guard hardware readiness tool. In a Windows PowerShell administrator session, run the following command:

```
DG_Readiness_Tool_v3.6.ps1 -Disable -AutoReboot
```

A final option for disabling Windows Defender Credential Guard is by disabling Hyper-V.

NTLM blocking

The NTLM authentication protocol is less secure than the Kerberos authentication protocol. You should block the use of NTLM for authentication and use Kerberos instead.

Audit NTLM traffic

Prior to blocking NTLM, you need to ensure that existing applications are no longer using the protocol. You can audit NTLM traffic by configuring the following Group Policy settings under Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options:

- **Network security: Restrict NTLM: Outgoing NTLM Traffic to remote servers.** Configure this policy with the **Audit All** setting.
- **Network security: Restrict NTLM: Audit Incoming NTLM Traffic.** Configure this policy with the **Enable auditing for all accounts** setting.

- **Network security: Restrict NTLM: Audit NTLM authentication in this domain.** Configure this policy with the **Enable for domain accounts to domain servers** setting on domain controllers. You should not configure this policy on all computers.

Block NTLM

After you have determined that you can block NTLM in your organization, you need to configure the Restrict NTLM: NTLM authentication in this domain policy in the previous Group Policy node. The configuration options are:

- **Deny for domain accounts to domain servers.** This option denies all NTLM authentication sign-in attempts for all servers in the domain that use domain accounts, unless the server name is listed in the **Network Security: Restrict NTLM: Add server exceptions for NTLM authentication** setting in this domain policy.
- **Deny for domain accounts.** This option denies all NTLM authentication attempts for domain accounts unless the server name is listed in the **Network Security: Restrict NTLM: Add server exceptions** for NTLM authentication in this domain policy.
- **Deny for domain servers.** This option denies NTLM authentication requests to all servers in the domain unless the server name is listed in the **Network Security: Restrict NTLM: Add server exceptions** setting for NTLM authentication in this domain policy.
- **Deny all.** This option ensures that all NTLM pass-through authentication requests for servers and accounts will be denied unless the server name is listed in the **Network Security: Restrict NTLM: Add server exceptions** setting for NTLM authentication in this domain policy.

Locate problematic accounts

You should check your AD DS environment for accounts that have not signed in for a specific period of time, or that have passwords with no expiration date.

Inactive user accounts usually indicate a person that has left the organization and organization processes have failed to remove or disable the account. The account might also have originally been shared by IT staff, but is no longer in use. These extra accounts represent additional opportunities for unauthorized users to gain access to your network resources.

Accounts with fixed passwords are less secure than accounts that are required to update their password periodically. If a third-party user obtains a user's password, that knowledge is only valid until the user updates the password. If you configure an account with a password that the user doesn't have to update periodically, then a potential cybercriminal could have access to your network indefinitely. Ensuring regular password updates is especially important for highly privileged accounts.

When you find accounts that haven't signed in for a specified number of days, you can disable those accounts. Disabling them allows you to reenable them should the person return. After you've located accounts that are configured with passwords that don't expire, you can take steps to ensure that an appropriate password update policy is enforced.

Note: User accounts with credentials shared by multiple IT staff members should be avoided, even if they have a strong password policy. Shared accounts make it hard to track which individual performed a specific administrative task.

You can use Windows PowerShell or the AD DS Administrative Center to find problematic users. To use Windows PowerShell to find active users with passwords set to never expire, use the following command:

```
Get-ADUser -Filter {Enabled -eq $true -and PasswordNeverExpires -eq $true}
```

Use the following Windows PowerShell command to find users that have not signed in within the last 90 days, using Windows PowerShell:

```
Get-ADUser -Filter {LastLogonTimeStamp -lt (Get-Date).Adddays(-(90)) -and enabled -eq $true} -Properties LastLogonTimeStamp
```

Demonstration: Locate problematic accounts

In this demonstration, you will learn how to locate problematic user accounts.

Demonstration steps

1. Sign in to **SEA-ADM1** as **Contoso\Administrator**.
2. Use the Windows PowerShell cmdlet **Get-ADUser** to find users whose passwords are set to never expire.
3. Review the returned list of users.
4. Use the Windows PowerShell cmdlet **Get-ADUsers** to find users who have not signed in within the last 90 days.
5. Review the returned list of users, if any.

Test your knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

Which security setting should not be enabled when configuring administrative user accounts?

- Logon Hours
- Account is sensitive and cannot be delegated
- This account supports Kerberos AES (Advanced Encryption Standard) 256-bit encryption
- Do not require Kerberos preauthentication

Question 2

Which feature allows you to configure TGT (Ticket-granting tickets) lifetime and access-control conditions for a user?

- Protected Users group
- Authentication policies
- Authentication policy silos
- NTLM blocking

Question 3

Which is not a valid way to enable Windows Defender Credential Guard on a server?

- Group policy
- Adding server role
- Updating the registry
- Using a Windows PowerShell script

Question 4

What are two types of problematic user accounts you should check for regularly?

- Users with passwords that do not expire
- Users that have not signed in recently
- Users with complex passwords
- Users with few administrative permissions

Hardening Windows Server

Lesson overview

Another integral part of securing your Windows Server environment is making sure the servers themselves are secure, or *hardened*. It's also important that any client devices that are used to access those servers are also hardened. In this lesson, we'll review a variety of tools that help you harden your servers and devices.

Lesson objectives

After completing this lesson, you will be able to:

- Use the Local Administrator Password solution to manage local Administrator passwords.
- Explain the need to limit administrative access to secure hosts and know how to manage that access.
- Describe how to secure domain controllers.
- Describe how to use the Microsoft Security Compliance Toolkit to harden servers.

What is Local Administrator Password Solution?

Each computer that is member of a domain keeps a local Administrator account. This is the account that you configure when you first deploy the computer manually, or which is configured automatically when you use software deployment tools such as Microsoft Endpoint Configuration Manager. The local Administrator account allows IT staff to sign in to the computer if they cannot establish connectivity to the domain.

Managing passwords for the local Administrator account for every computer in the organization can be extremely complicated. An organization with 5,000 computers has 5,000 separate local Administrator accounts to manage. What often happens is that organizations assign a single, common local Administrator account password to all local Administrator accounts. The drawback to this approach is that people beyond the IT operations team often figure out this password, and then use it to gain unauthorized local Administrator access to computers in their organization.

Local Administrator Password Solution (LAPS) provides organizations with a central local administrator passwords repository for domain-member machines, and provides several features:

- Local administrator passwords are unique on each computer that LAPS manages.
- LAPS randomizes and changes local administrator passwords regularly.
- LAPS stores local administrator passwords and secrets securely within AD DS (Active Directory).
- Configurable permissions control access to passwords in AD DS.
- Passwords that LAPS retrieves are transmitted to the client in a secure, encrypted manner.

Prerequisites

- LAPS supports all currently supported Windows operating system versions.
- LAPS requires an update to the AD DS schema. You perform this update by running the **Update-AdmPwdADSchema** cmdlet, which is included in a Windows PowerShell module that's made available when you install LAPS on a computer. However, the person running this cmdlet must be a member of

the Schema Admins group, and you should run this cmdlet on a computer that's in the same AD DS site as the computer containing the Schema Master role for the forest.

You configure the LAPS agent through a Group Policy client-side extension. You install the LAPS client using an .msi file on client computers that you will manage.

How LAPS works

The LAPS (Local Administrator Password Solution) process occurs each time Group Policy refreshes. When a Group Policy refresh occurs, the following steps take place:

1. LAPS determines if the password of the local Administrator account has expired.
2. If the password hasn't expired, LAPS does nothing.
3. If the password has expired, LAPS performs the following steps:
 1. Changes the local Administrator password to a new, random value based on the configured parameters for local Administrator passwords.
 2. Transmits the new password to AD DS, which stores it in a special, confidential attribute associated with the computer account of the computer that has had its local Administrator account password updated.
 3. Transmits the new password-expiration date to AD DS, where it's stored in a special, confidential attribute associated with the computer account of the computer that has had its local Administrator account password updated.

Authorized users can read passwords from AD DS, and an authorized user can trigger a local Administrator password change on a specific computer.

Configure and manage passwords using LAPS

There are several steps that you need to take to configure and manage passwords by using LAPS (Local Administrator Password Solution). The first set of steps involve configuring AD DS (Active Directory). First, you move the computer accounts of computers that you want to use LAPS to an OU (Organizational Unit). After you've moved the computer accounts into an OU, you use the **Set-AdmPwdComputerSelfPermission** cmdlet to assign the computers the ability to update their local Administrator account password when it expires.

For example, to allow computers in the Sydney OU with expired passwords to update their passwords by using LAPS, you would use the following command:

```
Set-AdmPwdComputerSelfPermission -Identity "Sydney"
```

By default, accounts that are members of the Domain Admins and Enterprise Admins groups can access and find stored passwords. You use the **Set-AdmPwdReadPasswordPermission** cmdlet to provide additional groups the ability to find the local administrator password.

For example, to assign the Sydney_ITOps group the ability to find the local administrator password on computers in the Sydney OU, you would use the following command:

```
Set-AdmPwdReadPasswordPermission -Identity "Sydney" -AllowedPrincipals  
"Sydney_ITOps"
```

The next step is to run the LAPS installer to install the Group Policy Object (GPO) templates into AD DS. After you have installed the templates, you can configure the following policies:

- **Enable local admin password management.** This policy enables LAPS and enables you to manage the local Administrator account password centrally.
- **Password settings.** This policy allows you to configure the complexity, length, and maximum age of the local Administrator password. The default password requirements are:
 - Uppercase and lowercase letters
 - Numbers
 - Special characters
 - 14-character password length
 - 30 days maximum password age
- **Do not allow password expiration time longer than required.** When enabled, the password updates according to the domain password expiration policy.
- **Name of administrator account to manage.** Use this policy to identify custom local Administrator accounts.

You can study passwords assigned to a computer by using one of the following methods:

- Use **Advanced Features** to study the computer account properties that are enabled in AD DS Users and Computers by examining the:
 - **ms-Mcs-AdmPwd** attribute.
 - LAPS GUI application.
 - **Get-AdmPwdPassword** cmdlet, available from the `AdmPwd.PS` module, which is available when you install LAPS.

Demonstration: Configure and deploy LAPS

Demonstration steps

1. Sign in to **SEA-ADM1** as **Contoso\Administrator**.
2. Use the Windows PowerShell cmdlet **New-ADOrganizationalUnit** to create an OU (Organizational Unit) named **Seattle_Servers**.
3. Use the Windows PowerShell cmdlet **MOve-ADObject** to add **SEA-SVR1** to the new OU.
4. Install the LAPS (Local Administrator Password Solution) tool, including management components, using the installation file located at **C:\Labfiles\Mod08\LAPS.x64.msi**.
5. After installation is complete, in Windows PowerShell, import the `admpwd.ps` module.
6. Update the AD DS (Active Directory) schema using the cmdlet **Update-AdmPwsADSchema**.
7. Use the cmdlet **Set-AdmPwdComputerSelfPermission** to give **Seattle_Servers** OU computers write access to their passwords.
8. Create a Group Policy named **LAPS_GPO** linked to the **Seattle_Servers** OU, which enables LAPS and sets the password age to 30 days and the password length to 20 characters.
9. Switch to **SEA-SVR1** and sign in as **Contoso\Administrator**.

10. Install the LAPS tool without management components using the installation file located at **C:\Labfiles\Mod08\LAPS.x64.msi**.
11. Update Group Policy settings on **SEA-SVR1**.
12. Switch to **SEA-ADM1**.
13. Open the **LAPS UI** and review the **Password** and **Password expires** values.
14. In the Windows PowerShell window, use the **Get-AdmPwsPassword** cmdlet to review the **Password** and the **Password expires** values for **SEA-SVR1**.

Limit administrative access to secure hosts

Another important part of securing an environment is to ensure that the computers that IT staff use to connect to secure servers are themselves secure. Your business-critical servers are only as secure as the computers that your IT staff use to perform administrative tasks.

IT staff should not use computers that they use to perform administrative tasks on servers for daily tasks such as answering email and browsing the internet. Those daily tasks provide another avenue for compromising administrative workstations.

IT staff are high-value targets for unauthorized users. By gaining access to IT computers or accounts they can have all the privileges and resources that the team member has. Without protection solutions such as Windows Defender Credential Guard, in addition to the original compromised account a sophisticated unauthorized user can extract any other account credentials that have been used to sign in to the same computer. This is particularly problematic if the same device is used to sign in with both credentials for daily tasks and for highly privileged credentials.

Even if an unauthorized user can't harvest credentials directly, they could install malware such as a keystroke logger that lets them indirectly discover usernames and passwords used for administrative tasks. You can help make sure external attacks can't infect devices by ensuring IT staff perform administrative tasks only on secure administrative hosts, also called a *privileged access workstation (PAW)*.

Privileged access workstation configuration

When configuring a PAW, you should:

- Ensure that only authorized users can sign in to the PAW. Standard user accounts should not be able to sign in.
- Enable Windows Defender Credential Guard to help protect against credential theft.
- Enable BitLocker Drive Encryption to help protect the boot environment and the hard disk drives from tampering.
- Use Windows Defender Device Guard policies to restrict application execution to only trusted applications that your organization's employees use for performing administrative tasks.
- Block PAWs from accessing the internet.
- Install all the tools your administrative tasks require, so your IT staff are not tempted or required to use other workstations to perform their administrative tasks.
- Limit physical access to the PAWs

After you have PAWs configured, you should then perform the following configuration tasks to maximize their value in securing your environment:

- Block Remote Desktop Protocol (RDP), Windows PowerShell, and management console connections to your servers that come from any computer that is not a PAW.
- Implement connection security rules so that traffic between servers and PAWs is encrypted and protected from replay attacks.
- Configure sign-in restrictions for administrative accounts so that those accounts can only sign in to a PAW.

Combining a daily-user workstation and a PAW on the same device is a common practice. You do this by hosting one of the operating systems in a virtual environment. However, if you do this you should host the daily-use workstation virtual machine within the PAW host, and not a PAW virtual machine within a daily-user host. If the PAW is hosted in the daily user workstation and the workstation is compromised, the PAW could be compromised as well.

For more information about Device Guard, refer to **Windows Defender Application Control and virtualization-based protection of code integrity**⁵.

Jump servers and secure bastion hosts

Jump servers and secure bastion hosts are similar to PAWs. When performing administrative tasks, you use RDP to connect to a dedicated server that's configured similarly to a PAW. The difference is that you sign in to PAWs locally, while you sign in to jump servers remotely. In this scenario, there is a risk that the host you sign in from might be compromised.

Note: You can combine jump servers with PAWs.

Secure domain controllers

Domain controllers represent one of the most valuable targets on a network. This is because an unauthorized user who can compromise a domain controller has control of all domain identities. You can help secure your domain controllers by taking the following precautions:

- Ensure that domain controllers are running the most recent version of the Windows Server operating system and have current security updates.
- Deploy domain controllers by using the **Server Core** installation option rather than the **Desktop Experience** option. Deploying the **Server Core** installation option reduces the domain controller's attack surface and minimizes the chance that someone might install malware inadvertently because they signed in to the domain controller on the same computer they used to navigate to an unsafe website.
- Keep physically deployed domain controllers in dedicated, secure racks that are separate from other servers.
- Deploy domain controllers on hardware that includes a Trusted Platform Module (TPM) chip and configure all volumes with BitLocker Drive Encryption. If you cannot physically isolate and secure domain controllers at a branch-office location, you should configure them as read-only domain controllers (RODC).
- Run virtualized domain controllers either on separate virtualization hosts or as a shielded virtual machine on a guarded fabric, which helps protect the domain controller.

⁵ <https://aka.ms/device-guard-virtualization-based-security-and-windows-defender-application-control>

- Review Center for Internet Security (CIS) benchmark for Windows Server operating systems, for security guidance specific to domain controllers.
- Use Windows Defender Device Guard to control the execution of scripts and executables on the domain controller. This minimizes the chance that unauthorized executables and scripts can run on the computer.
- Configure RDP (Remote Desktop Protocol) through Group Policy assigned to the Domain Controllers' OU to limit RDP connections so that they can occur only from jump servers and privileged access workstations.
- Configure the perimeter firewall to block outbound connections to the internet from domain controllers. If an update management solution is in place, it might also be prudent to block domain controllers from communicating with hosts on the internet entirely.

Additional reading: For more information, refer to **Securing Domain Controllers Against Attack**⁶.

Additional reading: For more information about the CIS benchmarks for Windows Server, refer to **CIS Microsoft Windows Server Benchmarks**⁷.

Overview of the Security Compliance Toolkit

The Microsoft Security Compliance Toolkit (SCT) is a set of tools provided by Microsoft that you can use to download and implement security configuration baselines, typically referred to as *simply security baselines*, for Windows Server and other Microsoft products, including Windows 10, Microsoft 365 Apps for enterprise, and Microsoft Edge. You implement the security configuration baselines by using the toolkit to manage your Group Policy Objects (GPOs).

You can also use the SCT to compare your current GPOs to the recommended GPO security baselines. You can then edit the recommended GPOs and apply them to devices in your organization.

In addition to security baselines for Windows Server, the SCT also includes the Policy Analyzer and Local Group Policy Object (LGPO) tools, which also help you manage your GPO settings.

Policy Analyzer tool

The Policy Analyzer tool helps you to compare sets of GPOs. The Policy Analyzer:

- Highlights redundant or inconsistent settings.
- Highlights differences between sets of GPOs.
- Compares GPOs to local policy and registry settings.
- Exports results to Microsoft Excel.

Additional reading: For more information about the Policy Analyzer tool, refer to **New tool: Policy Analyzer**⁸.

LGPO tool

The LGPO tool (LGPO.exe) helps you verify the effects of GPO settings on a local host. You can also use it to help manage systems that are not domain joined. The LGPO tool can export and import Registry Policy

⁶ <https://aka.ms/securing-domain-controllers-against-attack>

⁷ <https://aka.ms/Securing-microsoft-windows-server>

⁸ <https://aka.ms/new-tool-policy-analyzer>

settings files, security templates, Advanced Auditing backup files, and from LGPO files, text files with a special formatting.

Additional reading: For more information about the LGPO tool, refer to **LGPO.exe - Local Group Policy Object Utility, v1.0**⁹.

CIS-level hardening

The CIS (Center for Internet Security) issues benchmarks for hardening Windows Server and other operating systems. These benchmarks are considered by many to be the industry standard for OS security configuration. The Windows Server security baselines provided by Microsoft align closely to CIS Level 1 hardening for member servers.

Additional reading: For more information about the CIS benchmarks for Windows Server, refer to “CIS Microsoft Windows Server Benchmarks” at **Securing Microsoft Windows Server**¹⁰.

Test your knowledge

Use the following questions to check what you’ve learned in this lesson.

Question 1

Which of these is a capability of LAPS (Local Administrator Password Solution)?

- Verify the local administrator password is the same on all managed servers.
- Store local administrator passwords in Microsoft Exchange.
- Prevent local administrator passwords from expiring.
- Ensure that local administrator passwords are unique on each managed server.

Question 2

When configuring a PAW (Privileged Access Workstation), which of these should you not do?

- Ensure that only authorized users can sign in to the PAW. Standard user accounts should not be able to sign in.
- Enable Windows Defender Credential Guard to help protect against credential theft.
- Ensure the PAW can access the internet.
- Limit physical access to the PAW.

⁹ <https://aka.ms/lgpo.exe-local-group-policy-object-utility-v1-0>

¹⁰ <https://aka.ms/Securing-microsoft-windows-server>

Question 3

Which options are valid ways to secure a domain controller? Select all that apply.

- Ensure that domain controllers run the most recent version of the Windows Server operating system and have current security updates.
- Deploy domain controllers by using the "Server Core" installation option.
- Configure RDP (Remote Desktop Protocol) through Group Policy to limit RDP connections to domain controllers, so they can occur only from PAWs.
- Configure the perimeter firewall to block outbound connections to the internet from domain controllers.

Question 4

What CIS hardening level maps to the security configuration baselines included in the SCT (Microsoft Security Compliance Toolkit)?

- Level 0
- Level 1
- Level 2
- None

Just Enough Administration in Windows Server

Lesson overview

Just Enough Administration (JEA) is an administrative technology that allows you to apply role-based access control (RBAC) and the least privilege principles to Windows PowerShell remote sessions. Instead of assigning users broad roles that enable them to perform tasks that are not directly related to a specific work requirement, JEA allows you to configure special Windows PowerShell endpoints that provide only the functionality necessary to perform a specific task.

Lesson objectives

After completing this lesson, you will be able to:

- Describe JEA.
- Explain the limitations of JEA.
- Describe role capabilities files and their use in JEA.
- Describe sessions configuration files and their use in JEA.
- Register JEA endpoints.
- Connect to JEA endpoints.

What is JEA?

JEA (Just Enough Administration) provides Windows Server and Windows client operating systems with RBAC functionality built on Windows PowerShell remoting. *Windows PowerShell remoting* is when a Windows PowerShell remote session is initiated on one computer and the activities are performed on another computer.

When you configure JEA, an authorized user connects to a specially configured endpoint and uses a specific set of Windows PowerShell cmdlets, parameters, and parameter values. You can also configure a JEA endpoint to allow certain scripts and commands to be run, providing these run from within a Windows PowerShell session. For example, you can configure a JEA endpoint to allow an authorized user to restart specific services, such as the Domain Name System (DNS) service, but not restart any other service or perform any other tasks on the system on which the endpoint is configured. You can also configure JEA endpoint to enable an authorized user to run a command such as **whoami.exe** to determine which account is being used with the session.

When connected to the endpoint, JEA uses a special, privileged, virtual account rather than the user's account to perform tasks. The advantages of this approach include:

- The user's credentials are not stored on the remote system. If the remote system is compromised, the user's credentials are not subject to credential theft and cannot be used to traverse the network and gain access to other hosts.
- The user account that's used to connect to the endpoint doesn't need to be privileged. The endpoint simply needs to be configured to allow connections from specified user accounts.
- The virtual account is limited to the system on which it is hosted. The virtual account cannot be used to connect to remote systems. This means that unauthorized users cannot use a compromised virtual account to access other protected servers.

- The virtual account has local administrator privileges but is limited to performing only the activities defined by JEA. You can configure the virtual account with membership of a group other than the local Administrators group to further reduce privileges.

JEA works on the following operating systems directly:

- Windows Server 2016 or later
- Windows 10

JEA limitations

Configuring JEA (Just Enough Administration) can be a complicated process. The person who's configuring the capabilities for JEA roles must understand precisely which cmdlets, parameters, aliases, and values are needed to perform administrative tasks. Because of this, JEA is suitable for routine configuration tasks such as restarting a service or deploying a container or virtual machine (VM).

JEA is not suitable for tasks where the problem and solution are not clearly defined, and therefore you don't know which tools you might need to solve the problem. If you don't know which tools are needed, you can't configure JEA with the necessary tools.

Also, JEA only works with Windows PowerShell sessions. While you can configure scripts and executable commands to be available in a JEA session, JEA requires that administrative tasks be performed from the Windows PowerShell command line. This will be challenging to staff who primarily use graphical user interface (GUI) tools.

Role capability files

What are role capability files?

Role capability files help you specify what can be done in a Windows PowerShell session. Anything that's not explicitly allowed in a role capability file or a session configuration file is not allowed.

You can create a new, blank, role capability file by using the **New-PSRoleCapabilityFile** cmdlet. (Role capability files use the .psrc extension.) You then edit the role capability file, adding cmdlets, functions, and external commands as necessary. You can allow entire Windows PowerShell cmdlets or functions, or you can list which parameters and parameter values can be used.

When you create a role capability file, you can define the following limitations for the Windows PowerShell session:

- **VisibleAliases.** This setting lists which aliases to make available in the JEA (Just Enough Administration) session.
- **VisibleCmdlets.** This setting lists which Windows PowerShell cmdlets are available in the session. You can choose either to list cmdlets, allowing all parameters and parameter values to be used, or limit cmdlets to particular parameters and parameter values.
- **VisibleFunctions.** This setting lists which Windows PowerShell functions are available in the session. Again, you can choose to list functions, allowing all parameters and parameter values to be used, or you can limit functions to particular parameters and parameter values.
- **VisibleExternalCommands.** This setting allows users who are connected to the session to run external commands. For example, you can use this field to allow access to **c:\windows\system32\whoami.exe** so that users connected to the JEA session can identify their security context.
- **VisibleProviders.** This setting lists Windows PowerShell providers that are visible to the session.

You can also configure other settings such as which modules to import, which assemblies are loaded, and data types that are available. For a list of all the options when creating a role capabilities file, refer to **New-PSRoleCapabilityFile**¹¹.

Which commands should you allow?

It's important to properly configure your role capability files. If you give users too few tools, they can't get their work done. If you give them too many tools, then you increase the attack surface of the Windows PowerShell session.

Using the following process can help you decide how to configure your role capabilities files:

1. Working with your IT team, survey your current tools and processes to identify which commands are needed.
2. Whenever possible, move away from command-line tools to PowerShell cmdlets.
3. Restrict which cmdlet parameters and values can be used to only those that are necessary to complete specific tasks.
4. Prevent the use of commands that let users elevate their permissions or that allow them to run arbitrary code.
5. Create custom functions with validation logic to replace complex commands.
6. Test and monitor the list of allowed commands over time and modify as needed.

Session configuration files

Session configuration files are used to register a JEA (Just Enough Administration) endpoint. The session configuration file is responsible for naming the JEA endpoint. It also controls:

- Who can access the JEA endpoint.
- What roles the user is assigned.
- Which identity is used by JEAs virtual account.

Session configuration files use the .pssc file extension, and you create new session configuration files by using the **New-PSSessionConfigurationFile** cmdlet.

You can also use session configuration files to define which cmdlets and functions are available in a JEA session, just like you can with role capabilities files. In addition, you can configure the following settings unique to session configure files:

- **Sessiontype**. This setting allows you to configure the session's default settings. The **Sessiontype** of **RestrictedRemoteServer** is used for sessions used by JEA for secure management.
- **RoleDefinitions**. This setting is used to assign role capabilities to specific security groups.
- **RunAsVirtualAccount**. This setting allows JEA to use the privileged virtual account created just for the JEA session. This account is a member of the local Administrators group and the Domain Admins group on domain controllers.
- **TranscriptDirectory**. This setting specifies where JEA activity transcripts are stored.
- **RunAsVirtualAccountGroup**. This setting allows you to specify the groups that the virtual account is a member of, instead of the default Administrators or Domain Admins groups.

¹¹ <https://aka.ms/new-psrolecapabilityfile>

Additional reading: For a list of all the options when creating session configuration files, refer to [New-PSSessionConfigurationFile¹²](#).

JEA endpoints

A *JEA (Just Enough Administration) endpoint* is a Windows PowerShell endpoint that is configured so only specific authenticated users can connect to it. Once connected, those users only have access predefined sets of Windows PowerShell cmdlets, parameters, and values, based on security group and role capability definitions.

Servers can have multiple JEA endpoints. Each JEA endpoint should be configured so it's used for a specific administrative task. For example, you might have a Domain Name System Operations (DNSOps) endpoint to perform DNS administrative tasks, and a DHCPOps endpoint to perform Dynamic Host Configuration Protocol (DHCP) administrative tasks.

With the JEA endpoints, your IT staff doesn't need to have privileged accounts that are members of groups such as the local Administrators group, to connect to an endpoint. Instead, users have the privileges assigned to the virtual account, which is configured in the session configuration file and could include the privileges of a local administrator or Domain Admin.

Registering JEA on a single machine

On a single computer, you can create JEA endpoints by using the **Register-PSSessionConfiguration** cmdlet. When using this cmdlet, you specify an endpoint name and a session configuration file located on the local machine. However, prior to creating the JEA endpoint you must ensure that the following prerequisites are met:

- You must have defined one or more roles, and the role capabilities file (or files) must be placed in the **RoleCapabilities** folder of a Windows PowerShell module.
- You have created a session configuration file.
- The user registering JEA must be an administrator on the machine.
- You have decided on a name for the JEA endpoint

Windows Server ships with some predefined JEA endpoints, which have a name starting with Microsoft. You can find existing JEA endpoints using the following Windows PowerShell command:

```
Get-PSSessionConfiguration | Select-Object Name
```

For example, to register the endpoint DNSOps using the DNSOps.pssc session configuration file, use the following command:

```
Register-PSSessionConfiguration -Name DNSOps -Path .\DNSOps.pssc
```

Registering JEA on multiple machines

You can register JEA on multiple machines by using Desired State Configuration (DSC). To use DSC to deploy JEA, the following prerequisites must be met:

- You must have defined one or more roles, and the role capabilities file (or files) must be placed in the **RoleCapabilities** folder of a Windows PowerShell module.

¹² <https://aka.ms/new-pssessionconfigurationfile>

- The PowerShell module must be stored on a read-only file share accessible by the machines.
- You have determined the session configuration settings. (You don't need to create a session configuration file though.)
- You have account credentials that have administrative access to each machine.
- You have downloaded the JEA DSC resource from **JEA/DSC Resource at master PowerShell/JEA**¹³
- You have decided on a name for the JEA endpoint

You can apply the DSC configuration using the Local Configuration Manager or by updating the pull server configuration.

Additional reading: For more information about Registering JEA on multiple machines, refer to the GitHub page **JEA/DSC Resource**/¹⁴.

Connect to a JEA endpoint

You connect to a JEA (Just Enough Administration) endpoint by connecting interactively, using implicit remoting, programmatically, or through PowerShell Direct.

Interactive JEA connections

You can use JEA the same way you would connect with a regular PowerShell remoting session. To use JEA interactively, you need:

- The remote computer name.
- The JEA endpoint name.
- An account with access to the desired endpoint.

For example, if you have access to the JEA endpoint named DNSOps on the local server, you can connect to the JEA endpoint using the following PowerShell command:

```
Enter-PSSession -<ComputerName> localhost -ConfigurationName DNSOps
```

After you are connected, your command prompt will change to `[localhost]: PS>`. If you're not sure what commands are available, you can use the **Get-Command** cmdlet to review which ones are available.

One limitation of interactive JEA sessions is that they operate in `NoLanguage` mode. This means you can't use variables to store data. For example, the following commands to start a virtual machine will not work because of the user of variables:

```
$myvm = Get-VM -Name 'MyVM'
Start-VM -vm $myvm
```

However, you can use piping to direct output of one command to another. This means that the following command would be the equivalent of the previous commands:

```
Get-VM -Name 'MyVM' | Start-VM
```

¹³ <https://aka.ms/PowerShell/JEA/tree/master/DSC>

¹⁴ <https://aka.ms/multi-machine-configuration-with-dsc>

Implicit remoting and JEA

Implicit remoting lets you import proxy versions of cmdlets from a remote machine to your local Windows PowerShell environment. This lets you use Windows PowerShell features such as tab completion, variables, or even local scripts.

You can even prefix PowerShell commands with a unique string so you can differentiate between the remote commands and local ones. For example, you could use the following commands to import the DNSOps JEA session and prefix the commands with DNSOps:

```
$DNSOpssession = New-PSSession -ComputerName 'MyServer' -ConfigurationName  
'DNSOps'  
Import-PSSession -Session $DNSOpssession -Prefix 'DNSOps'  
Get-DNSOpsCommand
```

Programmatic access to JEA

You can connect to JEA endpoints programmatically the same way you connect to other PowerShell endpoints programmatically.

For more information about connecting to JEA endpoints programmatically, refer to **Using JEA programmatically**¹⁵

JEA and PowerShell Direct

PowerShell Direct allows Hyper-V administrators to connect to VMs from the Hyper-V host. By doing this, they can ignore any network or remote management settings on the VM.

The Hyper-V administrator connects to the VM the same way they would connect to any other server using PSRemoting, only specifying the **-VMName** or **-VMId** parameter. Whenever using JEA to manage VMs, you should create a dedicated JEA user account for the Hyper-V administrator, and the account's ability to sign-in locally to the VM.

Demonstration: Connect to a JEA endpoint

In this demonstration you will create a JEA (Just Enough Administration) endpoint and connect to it.

Demonstration steps

1. Sign in to **SEA-ADM1** as **Contoso\Administrator**.
2. In Windows PowerShell, use the **New-ADgroup** cmdlet to create a security group named **DNSOps**, and add **Contoso\Administrator** to it.
3. Open the Windows Admin Center.
4. Connect to **SEA-SVR1**.
5. Use remote PowerShell connection to **SEA-SVR1** to create the directory **c:\Program Files\WindowsPowerShell\ModulesPowerShell\DNSOps**.
6. In the **DNSOps** folder, create a module manifest named **DNSOps.psd1**.
7. In the **DNSOps** folder, create a folder named **RoleCapabilities**.

¹⁵ <https://aka.ms/using-jea-programmatically>

8. In the **RoleCapabilities** folder, create a Role Capabilities file named **DNSOps.psrc**.
9. From **SEA-ADM1**, use Notepad to edit the file **DNSOps.psrc** located on **SEA-SRV1**.
10. Replace the line that starts with **# VisibleCmdlets =** with the following text:

```
VisibleCmdlets = @{ Name = 'Restart-Service'; Parameters = @{ Name='Name'; ValidateSet = 'DNS'}}
```
11. Replace the line that starts with **# VisibleFunctions =** with the following text:

```
VisibleFunctions = 'Add-DNSServerResourceRecord', 'Clear-DNSServerCache','Get-DNSServerResourceRecord','Remove-DNSServerResourceRecord'
```
12. Replace the line that starts with **# VisibleExternalCommands =** with the following text:

```
VisibleExternalCommands = 'C:\Windows\System32\whoami.exe'
```
13. Save the file and close Notepad.
14. In the **SEA-SRV1** Windows Admin Center's PowerShell window, enter the following command:

```
New-PSSessionConfigurationFile -Path .\DNSOps.pssc -Full
```
15. From **SEA-ADM1**, use Notepad to edit the file **DNSOps.pssc** located on **SEA-SRV1**.
16. Replace the line that starts with ****SessionType = 'Default'** with the following text:

```
SessionType = 'RestrictedRemoteServer'
```
17. Replace the line that starts with **#RunAsVirtualAccount = \$true** with the following text:

```
RunAsVirtualAccount = $true
```
18. Replace the line that starts with **# RoleDefinitions** with the following text:

```
RoleDefinitions = @{ 'Contoso\DNSOps' = @{ RoleCapabilities = 'DNSOps' };
```
19. Save the file.
20. In **Windows Admin Center**, in the **SEA-SRV1** Windows PowerShell window, use the cmdlet **Register-PSSessionConfiguration** to register the DNSOps JEA endpoint:
21. On **SEA-ADM1**, open Windows PowerShell as an administrator.
22. Connect to the DNSOps JEA endpoint using the **Import-PSSession** cmdlet, adding the prefix **DN-
SOps**.
23. Display the Windows PowerShell commands available from the DNSOps JEA endpoint.

Test your knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

What security benefit does JEA (Just Enough Administration) provide?

- Enables RBAC functionality for Windows PowerShell remoting
- Ensures only privileged user accounts can connect remote servers
- Allows remote users to perform all the same actions as a local administrator
- Prevents remote users from running any scripts on a remote server

Question 2

What file allows you to define which commands are available from a JEA endpoint?

- Role capability file
- Session configuration file
- Endpoint configuration file
- Session capability file

Question 3

When connected to remote Windows PowerShell session with the prefix DNSOps, which of the following commands would provide the available cmdlets?

- Get-DNSOpsCommand
- Get-Command -Noun DNSOps
- Get-Command -Name DNSOps
- List-Command -Name DNSOps

Securing and analyzing SMB traffic

Lesson overview

Server Message Block (SMB) protocol is a network protocol primarily used for file sharing. Along with its common file-sharing use, it's also frequently used by printers, scanners, and email servers. The original version of SMB, SMB 1.0 does not support encryption. SMB encryption was introduced with version 3.0.

Encryption is important whenever sensitive data is moved by using the SMB protocol. SMB encryption also lets file services provide secure storage for server applications such as Microsoft SQL Server and is generally simpler to use than dedicated hardware-based encryption.

In this lesson you'll learn about the security features of SMB 3.1.1, the latest and most secure version of SMB.

Lesson objectives

After completing this lesson, you will be able to:

- Describe SMB 3.1.1 protocol security.
- Describe the requirements for implementing SMB 3.1.1.
- Describe how to configure SMB encryption on SMB shares.
- Disable SMB 1.0.

What is SMB 3.1.1 protocol security?

SMB 3.0 introduced end-to-end encryption to the SMB (Server Message Block) protocol. SMB encryption provides for data packet confidentiality and helps prevent a malicious hacker from tampering with or eavesdropping on any data packet.

SMB 3.1.1, introduced in Windows Server 2016, provides several enhancements to SMB 3.0 security, including preauthentication integrity checks and encryption improvements. The version of SMB included with Windows Server 2019 is SMB 3.1.1.c.

Preauthentication integrity

With preauthentication integrity, while a session is being established the "negotiate" and "session setup" messages are protected by using a strong (SHA-512) hash. This helps prevent man-in-the-middle attacks that tamper with the connection. The resulting hash is used as input to derive the session's cryptographic keys, including its signing key. The final session setup response is signed with this key. If any tampering has occurred to the initial packets, the signature validation fails, and the connection would not be established. This enables the client and server to mutually trust the connection and session properties.

SMB encryption improvements

SMB 3.1.1 provides improvements to the following security features:

- **SMB encryption.** Introduces support for Advanced Encryption Standard (AES) Galois/Counter Mode (GCM) 128-bit encryption, along with continued support for AES 128-bit encryption.
- **Directory Caching.** Allows Windows to cache larger directories—up to 500,000 entries—and attempts to directory queries with 1-megabyte (MB) buffers to improve performance.

- Rolling cluster upgrade support. Lets SMB appear to support different max versions of SMB for clusters during upgrade.
- Support for **FileNormalizedNameInformation** API calls. Adds native support for querying the normalized name of a file. The normalized name is the exact name, including letter casing of files as stored on the disk.

SMB 3.1.1.c provides the following improvements to SMB 3.0 encryption:

- Write-through to disk. This feature allows write operations to ensure that writes to a file share make it to the physical disk. This feature is new to SMB 3.1.1.c.
- Guest access to file shares. The SMB client no longer allows Guest accounts to access a remote server or Fallback to Guest account when invalid credentials are provided.
- SMB global mapping. Maps remote SMB shares to drive letters accessible to all users on the local host, including containers. This allows containers to write to remote shares.
- SMB dialect control. Allows administrators to set the minimum and maximum SMB versions (also known as *dialect*), used on the system.

SMB 3.1.1 encryption requirements

Windows Server systems support multiple versions of SMB (Server Message Block). This enables them to communicate with servers and clients running other operating systems and other Windows versions. To use SMB 3.1.1, both your host server and the system it communicates with must support SMB 3.1.1.

Preauthentication with SMB 3.1.1 is not compatible with devices that modify SMB packets, such as some wide area network (WAN) accelerators. Therefore, you might need to replace some network equipment to use SMB 3.1.1.

When communicating with other supported operating systems, Windows Server 2019 will negotiate the following versions of SMB:

- Windows 10 and Windows Server 2016 or later – SMB 3.1.1
- Windows 8.1 and Windows Server 2012 R2 – SMB 3.02
- Windows 8 and Windows Server 2012 – SMB 3.0

Note: SMB encryption is not related to Encrypted File System (EFS) or BitLocker Drive Encryption.

Configure SMB encryption on SMB shares

SMB (Server Message Block) encryption `Verify` is not enabled by default, because encryption adds some overhead to network communications. If you want to use SMB encryption, you can configure it on a per-share basis or for an entire file server. Where you enable SMB encryption depends on your security needs. Enabling it restricts the file share or file server to only SMB 3.x clients for the protected shares.

To use Windows PowerShell to enable SMB encryption on an existing file share, use the following command from the server hosting the share:

```
Set-SmbShare -Name <sharename> -EncryptData $true
```

To encrypt all shares on a file server, from the server use the following command:

```
Set-SmbServerConfiguration -EncryptData $true
```

To create a new SMB file share on a server and enable SMB encryption at the same time, use the following command:

```
New-SmbShare -Name <sharename> -Path <pathname> -EncryptData $true
```

To allow connections that don't use SMB 3 encryption, such as when older servers and clients are still in your network, use the following command:

```
Set-SmbServerConfiguration -RejectUnencryptedAccess $false
```

Note: You can also configure SMB encryption using Server Manager.

Disable SMB 1.0

As an older protocol, SMB 1.0 exposes your server to a large, and possibly unnecessary attack surface. SMB 2 (Server Message Block 2), first introduced in Windows Vista and Windows Server 2008, is installed as a separate server component and supports SMB protocol versions 2.x and higher, including SMB 3.x. When an SMB connection is negotiated, the negotiation process prevents the connection from downgrading from SMB 3.0 to SMB 2.0. However, it can still downgrade to from SMB 3.0 to SMB 1.0. If all the systems your server communicates with are currently supported versions of Windows operating systems, you might no longer need support for SMB 1.0.

By default, SMB 1.0 is disabled in Windows Server 2019. If you have enabled it to support older printers or network-attached storage (NAS) devices that require SMB 1.0, you should disable SMB 1.0 as soon as it's no longer needed to ensure that clients always make an encrypted SMB connection.

To determine whether SMB 1.0 support is installed on your server, use the following Windows PowerShell command:

```
Get-WindowsFeature FS-SMB1
```

To uninstall SMB 1.0, use the following command:

```
Disable-WindowsOptionalFeature -Online -FeatureName smb1protocol
```

To disable SMB 1.0, use the following command:

```
Set-SmbServerConfiguration -EnableSMB1Protocol $false
```

If you need to reinstall SMB 1.0 later, you can use the **Enable-WindowsOptionalFeature** cmdlet.

Demonstration: Disable SMB 1.0 and configure SMB encryption on shares

In this demonstration, you will learn how to disable SMB 1.0 (Server Message Block 1.0) on servers where it has been enabled, and you will configure an SMB share.

Demonstration steps

Disable SMB 1.x on Windows Server

1. Sign-in to **SEA-ADM1**.
2. Open **Windows Admin Center**, and then connect to **SEA-SVR1**.
3. In **Windows Admin Center**, open a remote Windows PowerShell session.
4. Use the **Set-SmbServerConfiguration** cmdlet to disable SMB 1.0.

Configure a share for SMB encryption

1. In the remote PowerShell session, create the folder **c:\labfiles\mod08**.
2. Use the **New-SmbShare** cmdlet to create the share **Mod08**, and then encrypt it.
3. Use the **Grant-FileShareAccess** cmdlet to grant **Everyone** full access to the share **Mod08**.
4. In File Explorer, on **SEA-ADM1**, navigate to **\\SEA-SVR1\mod08**.

Test your knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

What SMB (Server Message Block) version is enabled in Windows Server 2019 by default?

- SMB 3.1.1.c
- SMB 3.2.2.c
- SMB 1.0
- SMB 1.1.2

Question 2

Which cmdlet would you use to create a new, encrypted SMB file share?

- `New-SmbShare -Name <sharename> -Path <pathname> -EncryptData $true`
- `Set-SmbShare -Name <sharename> -EncryptData $true`
- `Set-SmbServerConfiguration -EncryptData $true`
- `Set-SmbServerConfiguration -EnableSMB1Protocol $false`

Windows Server Update Management

Lesson overview

Another important security task is ensuring that your server has updates to software applied in a timely fashion. Windows Server Update Services (WSUS) provides infrastructure to download, test, and approve updates. When you install updates quickly—especially security updates—you help block attacks based on the vulnerabilities addressed by the update.

Lesson objectives

After completing this lesson, you will be able to:

- Describe the role of WSUS.
- Describe the WSUS update management process.
- Deploy updates with WSUS.

Overview of Windows Update

Windows Update is a Microsoft service that provides updates to Microsoft software. This includes service packs, security patches, drive updates, and even firmware updates.

Orchestrator software on a Windows device scans for and downloads updates. You can configure the orchestrator to get updates from a Windows Server Update Services by using Group Policy.

What is WSUS?

WSUS (Windows Server Update Services) is a server role that helps you download and distribute updates to Windows clients and servers. WSUS can obtain updates that are applicable to the operating system, and to common Microsoft products such as Microsoft Office and Microsoft SQL Server.

WSUS

WSUS provides a central management point for updates to your computers running Windows operating systems. By using WSUS, you can create a more efficient update environment in your organization and stay better informed about the overall update status of your network's computers.

In the simplest configuration, a small organization can have a single WSUS server that downloads updates from Microsoft Update. The WSUS server then distributes the updates to computers that are configured to obtain automatic updates from the WSUS server. You can choose whether updates need approval before clients can download them.

Larger organizations might want to create a hierarchy of WSUS servers. In this scenario, a single, centralized WSUS server obtains updates from Microsoft Update, and other WSUS servers obtain updates from the centralized WSUS server.

You can organize computers into groups, or *deployment rings*, to manage the process of deploying and approving updates. For example, you can configure a pilot group to be the first set of computers used for testing updates.

WSUS can generate reports to help monitor update installations. These reports can identify which computers have not yet applied recently approved updates. Based on these reports, you can investigate why updates are not being applied.

Prerequisites

To install the WSUS server role on a server, in addition to the requirements of the Windows Server 2019 operating system, it must meet the following requirements:

- Memory. An additional 2 gigabytes (GB) of random access memory (RAM) beyond that required for the server and all other services.
- Available disk space. 40 GB or greater available disk space.
- Reporting. Installation of the Microsoft Report Viewer 2012 Runtime.

The WSUS database requires either a Windows Internal Database (WID) or a SQL Server database. When using a SQL Server database, the database can live on another computer.

WSUS server deployment options

Before installing and configuring WSUS (Windows Server Update Services) servers, you must consider how to deploy WSUS in your environment. WSUS implementations vary in size and configuration depending on your network environment and how you want to manage updates. You could have a single WSUS server for your entire organization, multiple WSUS servers acting independently, or multiple WSUS servers connected to each other in a hierarchy.

Single WSUS server

The most basic implementation of WSUS uses a single WSUS server inside your network. This server connects to Microsoft Update and downloads updates through the firewall. The WSUS server uses port 8530 for HTTP communication and port 8531 for HTTPS, instead of the default 80 and 443, respectively. You need to make sure your firewall has the necessary rules to allow the server to connect to Microsoft Update. This basic scenario is commonly used for small networks with a single physical location.

Multiple WSUS servers

If your environment is composed of several isolated physical locations, you might need to implement a WSUS server in each location. In this scenario, each WSUS server has its own connection to the Internet to download updates from Microsoft Update.

Although this is a valid option, it requires substantially more administrative effort—especially as the number of physical locations grows—because you must manage each individual WSUS server independently. You have to download updates to each server separately, approve updates on each server individually, and manage WSUS clients so that they receive updates from the correct WSUS server.

Individual WSUS servers work well for organizations that have a small number of physical locations, where each physical location has its own IT management team. You can also use this scenario for a single physical location that has too many clients for a single WSUS server to manage, by placing multiple WSUS servers in a Network Load Balancing (NLB) cluster.

Disconnected WSUS servers

A *disconnected WSUS server* is a server that doesn't connect to Microsoft Update over the internet or receive its updates from any other server in the network. Instead, this server receives its updates from removable media generated on another WSUS server.

A disconnected WSUS server is most common in isolated network environments without internet access, such as in some high security environments. You can use a WSUS server in a different location to syn-

chronize with Microsoft Update, then export the updates to portable media, and then transport the portable media to the remote location to be imported into the disconnected WSUS server.

WSUS server hierarchies

All the scenarios we have discussed so far deal with an independently managed WSUS server that connects directly to Microsoft Update or receives its updates in a disconnected manner. However, in larger organizations with multiple physical locations you might want to have the ability to synchronize with Microsoft Update on one server. You might also want to push the updates to servers in various locations over your network and approve updates from a single location.

WSUS server hierarchies allow you to:

- Download updates to servers that are closer to clients, such as servers in branch offices.
- Download updates once to a single server, and then replicate the updates over your network to other servers.
- Separate WSUS servers based on the language used by their clients.
- Scale WSUS servers for a large organization that has more client computers than a single WSUS server can manage.

In a WSUS server hierarchy, there are two types of servers:

- Upstream servers. Upstream servers connect directly to Microsoft Update to retrieve updates, or are disconnected and receive updates by using portable media.
- Downstream servers. Downstream servers receive updates from a WSUS upstream server.

you can configure downstream servers in one of two modes:

- Autonomous mode. Autonomous mode, or *distributed administration*, allows a downstream server to receive updates from an upstream server, but enables administrators to maintain administration of the updates locally. In this scenario, the downstream server maintains its own set of computer groups, and updates can be approved independently of approval settings in the upstream servers. This allows a different group of administrators to manage updates at their own locations, and only use the upstream server as a source of downloadable updates.
- Replica mode. Replica mode, or *centralized administration*, allows a downstream server to receive updates, computer group membership information, and approvals from an upstream server. In this scenario, a single group of administrators can manage updates for the entire organization. In addition, downstream servers can be placed in different physical offices and receive all updates and management data from an upstream server.

You can have multiple layers in your WSUS hierarchy. You can configure some of your downstream servers to use autonomous mode, while you can use replica mode to configure other servers. For example, you can have a single upstream server connected to Microsoft Update, downloading updates for your entire organization. You can have two other downstream servers in autonomous mode, one that manages updates for all computers running software in English, and another for all computers running software in Spanish. Finally, you can have another set of downstream servers receiving their updates from the middle-tier WSUS servers configured in replica mode. These are the actual servers that clients receive updates from, but all the management is done at the middle tier.

Note: You can configure downstream servers to download the update information metadata from an upstream server, but to download the actual updates themselves from Microsoft Update. This is a common configuration when the downstream servers have good internet connectivity and you want to reduce WAN traffic.

The WSUS update management process

The update management process enables you to manage and maintain WSUS (Windows Server Update Services) and the updates retrieved by WSUS. This process is a continuous cycle during which you can reassess and adjust the WSUS deployment to meet changing needs. The four phases in the update management process are:

- Assess
- Identify
- Evaluate and plan
- Deploy

The assess phase

The goal of the assess phase is to set up a production environment that supports update management for routine and emergency scenarios. After initial setup, the assess phase becomes an ongoing process that you use to determine the most efficient topology for scaling the WSUS components. As your organization changes, you might identify the need to add more WSUS servers in different locations.

As part of the Assess phase, you will decide how you will synchronize updates from Windows Update, and which WSUS servers will download those updates. You can choose to synchronize updates based on:

- Product or product family. For example, you could select updates for:
 - All Windows operating systems.
 - All editions of a specific version, such as Windows Server 2019.
 - A specific edition, such as Windows Server 2019 Datacenter edition.
- Classification. For example, you can choose critical updates or security updates.
- Language. You can choose all languages or choose from a subset of languages.

You will also decide whether your WSUS servers will get updates directly from Windows Update or from another WSUS server.

The identify phase

During the identify phase, you identify new updates that are available and determine whether they're relevant to your organization. WSUS automatically identifies which updates are relevant to registered computers.

The evaluate-and-plan phase

After you've identified the relevant updates, you need to evaluate whether they work properly in your environment. It's always possible that the specific combination of software in your environment might have problems with an update.

To evaluate updates, you should have a test environment in which you can apply updates to verify proper functionality. During this time, you might identify dependencies required for an update to function properly, and you can then plan for any changes that you need to make.

To accomplish this, you should use one or more computer groups for testing purposes. For example, you can create a computer group of non-production servers that run the different applications and operating

systems that are updated by WSUS. Before you deploy updates to the production environment, you can push updates to these computer groups, test them, and after making sure they work as expected, deploy these updates to the organization.

The Deploy Phase

After you have thoroughly tested an update and determined any dependencies, you can approve it for deployment in the production network. Ideally, you should approve the update for a pilot group of servers that run the least-critical applications before approving the update for the entire organization and for business-critical servers. You can configure WSUS to approve updates automatically, but that isn't recommended.

Troubleshooting WSUS

After your WSUS environment is configured and in use, you might still find problems. Some problems are easier to manage, while others might require the use of special debugging tools. Here's a list of common problems you could encounter when managing a WSUS environment:

- Computers not displaying in WSUS. This is typically a result of client computer misconfiguration, or a Group Policy Object (GPO) not applied to the client computer.
- WSUS server stops with a full database. When this happens, you'll notice a SQL Server dump file (SQLDumpnnnn.txt) in the **LOGS** folder for SQL Server. This is usually a result of index corruption in the database. You might need help from a SQL Server database administrator (DBA) to recreate indexes, or you might simply need to reinstall WSUS to fix the problem.
- You cannot connect to WSUS. Verify network connectivity and ensure the client can connect to the ports used by WSUS by using the **Test-NetConnection** cmdlet.

Microsoft also provides tools and utilities that you can use to help troubleshoot issues with WSUS.

Additional reading: For more information, refer to **Windows Server Update Services Tools and Utilities**¹⁶.

Azure Update Management

If you have an Azure account, you can also use Microsoft Azure Update Management, a feature of Azure Automation, in conjunction with WSUS (Windows Server Update Services) or instead of WSUS to manage updates on your servers.

What is Azure Automation?

Azure Automation is a cloud-based service that provide process automation, configuration management, update management, and other management features for both Azure and non-Azure environments, including on-premises environments.

What is Update Management?

Update Management is a free service within Azure Automation that helps you manage operating system update for both Windows and Linux machines, both in the cloud and on-premises. The only cost associated with using Update Management is the cost of log storage in Azure Log Analytics.

¹⁶ <https://aka.ms/wsus-tools>

Azure Update Management does not require configuring Group Policies for updates, making it simpler to use than WSUS in many cases. As previously stated, you can use it to manage updates for both Windows and Linux servers, making it a good choice for mixed server environments. Also, because you can use it with cloud-based servers, it's also a good option for managing updates in hybrid environments.

On-premises servers are managed via a locally installed agent on the server that communicates with the cloud service.

Update Management capabilities

Update Management includes the following capabilities related to on-premises servers:

- Check status of updates on your servers. The Update Management service includes a cloud-based console from where you can review the status of updates across your organization and for a specific machine.
- Ability to configure dynamic groups of machines to target. This service allows you to define a query based on *computer group*, which is a group of computers that are defined based on another query or imported from another source such as WSUS or Endpoint Configuration Manager.
- Ability to search Azure Monitor logs. Records collected by Update Management are listed in Azure Monitor Logs.

Onboarding your on-premises server

You must add your on-premises servers to Update Management in Azure Automation manually. After you have an Azure Automation account, you then need to enable either the **Inventory** or **Change tracking** features within your Automation account. After either of those solutions is running, then you can enable the **Update management** solution with your account.

After you enable Update Management, you then download and install the Log Analytics agent for Windows to your on-premises server. After the agent is installed and information about your server is reported to your automation account, you can then onboard the machines in the update workspace within your Automation account.

Additional reading: For more information about Azure Update Management, refer to **Update Management overview**¹⁷.

Additional reading: For more information about installing the Log Analytics agents, refer to **Connect Windows computers to Azure Monitor**¹⁸.

Test your knowledge

Use the following questions to check what you've learned in this lesson.

¹⁷ <https://aka.ms/update-management>

¹⁸ <https://aka.ms/agent-windows>

Question 1

What are the options for a WSUS (Windows Server Update Services) database? Choose two:

- Windows Internal Database (WID)
- SQL Server
- MariaDB
- MySQL

Question 2

Which is not a valid WSUS server deployment option?

- Single WSUS server
- Multiple WSUS servers
- Disconnected WSUS servers
- Autonomous WSUS servers

Question 3

Which are steps in the update management process? Choose three.

- Assess
- Identify
- Classify
- Deploy

Question 4

Azure Update Management is part of what Azure service?

- Azure Automation
- Azure Sentinel
- Azure Monitor
- Azure AD DS (Active Directory)

Module 08 lab and review

Lab: Configuring security in Windows Server

Scenario

Contoso Pharmaceuticals is a medical research company with about 5,000 employees worldwide. They have specific needs for ensuring that medical records and data remain private. The company has a headquarters location and multiple worldwide sites. Contoso has recently deployed a Windows Server and Windows client infrastructure. You have been asked to implement improvements in the server security configuration.

Objectives

After completing this lab, you will be able to:

- Configure Windows Defender Credential Guard.
- Locate problematic user accounts.
- Implement and verify LAPS (Local Administrator Password Solution)

Estimate time: 40 minutes

Module review

Use the following questions to check what you've learned in this module.

Question 1

What should an organization do before it institutes NTLM blocking?

- Audit NTLM usage
- Configure the Restrict NTLM: NTLM Authentication Group Policy
- Enable Kerberos authentication

Question 2

Which Windows PowerShell cmdlet do you use to configure a specific OU so that computers within that OU can use LAPS (Local Administrator Password Solution)?

- Disable-ADAccount
- Update-AdmPwdADSchema
- Get-AdmPwdPassword
- Set-AdmPwdComputerSelfPermission

Question 3

Which SMB (Server Message Block) version is negotiated by Windows Server 2019 when communicating with Windows Server 2012 R2?

- SMB 1.0
- SMB 2.0
- SMB 3.02
- SMB 3.1.1

Answers

Question 1

Which security setting should not be enabled when configuring administrative user accounts?

- Logon Hours
- Account is sensitive and cannot be delegated
- This account supports Kerberos AES (Advanced Encryption Standard) 256-bit encryption
- Do not require Kerberos preauthentication

Explanation

"Do not require Kerberos preauthentication" is the correct answer. Kerberos preauthentication reduces the risk of replay attacks. Therefore, you should not enable this option. All other answers are valid ways to configure additional security for administrative user accounts.

Question 2

Which feature allows you to configure TGT (Ticket-granting tickets) lifetime and access-control conditions for a user?

- Protected Users group
- Authentication policies
- Authentication policy silos
- NTLM blocking

Explanation

"Authentication policies" is the correct answer. Authentication policies allow you to configure TGT lifetime and access-control conditions for a user, service, or computer account. The AD DS (Active Directory) security group Protected Users helps you protect highly privileged user accounts against compromise. Authentication policy silos allow administrators to assign authentication policies to user, computer, and service accounts. NTLM blocking prevents the use of the NTLM authentication protocol, which is less secure than the Kerberos authentication protocol.

Question 3

Which is not a valid way to enable Windows Defender Credential Guard on a server?

- Group policy
- Adding server role
- Updating the registry
- Using a Windows PowerShell script

Explanation

"Adding server role" is the correct answer. You cannot enable Windows Defender Credential Guard through a server role. However, you can enable Windows Defender Credential Guard by using a Group Policy object, by updating the registry on the server, or by running the Hypervisor-Protected Code Integrity and Windows Defender Credential Guard hardware readiness tool, which is a Windows PowerShell script.

Question 4

What are two types of problematic user accounts you should check for regularly?

- Users with passwords that do not expire
- Users that have not signed in recently
- Users with complex passwords
- Users with few administrative permissions

Explanation

Users with passwords that do not expire or who have not signed in for an extended period of time are both problematic accounts that you should identify and remediate on a regular schedule. Passwords that do not expire are considered insecure. Therefore, you should disable user accounts that are not being used to limit a potential avenue of attack. Complex passwords are not considered insecure, and limiting user permissions to only those needed (the principle of least privilege) is considered a best practice.

Question 1

Which of these is a capability of LAPS (Local Administrator Password Solution)?

- Verify the local administrator password is the same on all managed servers.
- Store local administrator passwords in Microsoft Exchange.
- Prevent local administrator passwords from expiring.
- Ensure that local administrator passwords are unique on each managed server.

Explanation

"Ensure local administrator passwords are unique on each managed server" is the correct answer. LAPS doesn't verify the local administrator password is the same on all managed servers, but it does make sure they are unique. LAPS doesn't store local administrator passwords in Exchange, it stores them in AD DS. Finally, LAPS doesn't prevent local administrator passwords from expiring, but it does set an expiration date and automatically changes the password before that date.

Question 2

When configuring a PAW (Privileged Access Workstation), which of these should you not do?

- Ensure that only authorized users can sign in to the PAW. Standard user accounts should not be able to sign in.
- Enable Windows Defender Credential Guard to help protect against credential theft.
- Ensure the PAW can access the internet.
- Limit physical access to the PAW.

Explanation

"Ensure the PAW can access the internet" is the correct answer. You should not enable PAWs to access the internet, because it's a significant source of cyberattacks. All the other options are valid ways to secure PAWs.

Question 3

Which options are valid ways to secure a domain controller? Select all that apply.

- Ensure that domain controllers run the most recent version of the Windows Server operating system and have current security updates.
- Deploy domain controllers by using the "Server Core" installation option.
- Configure RDP (Remote Desktop Protocol) through Group Policy to limit RDP connections to domain controllers, so they can occur only from PAWs.
- Configure the perimeter firewall to block outbound connections to the internet from domain controllers.

Explanation

All of these are valid options for securing domain controllers.

In addition, you should keep physically deployed domain controllers in dedicated, secure racks that are separate from other servers. You should run virtualized domain controllers either on separate virtualization hosts or as a shielded virtual machine on a guarded fabric. You should also review CIS (Center for Internet Security) benchmarks for Windows Server operating systems for security guidance specific to domain controllers and use Device Guard to control the execution of scripts and executables on the domain controller.

Question 4

What CIS hardening level maps to the security configuration baselines included in the SCT (Microsoft Security Compliance Toolkit)?

- Level 0
- Level 1
- Level 2
- None

Explanation

"Level 1" is the correct answer. The security baselines included in the SCT align closely to CIS Level 1 benchmark hardening guidelines.

Question 1

What security benefit does JEA (Just Enough Administration) provide?

- Enables RBAC functionality for Windows PowerShell remoting
- Ensures only privileged user accounts can connect remote servers
- Allows remote users to perform all the same actions as a local administrator
- Prevents remote users from running any scripts on a remote server

Explanation

"RBAC functionality for Windows PowerShell remoting" is the correct answer. JEA provides Windows Server and Windows client operating systems with RBAC functionality built on Windows PowerShell remoting. It also allows user accounts that are not privileged to connect to a JEA endpoints and perform administrative tasks. While JEA gives a user local administrator privileges on a remote server, JEA endpoints limit users only to specific activities defined by JEA. JEA endpoints can be configured to allow remote users to run some scripts, providing they run them from within Windows PowerShell.

Question 2

What file allows you to define which commands are available from a JEA endpoint?

- Role capability file
- Session configuration file
- Endpoint configuration file
- Session capability file

Explanation

"Role capability file" is the correct answer. Role capability files help you specify what can be done in a Windows PowerShell session. Session configuration files are used to register a JEA endpoint, and there are no Endpoint configuration files or Session capability files in JEA.

Question 3

When connected to remote Windows PowerShell session with the prefix DNSOps, which of the following commands would provide the available cmdlets?

- Get-DNSOpsCommand
- Get-Command -Noun DNSOps
- Get-Command -Name DNSOps
- List-Command -Name DNSOps

Explanation

"Get-DNSOpsCommand" is the correct answer. The following command will add the prefix DNSOps to the commands available in a remote PowerShell session: `Import-PSSession -Session MySessionObject -Prefix 'DNSOps'`. "Get-Command -Noun DNSOps" would retrieve any cmdlets that have the noun DNSOps in their name. "Get-Command -Name DNSOps" would retrieve a cmdlet named DNSOps, which would be a non-standard cmdlet name, and "List-Command" is not a valid PowerShell command.

Question 1

What SMB (Server Message Block) version is enabled in Windows Server 2019 by default?

- SMB 3.1.1.c
- SMB 3.2.2.c
- SMB 1.0
- SMB 1.1.2

Explanation

"SMB 3.1.1.c" is the correct answer. Windows Server 2019 supports SMB 3.x and SMB 2.x, but SMB 2.x is not listed. The default server configuration for Windows Server 2019 does not install support for SMB 1.x, but it is available.

Question 2

Which cmdlet would you use to create a new, encrypted SMB file share?

- New-SmbShare -Name <sharename> -Path <pathname> -EncryptData \$true
- Set-SmbShare -Name <sharename> -EncryptData \$true
- Set-SmbServerConfiguration -EncryptData \$true
- Set-SmbServerConfiguration -EnableSMB1Protocol \$false

Explanation

"New-SmbShare -Name <sharename> -Path <pathname> -EncryptData \$true" is the correct answer. "Set-SmbShare -Name <sharename> -EncryptData \$true" encrypts an existing SMB share. "Set-SmbServerConfiguration -EncryptData \$true" encrypts all existing SMB shares on a server. "Set-SmbServerConfiguration -EnableSMB1Protocol \$false" disables SMB 1.x support if was previously enabled.

Question 1

What are the options for a WSUS (Windows Server Update Services) database? Choose two:

- Windows Internal Database (WID)
- SQL Server
- MariaDB
- MySQL

Explanation

"Windows Internal Database (WID)" and "SQL Server" are both correct answers. MySQL and MariaDB are not valid database options for the WSUS database.

Question 2

Which is not a valid WSUS server deployment option?

- Single WSUS server
- Multiple WSUS servers
- Disconnected WSUS servers
- Autonomous WSUS servers

Explanation

"Autonomous WSUS servers" is the correct answer. "Autonomous WSUS servers" is not a valid deployment option. Downstream WSUS servers can be deployed in "Autonomous mode". The remaining options are all valid options for deploying WSUS servers.

Question 3

Which are steps in the update management process? Choose three.

- Assess
- Identify
- Classify
- Deploy

Explanation

"Assess", "Identify", and "Deploy" are the correct answers. The update management process includes the following steps: Access, Identify, Evaluate and Plan, and Deploy. Classify is not a step in the update management process. However, during the Assess phase you will decide what classification of updates you want to deploy.

Question 4

Azure Update Management is part of what Azure service?

- Azure Automation
- Azure Sentinel
- Azure Monitor
- Azure AD DS (Active Directory)

Explanation

"Azure Automation" is the correct answer. Update Management is a free service within Azure Automation that helps you manage operating system updates for both Windows and Linux machines, both in the cloud and on-premises. Update Management is not included with the other services listed:

Question 1

What should an organization do before it institutes NTLM blocking?

- Audit NTLM usage
- Configure the Restrict NTLM: NTLM Authentication Group Policy
- Enable Kerberos authentication

Explanation

Prior to blocking NTLM, you should ensure that existing applications are no longer using the protocol. You can audit NTLM traffic by enabling policies in the Computer Configuration\Policies\Windows Settings\Security Settings\Local Policies\Security Options node. After you perform the audit and determine there are no existing applications that use the protocol, you will configure the Restrict NTLM: NTLM Authentication Group Policy to enable NTLM blocking. Kerberos authentication is already enabled in Windows Server.

Question 2

Which Windows PowerShell cmdlet do you use to configure a specific OU so that computers within that OU can use LAPS (Local Administrator Password Solution)?

- Disable-ADAccount
- Update-AdmPwdADSchema
- Get-AdmPwdPassword
- Set-AdmPwdComputerSelfPermission

Explanation

You use the `Set-AdmPwdComputerSelfPermission` cmdlet to configure a specific OU so that computers within that OU can use LAPS.

The `Get-AdmPwdPassword` cmdlet retrieves a local administrator password assigned to a computer.

The `Update-AdmPwdADSchema` cmdlet updates the AD DS (Active Directory) schema in preparation for using LAPS.

The `Disable-ADAccount` disables user accounts in AD DS.

Question 3

Which SMB (Server Message Block) version is negotiated by Windows Server 2019 when communicating with Windows Server 2012 R2?

- SMB 1.0
- SMB 2.0
- SMB 3.02
- SMB 3.1.1

Explanation

When communicating with a Windows Server 2012 R2 server, Windows Server 2019 (and Windows Server 2016) negotiate using SMB 3.02.

Windows Server 2019 uses SMB 3.1.1 when communicating with Windows Server 2016 or later.

Windows Server 2019 uses SMB 2.0 for communicating with operating systems prior to Windows 8.

After disabling SMB 1.0, as recommended in this course, Windows Server 2019 will not use it to communicate with any device.



Module 9 RDS in Windows Server

Overview of RDS

Lesson overview

Remote Desktop Services (RDS) is a Windows Server role that provides much more than just remote desktops. Similar functionality was known as *Terminal Services (TS)*, but RDS has evolved dramatically since then. RDS includes six role services that enable you to create a scalable and fault-tolerant RDS deployment. You can manage an RDS deployment centrally and in the same way, regardless of the number of servers in an RDS deployment.

In this lesson, you'll learn about RDS, explore how it provides an enhanced user experience for remote users, and compare it with the Microsoft Remote Desktop feature. You'll also learn how to connect to RDS and about the licenses that are required to use it. This lesson focuses on a session-based RDS desktop deployment, and will briefly discuss options for Remote Desktop Services in Microsoft Azure.

Lesson objectives

After completing this lesson, you'll be able to:

- Describe and understand RDS.
- Explain the benefits of using RDS.
- Understand the **Client Experience** features in RDS.
- Explain Remote Desktop features with RDS.
- Plan an RDS deployment.
- Understand how to use RDS.
- Explain Remote Desktop Gateway.
- Understand RDS licensing.
- Explain the options for RDS in Azure.

RDS overview and benefits

Remote Desktop Services (RDS) is a server role in the Windows Server OS that enables you to use session-based deployments to provide multiple users with a virtual desktop from a single server. Users connect to the server to run their applications and access other network resources. These applications are installed on the server, referred to as the *Remote Desktop Session Host* (RD Session Host).

Each user has an independent desktop when they connect to an RD Session Host, and users can't observe other users' desktops. However, each RD Session Host has limited resources, such as memory and disk throughput. Therefore, the limited resources on each RD Session Host limit the number of users that can connect and run applications simultaneously.

When users connect to an RD Session Host, screen display information, mouse movements, and key-strokes are sent across the network between the RD Session Host and the client computer. RDP is used between client computers and the RD Session Host, which is very efficient and consumes little network bandwidth. This makes it possible to use a session-based desktop deployment from remote locations over the Internet and other slow networks. In some cases, a session-based desktop deployment of RDS is a replacement for accessing files directly over a virtual private network (VPN) for mobile users. Remote desktop clients that use RDP are available for many OSs. Microsoft provides an RDP client for Windows clients, MacOS, iOS, and Android.

Note: Some software is not compatible with session-based RDS desktop deployment because it's not designed for multiple session use. In such cases, you might be able to mitigate compatibility issues by using virtual machine (VM)-based desktops instead because they use Windows 10.

RDS role services

RDS has several different role services that you can install. An RDS deployment always includes at least three RDS role services, which are required to manage a deployment. You can install additional individual RDS role services, but you won't be able to manage them unless they are part of an RDS deployment. Depending on your implementation goals, an RDS deployment can include additional RDS role services, which you can install on multiple servers for scalability and high availability.

RDS in Windows Server includes the following role services:

- **RD Session Host.** With this role service, a server can host Windows-based programs or a full Windows desktop. Users can connect to an RD Session Host server, run applications, and use the network resources that the RD Session Host offers. However, access is dependent on a Remote Desktop Connection (RDC) client or any other RDP client. RD Session Host is a required role service in a session-based RDS desktop deployment.
- **Remote Desktop Virtualization Host (RD Virtualization Host).** This role service integrates with the Microsoft Hyper-V role in Windows Server to provide VMs that can be used as virtual desktops. It also monitors and reports on established client sessions to a Remote Desktop Connection Broker (RD Connection Broker) server. If a VM has no connection for a configured interval, providing it's being used as a virtual desktop and is in a saved state, it will start. For pooled VMs, the RD Virtualization Host will revert them to their initial state when users sign out. RD Virtualization Host is a required role service in an RDS VM-based desktop deployment.
- **RD Connection Broker.** This role service manages connections to RemoteApp programs and virtual desktops, and directs client connection requests to an appropriate endpoint. It also provides session reconnection and session load balancing. For example, when a user disconnects from a session and later reestablishes a connection, the RD Connection Broker role service ensures that the user reconnects to their existing session. Although this role service is mandatory in each RDS deployment, it does not require large amounts of server resources.

- **Remote Desktop Web Access (RD Web Access).** This role service provides users with links to RDS resources, which can be RemoteApp programs, remote desktops, or virtual desktops, through a web browser. A webpage provides a user with a customized view of all RDS resources that have been published to that user. This role service supports organizing resources in folders, which allows administrators to group remote applications in a logical manner. It also publishes available RDS resources in an RDWeb feed, which can integrate with the Start screen on client devices. RD Web Access is a mandatory role service for each RDS deployment.
- **Remote Desktop Licensing (RD Licensing).** This role service manages RDS client access licenses (RDS CALs) that are required for each device or user to connect to an RD Session Host server. You use RD Licensing to install, issue, and track RDS CAL availability on an RD Licensing server.
- **Remote Desktop Gateway (RD Gateway).** With this 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 Remote Desktop connection authorization policies (RD CAPs) and 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 can connect.

RDS in Server Manager consolidates all aspects of RDS management into one location. The interface provides an overview of all servers in an RDS deployment and a management interface for each server. RDS in Server Manager uses a discovery process to detect the role services that are installed on each machine that is added to Server Manager.

Benefits of using RDS

RDS provides technologies that enable you to access session-based desktop deployments, VM-based desktop deployments, and remote applications that run on centralized servers. It also provides an enhanced desktop and application experience, and you can connect to it securely from managed or unmanaged devices. You can establish secure connections to an RDS deployment from both a local network and the internet.

RDS provides the following capabilities:

- You can observe and control applications, remote desktops, or virtual desktops from almost any client device that runs in a datacenter from a remote location.
- You can maintain installation and management on centralized servers in the datacenter. RDS sends an image of a remote desktop to a user's device, and the user's interaction with the desktop, keystrokes, and mouse movements is sent back to the RDS server.
- You can present users with a full desktop environment, or with individual application windows and the data within those applications that they require for their jobs.
- Remote RDS applications can integrate with users' local desktops. They show and behave as if they are a local application installed on a client desktop.
- RDS provides standard user desktops—either session-based or VM-based—that users can access from almost anywhere and from any device.
- RDS enables secure remote access to an entire desktop, remote application, or VM without establishing a VPN connection.
- You can centrally control which users can access RDS servers and which RDS servers can be accessed. You also can specify additional configurations, such as device redirection settings.

Running applications from an RDS deployment instead of installing them on each client computer provides several benefits, including:

- Quick application deployment. You can quickly deploy Windows-based programs to various devices across an enterprise. RDS is especially useful when you have programs that are frequently updated, infrequently used, or difficult to manage.
- Application consolidation. You can install and run programs from an RDS server and eliminate the need to update programs on each client computer.
- Settings and data sharing. Users can access the same settings and data from local devices and their RDS sessions.
- Remote access. Users can also access remote programs from multiple device types such as home computers, kiosks, and tablets. They can even access these programs from devices running legacy and non-Windows OSs.
- Data protection. Applications and data are stored centrally, not on client devices. Therefore, you can perform central backups and if a client device is compromised, it can be replaced without affecting a user's applications and data.
- Branch office access. RDS can provide better program performance for branch office users who need access to centralized data stores. Data-intensive programs often are not optimized for low-speed connections, and such programs often perform better over an RDS connection than running applications locally and accessing data remotely over a wide area network (WAN).
- Low utilization of client computers. Users can run applications that have high random access memory (RAM) and central processing unit (CPU) requirements on client desktops that have low computing power, because all of the data processing takes place on the RDS server.

You can deploy RDS in one two scenarios:

- A VM-based desktop deployment scenario provides more than just a connection to a remote computer; it's a connection to a hosted VM that runs a full Windows client OS. You can implement a VM-based desktop deployment as a pool of identical VMs that are temporarily assigned to users for the duration of a session, which then revert to their original configuration. You also can implement a VM-based desktop deployment as personal virtual desktops, in which a VM is permanently assigned to a specific user who always connects to the same virtual desktop.
- A session-based desktop deployment scenario allows users to access remote applications that run on the client's computing device as if they are installed locally. These application types are referred to as *RemoteApp programs*. It also allows users to connect to and access a client computer's full desktop and installed applications.

Note: The Remote Desktop feature is part of the Windows OS. RDS builds on top of the Remote Desktop feature and provides much greater functionality.

Client experience features with RDS

Users can connect to Remote Desktop Services (RDS) from various platforms. When a client connects to RDS it provides a graphical user interface (GUI) with either a full desktop or an application interface. Even if users are using different client software, they will all use Remote Desktop Protocol (RDP) to connect to RDS.

RDP provides rich, full-featured desktop functionality, and added security through encrypting network communication. It enables flexible connections by reducing the bandwidth used over slower connections, and offloading RDS servers by redirecting media streaming to the client. RDP can use either Transmission

Control Protocol (TCP) or User Datagram Protocol (UDP) as a transport mechanism, and its platform independent.

When a client connects to RDS, the user experience is similar to when they use local resources. Some RDS client-experience features are:

- **Bandwidth reduction.** When an RDP connection establishes, it uses various methods to reduce network bandwidth such as data compression and caching. Caching enables an adaptive user experience over local area networks (LANs) and wide area networks (WANs). Clients can detect available bandwidth and adjust the level of graphical detail being used.
- **Full desktop or application window only.** When a client connects to RDS, it can display either a full remote desktop or only the window of a remotely running application (RemoteApp program). With full desktops, users can perform remote administration or run multiple applications. However, the user must manage two desktops—both the local and remote. RemoteApp programs integrate with local desktops, but they still require network connectivity to RDS.
- **RemoteApp program appearance and behavior similar to installed applications.** When a user connects to a remote application that runs on the RDS RemoteApp program, the application's window displays.
- **RemoteApp program icons support pinning, tabbed windows, live thumbnails, and overlay icons.** Therefore, clients can add links to RemoteApp programs to their **Start** menu. RemoteApp windows can be transparent, and the content of a RemoteApp window displays while you are moving it. RemoteApp is now integrated with **Action Center** in Windows 10, which means support for notifications.
- **Automatic reconnect.** If a user disconnects from a remote desktop session, they can reconnect to the session and continue to work from the point at which they disconnected. The user can reconnect from the same device or connect from a different client device. If a session disconnects for a different reason, for example, because network connectivity is lost, the user automatically reconnects to the disconnected session when network connectivity restores.
- **Redirection of local resources.** Client resources such as drives, printers, clipboards, smart card readers, and USB devices can redirect to a remote desktop session. This enables users to use locally attached devices while working on RDS, and use a clipboard to copy content between a local and remote desktop. Users can even redirect USB devices that they plug in when the RDC is already established.
- **Windows media redirection.** This feature provides high-quality multimedia by redirecting Windows media files and streams from RDS to a client. Multimedia does not render on RDS—it is sent as a series of bitmaps to a client. However, audio and video content redirects in its original format, and all processing and rendering happens on the client. This offloads the RDS server, but also provides it with the same experience as when accessing multimedia content locally. RDP 1.0 has support for optimized H.264/AVC 444 codec playback.
- **Multiple monitor support.** This feature enables support for up to 16 monitors of any size, resolution, or layout. Applications function just as when you run them locally in configurations with multiple monitors.
- **Discrete device assignment.** This feature supports using physical graphics processing units (GPUs) in Windows Server 2016 or later Hyper-V hosts. It enables personal session desktops running in VMs to use the GPU from the Hyper-V host, and results in real-time performance from graphics-intensive applications running in Virtual Desktop Infrastructure (VDI).
- **Single sign-on (SSO).** When users connect to RDS, they have to provide their credentials again. With SSO, a user can connect to a remote desktop or start a RemoteApp program without having to reenter their credentials. SSO is also supported by the web client.

- CPU, Disk, and Network Fair Share. **Dynamic CPU Fair Share, Dynamic Disk Fair Share, Dynamic Network Fair Share, or Remote Desktop Services Network Fair Share** features are enabled by default on RDS to ensure even resource distribution among users. This prevents one user from monopolizing resources, thereby negatively affecting the performance of other users' sessions. Fair Share can dynamically distribute network, disk, and CPU resources among user sessions on the same Remote Desktop Session Host (RD Session Host) server. You can control these Fair Share settings through Group Policy.

Remote Desktop feature and RDS

Remote Desktop is part of the Windows operating system (OS). Remote Desktop Services (RDS) is a Windows Server role that depends on Remote Desktop and considerably expands it. They both have similar names and can provide similar user experiences, but RDS includes many additional features.

The Remote Desktop feature

Remote Desktop enables you to connect to a computer remotely and review its desktop just as when you sign in locally to that computer. The primary purpose for Remote Desktop is remote administration, which is why only the administrator who enables it can connect to the remote desktop. However, other users can connect to the remote desktop if the administrator grants them permission, and today users often use it to access their client desktop from home computers or when using mobile devices, by using either virtual private network (VPN) or DirectAccess. For example, it enables users to access their client desktop when using mobile devices or from home computers without establishing either VPN or DirectAccess connections.

You can enable Remote Desktop on both a client Windows OS and a server. When using it, you can typically only observe a full desktop. In general, you can't connect to individual applications that run on a remote computer—the exception to this is when you're using Remote Desktop in the Virtual Desktop Infrastructure (VDI) environment of a client computer running a Windows OS Enterprise edition.

The number of users who can use Remote Desktop simultaneously is limited. On a Windows client OS, only a single user can be connected, either by signing in locally or by using Remote Desktop. On a Windows Server without RDS, Remote Desktop allows two connections for administrative purposes only. Both users can be remote, or one user can sign in locally while the other establishes a remote connection. You don't need additional licenses to use Remote Desktop. You can establish as many remote desktop connections as you want from a single Windows-based computer.

RDS

RDS is a Windows Server role that's available only in the Windows Server OS. To deploy RDS, you need to install at least three role services and perform additional configuration. RDS provides a similar experience as Remote Desktop, in that it enables you to connect to a remote server and access the server's desktop. However, RDS can also present you with a window of the application that is running on the server (a *RemoteApp program*) or with a desktop of a VM that runs on the server.

The primary intention of RDS is to enable users to have a standard remote environment that is available from any device, and to use remote resources while integrating remote applications on the local user desktop. This provides users with an enhanced experience when they connect to RDS, which is similar to working with local applications. All rich client experience features are available with RDS, including:

- Hardware acceleration for video playback
- 4K downsampling

- Advanced device redirection for video cameras
- USB Redirection
- Multimedia redirection.

While you can provide users who connect to RDS with full desktops, you can also provide only RemoteApp programs).

There's no technical limit to how many users can connect to RDS. You're only limited to the available hardware resources and the number of RDS client access licenses (CALs) that are available. Each user or device that connects to RDS must have a valid RDS CAL, in addition to the license for their local device.

RDS includes several role services, including Remote Desktop Web Access and Remote Desktop Gateway. These services enable clients to connect securely over the internet to RDS, in addition to Remote Desktop. You can establish as many RDS connections as you want from a single Windows-based computer.

Overview of the RDC client

In addition to enabling users to connect to remote desktops, the RDC client can also connect to a physical desktop, a Remote Desktop Virtualization Host (RD Virtualization Host), or a Remote Desktop Session Host (RD Session Host). The RDC client is included with the Windows OS, and is also available as a separate download for additional platforms including Android, iOS, and macOS.

RDC uses RDP to transfer user actions, mouse movements, keyboard inputs, and redirected devices to the RD Session Host and graphical display from the RD Session Host to the RDC client. The RDC client can display an entire remote desktop or just the window of the running RemoteApp program.

Note: RDC can connect to a remote desktop if you use either the Transmission Control Protocol (TCP) or User Datagram Protocol (UDP). The protocol type is selected automatically based on the connectivity, available bandwidth, and speed of connection.

RDC has the following configuration tabs:

- **General.** You use this tab, to specify the computer to which you want to connect. You also can save RDC settings in a text file with an .rdp file name extension to initiate a connection later and to avoid configuring RDC settings again.
- **Display.** On this tab, you can choose the size of the remote desktop window, including the option to run the remote desktop in full-screen mode. You can select to use all local monitors for a remote session, select color depth, and enable a connection bar when the remote desktop is running in full-screen mode.
- **Local Resources.** From this tab you can set remote audio settings, such as whether you want to enable remote audio playback and recording. You also can specify a location where Windows key combinations (such as Alt+Tab) are applied, and whether local devices and resources in remote sessions are available. For example, you can select the option to make the clipboard, local drive, printers, and devices that you plug in available later in a remote session.
- **Advanced.** On this tab, you can configure server authentication and connect from anywhere settings.
- **Experience.** On this tab, you can select a connection speed to optimize performance. By default, RDC automatically detects connection quality and configures connection-quality-dependent features accordingly. From here you can also configure persistent bitmap caching and automatic reconnection if a connection drops. You can also enable different features, such as:
 - Desktop background

- Font smoothing or visual styles in RDC
- Automatic reconnection if a connection drops

RDC detects connection quality and available bandwidth between itself and the remote desktop computer. It displays the bandwidth with an icon on the connection bar that is similar to a signal strength meter, as described in the following table.

Table 1: Connection quality and bandwidth

Bandwidth	Latency	Icon
10 megabits per second (mbps) and higher	Any value	4 bars
2,000-9,999 kilobits per second (Kbps)	Any value	3 bars
512 kbps - 19,999 Kbps	Any value	2 bars
Less than 512 Kbps	Any value	1 bar
No bandwidth detected or older remote desktop host	Any value	No icon displays

You can also configure the connection settings for RD Gateway.

Plan RDS deployment

For some administrators, planning the installation of Windows roles is a straightforward process because they don't have specific requirements and they don't use many system resources. The Remote Desktop Services (RDS) role planning process is not as simple.

RDS roles include six role services and each RDS deployment should include at least three of them, preferably on separate servers. You should be aware of the functionality that each role service provides. You should also be aware how an RDS deployment uses each role service. You need to know role service requirements and which hardware resources are most critical for each role service. You need to understand the critical role services for a session-based desktop deployment of RDS and how to plan for them. You must also test and evaluate a proof of concept (POC) deployment before you perform an RDS deployment in a production environment.

Assess RDS infrastructure requirements

Before you implement RDS, you must determine your organization's needs and requirements to know whether RDS is the appropriate solution for your needs. You must then decide between session-based or virtual machine (VM)-based desktop deployments. (If necessary, an RDS deployment can include both session-based and VM-based desktop deployments.) You must also evaluate the existing server infrastructure and estimate the required server hardware, network bandwidth, client types and requirements, and connectivity needs for a successful RDS deployment.

Determine your RDS needs

To determine whether RDS is an appropriate solution for your needs, you should assess and analyze the types of users, hardware, and applications in your organization. Areas of consideration include:

- User types. Do you have users in remote locations, single-task users, contractors, or other types of users who would benefit from remote applications or virtual desktops?

- **Hardware.** What client hardware is deployed in your organization currently? Would it be beneficial for some users to move from traditional desktops to thin clients? (*Thin clients* are network devices that process information independently, but rely on servers for applications, data storage, and administration.)
- **Bring Your Own Device.** Do you allow users to bring their own devices into the organization's network? Do users want to use mobile devices to run certain applications?
- **Application compatibility.** Can the applications run in a multiuser environment? If not, will the applications run in a virtual environment?
- **Application performance.** How do the applications perform in a remote or virtual environment? Keep in mind that many applications perform better as RemoteApp programs on RDS because processing takes place on a server.
- **Application support.** Do vendors support the applications in a virtual or multiuser environment? Do vendors provide support to multiple users?
- **Licensing.** Can the applications be licensed for a virtual or multiuser environment?
- **Business benefits.** Are there justifiable business reasons to implement this solution? Potential benefits include cost savings, reduced deployment time, centralized management, and reduced administration costs.
- **Legal requirements.** Does your organization have legal requirements regarding application and data storage? For example, some financial and legal requirements mandate that applications and data remain on-premises. RDS enables users to connect to a standard virtual desktop, while organizational data and applications never leave the datacenter.

Choosing between session-based and VM-based desktop deployments

RDS offers two deployment types:

- **Session-based desktop deployment.** This deployment type enables users to connect to a Remote Desktop Session Host (RD Session Host), and either use a full desktop experience or run remote applications and present them as if they were installed locally.
- **VM-based desktop deployment.** This deployment type provides users with access to a full Windows client operating system (OS)—such as Windows 10—that runs on a VM.

To determine which RDS deployment type is best for your environment, you must consider your users' requirements. For example, do they need to be isolated or have administrative access? You should also consider whether the applications will work properly in a multiuser environment, and whether you can even install and run applications on Windows Server. Remember that a VM-based desktop deployment typically requires a more powerful server infrastructure and more disk storage than a session-based desktop deployment for the same number of users. For some applications, this might be the only viable solution.

Determine server hardware and network resource requirements

After you determine how deploying RDS can benefit your organization, you must consider hardware requirements to support your users. Areas to consider include:

- **Number of users.** How many users will use RDS, and where are they located?

- User types. How many users run CPU-intensive and bandwidth-intensive applications? Will you have to provide more bandwidth and server hardware to support expected usage?
- Connection characteristics. How many concurrent connections do you expect? Can your server and bandwidth resources handle peak usage times?
- Application silos. Will you have to create multiple server collections to support different applications that might not be able to run on the same server?
- Load balancing. How many servers will you need in a collection to spread the load among the servers and provides redundancy?
- High availability. What is the organization's tolerance for downtime? Do you need close to zero downtime, or could your organization tolerate the time it would take to restore from backups?
- Expansion considerations. What are the predicted growth expectations? At what point will new resources need to be brought online?

Determine client requirements

A large organization with multiple locations might have a number of user requirements to consider, such as:

- Languages. Can you install language packs on all of your RDS servers? Organizations with a global presence need to support multiple languages.
- Profile management. How will you store user states? Do users require the same user state when they sign in locally and to an RDS session? Which type of Windows user state virtualization will be used?
- Printing. Will existing printers function properly in a remote desktop environment? Will there be problems finding printer drivers to support existing printers? Is there a budget to replace older printer models?

Determine how clients access RDS

Client devices can connect to RDS in various ways. You will probably need to provide different access methods for different groups of users. Areas to consider include:

- Will you allow users to connect over the internet from remote locations? If so, you'll need to set up a Remote Desktop Gateway (RD Gateway) and obtain certificates.
- How will you manage Secure Sockets Layer (SSL) certificates—by using certificates from non-Microsoft certification authorities (CAs) or by using certificates that an internal CA issues?

Based on assessment results, you should identify RDS role services that are required and which you will deploy. You should also determine the number and hardware configuration of servers that are required, in addition to planning for required storage, connectivity, and firewall configurations.

Plan for RD Session Host

RD Session Host is role service that hosts Windows-based programs or full Windows desktops for RDS clients. This role service is mandatory for every RDS deployment that provides users with session-based desktops with full Windows desktops or with RemoteApp programs. An RD Session Host server accepts incoming RDP requests, and after a client authenticates, it provides a desktop-based or application-based session to the client. An *RD Session Host server* is the central location where remote applications are installed, accessed, and maintained.

An RD Session Host server accepts user connections and runs programs. To use an RD Session Host server, you must consider the number of installed applications and types, resource use, number of connected clients, and the type of user interaction. For example, on one RD Session Host, users might run a simple application that has low resource utilization and rarely runs, such as an old data entry application. On another RD Session Host, users might often run a resource-intensive graphical application that requires higher CPU usage, a considerable amount of random access memory (RAM), intensive disk I/O operations, and causes a lot of network traffic. If the hardware configuration on both of the RD Session Hosts is the same, the second server is considerably more utilized and can accept fewer user connections.

RD Session Host planning focuses on the number of concurrent users and the workload they generate. A server with a particular hardware configuration might support multiple, simultaneous users, or only a few, depending on their usage pattern and the applications that they are running on the RD Session Host. In general, on the same server hardware, you can support more users with session-based desktop deployments than with VM-based desktop deployments.

The main resources that you should consider when estimating RD Session Host utilization are:

- CPU. Each remote application that a user starts runs on an RD Session Host and utilizes CPU resources. In an environment where many users are connected to the same host, CPU and memory are typically the most critical resources.
- Memory. Additional memory must be allocated to an RD Session Host for each user who connects either to a full Windows desktop or runs a RemoteApp program.
- Disk. As user state typically is not stored on an RD Session Host, meaning that disk storage usually is not a critical resource. However, many applications run simultaneously on an RD Session Host, and the disk subsystem should be able to meet their disk I/O needs.
- Network. The network should provide enough bandwidth for connected users and for the applications that they run. For example, a graphically intensive application or a poorly written data entry application can cause a lot of network traffic.
- GPU. Remote applications that are graphically intensive, especially three-dimensional graphics, might require GPU support. Without such support, graphics will render on the server's CPU.

Note: Installing RD Session Host on a VM is fully supported because it integrates with Microsoft Hyper-V.

When estimating the required resources for an RD Session Host, you can use one of the following methods:

- Pilot deployment. This estimation method is a common and a simple approach. You first need to deploy RDS in a test environment and capture its initial performance. After that, you start increasing server load by increasing the number of users and monitoring response times and user feedback. You can find out how many users can connect to an RD Session Host and still have an acceptable user experience based on the number of users and the system response time. Based on the findings, you can estimate the number of servers that are needed for a production environment. This approach is dependable and simple, but it requires initial investments for the pilot deployment.
- Load simulation. This method also uses an initial RDS deployment in a test environment. You need to gather information on applications that users operate and how they interact with the applications. After that, you can use load simulator tools to generate various levels of typical user loads against an RDS deployment. When a load simulator tool runs, you need to monitor server utilization and responsiveness. This method is similar to the previous method, but it uses a load simulation tool to generate user load instead of real users. It also requires an initial investment, and its results depend on the initial estimation of actual user usage.
- Projection based on single-user systems. This method uses data collected from a single-user system, and then extrapolates it to determine expected utilization on an RD Session Host with multiple user

sessions. This method requires detailed knowledge of applications that are used, and is usually not very dependable because a single-user system has a different overhead than a multiuser system.

When planning an RDS deployment, you should consider its importance and how to provide high availability of desktops and RemoteApp programs that run on the RD Session Host. You can provide high availability for an RD Session Host by including multiple RD Session Hosts in each session collection in the RDS deployment.

Plan for RD Connection Broker

During RDS deployment planning, you must designate a server to host the Remote Desktop Connection Broker (RD Connection Broker) role service. This role service is a mandatory component of each RDS deployment, and provides users with access to RemoteApp programs and remote desktop connections, in addition to virtual desktops.

The RD Connection Broker role service also publishes an RDWeb feed. This role service manages all aspects of sessions, including load-balancing connections between multiple RD Session Host servers, and reconnecting user sessions to existing sessions on virtual desktops, remote desktops, and RemoteApp programs.

RD Connection Broker functions include:

- Determining the most appropriate RD Session Host or virtual desktop to send a connection request to, based on a user's identity and the current load on RD Session Host or Remote Desktop Virtualization Host (RD Virtualization Host) servers.
- Storing information about connections to VMs and sessions. This information is stored either in the Windows Internal Database (WID), a local Microsoft SQL server or and Microsoft Azure SQL Database.
- Configuring RDS servers in the same group, or *collection*. You configure settings once—for example, session settings or certificates—and RD Connection Broker applies the settings to servers in the collection.
- Managing VM creation, deletion, and assignments. In VM-based desktop deployments, RD Connection Broker manages VM creation and deletion for managed collections, and also assigns personal virtual desktops to users.
- Gathering RemoteApp information. RD Connection Broker acting as the resource for Remote Desktop Web Access (RD Web Access) to gather RemoteApp information from RD Session Host servers.
- Gathering information on RemoteApp programs and VMs. RemoteApp provides information on which RemoteApp programs and VMs are available through RD Web Access.

When a user initiates a session, RD Connection Broker receives the session request, which queries the database to determine whether there's an existing disconnected session for that user. If so, the user is directed to the disconnected session. If not, RD Connection Broker determines the server in the collection that's best able to manage the new connection based on the load balancing algorithm.

RD Connection Broker is an entry point to an RDS deployment, so it's critical that it's available at all times. RD Connection Broker only directs clients to available RD Session Host servers. If the RD Session Host role service is installed on multiple servers, it becomes highly available. RD Connection Broker uses WID for storing session information. So, if your RDS deployment only has one RD Connection Broker server, it represents a single point of failure. To make RD Connection Broker highly available, you need to add multiple RD Connection Broker servers to an RDS deployment and configure them for high availability. This requires all session information to be moved to either SQL Server data management software or Azure SQL databases. The computer that is running SQL Server should also be clustered to improve availability.

Use an RD Connection Broker server to establish sessions

When a client connects to a session collection, the connection is made between the client and one RD Session Host server in that collection. RD Connection Broker determines which RD Session Host server in the collection should accept the connection and directs the client to that server. The following steps describe the connection process:

1. A user selects the link in the RD Web Access portal to the RDS resource they want to access. This downloads the .rdp file, which contains information about the resource the user wants to connect to, and opens it in the Remote Desktop Connection (RDC) client.
2. RDC initiates the connection with RD Connection Broker.
3. The user authenticates to RD Connection Broker and passes the RDS resource request to which the user wants to connect.
4. RD Connection Broker examines the request to find an available RD Session Host server in the desired collection.
5. If the request matches a session that's already established for the associated user, RD Connection Broker redirects the client to the server in the collection where the session was established. If the user doesn't have an existing session in the collection, the client redirects the client to the server that's most appropriate for the user connection based on the RD Connection Broker load-balancing algorithm—for example, weight factor, fewest connections, and least utilized.
6. The client establishes a session with the RD Session Host server that RD Connection Broker provided.

Plan for RD Web Access

RD Web Access is the RDS role service that provides a single web portal where available remote desktops, RemoteApp programs, and virtual desktops are displayed. You can access RD Web Access from any web browser. When you select an RDS resource, the .rdp file downloads and RDC automatically opens it.

RD Web Access can also publish a list of available RDS resources as an RDWeb feed that can integrate with the **Windows Start** menu. The RD Web Access role service is a mandatory part of each RDS deployment, and it installs the Web server role and Internet Information Services (IIS) as a prerequisite.

RD Web Access benefits include:

- Authorized users can quickly access a list of available RemoteApp programs, remote desktops, and virtual desktops on a webpage from almost anywhere.
- A list of available RDS resources publishes automatically via an RDWeb feed, and it can integrate with the **Start** menu.
- Changes in available RDS resources update automatically on clients that have subscriptions to an RDWeb feed.
- Users can launch the RDC client from the RD Web Access portal, which enables them to connect remotely to the desktop of any computer to which they have Remote Desktop access.
- RD Web Access and RDWeb feeds are personalized and include only RDS resources for which users have permissions.
- Administrators can customize an RD Web Access portal without programming.

Note: RD Web Access only provides a link to RemoteApp launch programs, or to connect to a Remote Desktop session. RD Web Access doesn't proxy client requests. For a user to run an application or to connect to a VM or remote desktop, a client must be able to establish a connection to the target server.

Plan for an SSL certificate

When planning for the RD Web Access role service, you should focus on:

- How to acquire and distribute SSL certificates.
- How to distribute and configure a URL for the RDWeb feed.
- How to provide high availability for RD Web Access.

Note that SSL is required for both the RDWeb feed and the RD Web Access portal.

Important: Installing RDS automatically installs the Microsoft Internet Information Services (IIS) web server role, which generates a self-signed certificate for SSL. This certificate can be used for testing, but clients don't trust it. Clients can connect but will receive a warning that the certificate is not trusted.

You can use Server Manager to edit RDS deployment properties and to configure RD Web Access with an SSL certificate. You can also select an existing SSL certificate or generate a new certificate for RD Web Access. You can use both options to test a deployment in a test lab; however, try to avoid this in a production environment because by default, client computers don't trust an SSL certificate issued in this way. Instead, use an internal CA to issue a RD Web Access SSL certificate. Be aware though that only computers that are domain members by default trust such a certificate. Devices that are not domain members—such as devices used by contractors or home users—need a mechanism to distribute and install the CA certificate in the trusted root CA store on those devices. In addition, for a production environment this might also include devices that you have no management control over. The best option is to use an SSL certificate from a publicly trusted CA.

Plan for high availability

RD Web Access provides links to available RDS resources. If RD Web Access is not available, clients that don't have a local copy of the .rdp file are not able to initiate a remote desktop connection even if the RD Session Host or RD Virtualization Host servers accept connections. To provide high availability for RD Web Access, you need to install multiple servers and add them to the RDS deployment. You should also load-balance client requests between multiple RD Web Access servers, which means that you should configure Domain Name System (DNS) round robin or add servers to a Network Load Balancing (NLB) farm.

Plan for preserving user state

One of the challenges of supporting remote users is how to deal with user state data, specifically what settings and data needs to be stored, and where it is stored. Different technologies can address this challenge, such as roaming user profiles or Folder Redirection. When you deploy RDS, you have an additional option to use User Profile Disks (UPDs).

Roaming user profiles

You can configure a user with two profiles. They use a standard user profile when they sign in locally with their RDS user profile. A second profile is for when the user signs in to an RD Session Host. Creating the two profiles for a user enables them to have two different and independent user states. For roaming users, these profiles should be stored on a fault-tolerant file server. When the user signs out, any profile data can be removed from the RD Session Host server. However, to provide a smooth sign-in experience, roaming user profiles can be cached locally on an RD Session Host server.

Folder redirection

You can use Folder Redirection to change the target location of specific folders in a user's profile from a local hard disk to a network location. For example, the **Documents** folder can redirect from a user's profile to a central network location. Redirected data stores only on a network location, and a local copy is not saved. With **Offline Files** enabled, redirected data synchronizes to a user's computer, typically at sign-in and sign-out. In this way, data is available locally even if network connectivity is temporarily unavailable.

A user can transparently use a local copy of the data, and when connectivity restores, local changes synchronize with the network location. This decreases the amount of data that's stored in the profile and makes data available from any computer to which a user signs in, either locally or by using Remote Desktop. It also ensures that data can be backed up centrally, because it's stored on a network drive.

When planning for Folder Redirection, ensure that:

- The network location is accessible to users who store data in the redirected folders.
- Share-level and file permissions (NTFS file system permissions) are set correctly to allow users access to their redirected folder.
- The network location has enough storage for the redirected folders.

User profile disks

User profile disks are used in RDS either in session-based or VM-based sessions, to isolate user and application data for each user in a separate .vhdx file. This file must be stored on a network location. When a user signs in to an RD Session Host for the first time, their user profile disk is created and mounted into the user's profile path, **C:\Users%username%**. During a user's session, all changes to the profile write to the user's .vhdx file, and when the user signs out, their profile disk is unmounted.

Note: The administrator can limit the maximum size of a .vhdx file and can limit which files or folders are included or excluded from a user profile disks.

User profile disks have the following characteristics:

- You can use user profile disks to store all user data, or you can specifically designate folders to store on a user profile disk.
- You can control which files or folders are included or excluded from user profile disks. Only files and folders from a user profile can be included or excluded from that user's profile disks.
- User profile disk share-level permissions are set up automatically.
- User profile disks are available for both VM-based and session-based connections.
- User profile disks can be combined with Folder Redirection if you want to reduce the amount of data that is stored in a user profile.
- User profile disks can store roaming user profiles.
- When a user signs in to RDS, the user profile disks are mounted into the profile path, **C:\Users%username%**.
- User profile disk locations must be unique for each group (collection) of RDS servers. You cannot configure two collections with the same Universal Naming Convention (UNC) location to store a user profile disk.
- Distributed File System (DFS) is not supported for User Profile disks, meaning you cannot use a DFS share as a location.

- There is no offline caching of a user profile disk. If a profile disk cannot be mounted during sign-in, the user receives a temporary profile for the duration of the session.

Plan for user profile disks

When you plan for user profile disks, you must estimate the user profile size for an average user who connects to an RDS server. When you configure a user profile disk, you can limit its maximum size, but you cannot change that value later. By default, user profile disks are limited to 20 gigabytes (GB), and when a user signs in for the first time, a dynamically expanding .vhdx file is created. You should ensure that there's enough disk space available on the network location to store all user profile disks.

To configure a user profile disk, you must provide the following information:

- The network location to create user profile disks—a UNC path
- The maximum user profile disk size in gigabytes
- Whether to store all user settings and data on user profile disks, or only selected folders and files

Important: All RDS servers in a collection must have Full Control permissions on the user profile disk network share. Those permissions are added automatically when you configure a collection with a user profile disk UNC location.

Infrastructure testing prior to rollout

After you assess RDS infrastructure requirements and familiarize yourself with RDS and its role services, you should perform a POC deployment. A POC deployment is critical for a successful RDS deployment because it enables you to evaluate whether all RDS requirements are met. It does this by performing a load simulation (a test run) in a controlled environment. This simulation uses typical user actions to validate your estimates for capacity, application workloads, and usage patterns.

During testing, you should get answers to the following questions:

- Number of users that can connect, and average response time. Can a POC deployment support the expected number of RD users, and is the response time acceptable? How much are servers being utilized? How long does user sign-in and sign-out take, and is the user experience as expected?
- Application system resources consumption. Does the application consume resources in accordance with documented estimates? If the application uses hardware as expected, the rest of the deployment can continue based on initial estimates. If it is not as expected, you must recalculate capacity requirements to ensure accurate estimates.
- Number of user environments being tested. Are all of the potential user environment scenarios being tested? You should test the application by accessing it in all of the ways a user might use it. If there are access methods that you cannot replicate in a POC environment, they should be implemented in a controlled manner when performing the final deployment.
- Application and hardware performance. Are the applications and hardware running as expected? Is any additional performance tuning required? Do you need to perform any additional application configurations to run as expected in an RDS environment? Also, confirm that hardware performance is within estimated parameters.
- Usage or access performance. Are there any unexpected changes in usage or access? If any part of the presentation virtualization POC deployment does not reflect your production environment, alter the POC deployment so that it is as similar as possible to your final, planned infrastructure.

Tip: Using testing to eliminate errors in a deployment is highly important because problems with a presentation virtualization environment are much easier to resolve during testing than during full deployment.

Best Practices Analyzer for RDS

Windows Server includes a Best Practices Analyzer (BPA) for the Remote Desktop Services server role. BPA for RDS can analyze an RDS environment and check for changes that need to be made for RDS to perform optimally. You can access BPA in Server Manager, or by running the **Invoke-BpaModel** cmdlet in Windows PowerShell.

Access RDS

When internal clients want to connect to a remote desktop or run a RemoteApp program, they can access Remote Desktop Services (RDS) resources from the Remote Desktop Web Access (RD Web Access) portal. Alternatively, they can access RDS resources directly if they already have the .rdp file, which contains connection settings for a remote desktop connection.

In both cases, a Remote Desktop Connection Broker (RD Connection Broker) load-balances client requests and directs them to either a Remote Desktop Virtualization Host (RD Virtualization Host) or a Remote Desktop Session Host (RD Session Host). On an RD Virtualization Host, clients can access virtual desktops that run on virtual machines (VMs), or RemoteApp programs that publish on the VM. Clients can access full remote desktops or published RemoteApp programs on an RD Session Host.

Remote clients have two options to access an RDS deployment. They can contact the RD Web Access portal, where they can get links to internal RDS resources, or they can contact the Remote Desktop Gateway (RD Gateway) if they already have the .rdp file (which contains the address of the internal resource and other connection settings). Remote client communications are encapsulated and sent through RD Gateway to an RD Virtualization Host or an RD Session Host. The user experience is the same for both internal and external clients.

After you've created the collection and published RemoteApp programs, users can connect to a collection and run RemoteApp programs. They have three options to connect:

- Use Remote Desktop Connection (RDC)
- Sign in to the RD Web Access portal
- Subscribe to an RD Web feed and use included links

Note: If users subscribe to the RD Web feed, then links for published RemoteApp programs, collections, and virtual desktops are included on the **Start** menu of their Windows 10–based computers.

Connecting with RDC

A user can connect to a remote desktop if they use a standard RDP client. If the computer to which the user is trying to connect allows a Remote Desktop connection, and if the user is in the Remote Desktop Users group for that computer, the user can sign in. They are then presented with a full desktop and all the remote computer resources, just as if the user were sitting at the local console. The user can also start applications that are published as RemoteApp programs.

Note: When you initiate RDC manually, you connect to a specific RD Session Host server and not to a session collection. We recommend that you use RD Web Access to initiate a connection to RDS, and not use RDC directly.

Connecting to the RD Web Access page

RD Web Access is part of any RDS deployment. It provides a web view of available RDS resources and enables users to initiate RDP connections from the Web Access page.

Tip: The RD Web Access page is available at (link format not an actual link where fully qualified domain name is (FQDN)) `https://FQDNofRDWebAccessServer/rdweb`. So if the servers' fully qualified domain name (FQDN) is `SEA-RDS1.contoso.com`, the URL for the RD Web Access page would be: `https://SEA-RDS1.contoso.com/rdweb`.

The RD Web Access portal automatically uses a self-signed Secure Sockets Layer (SSL) certificate to secure network communication with the RD Web Access site. However, clients receive a security warning because they don't trust self-issued certificates. You should consider replacing the self-signed certificate with a certificate issued by a trusted certification authority (CA).

Before they can review available RDS resources, users must sign in to the RD Web Access portal. If SSO is configured for RDS and the device is domain-joined, users will be signed in automatically using their Windows sign in information. However, they can only review a list of available RDS resources to which they have permissions. Users initiate a connection to an RDS resource by selecting a link to the RD Web Access page, which downloads the .rdp file and initiates an RDP connection in RDC. Follow the instructions included in the Tip section to obtain the proper link.

Note: Users can connect to an RD Web Access portal by using Microsoft Edge, Google Chrome, Mozilla Firefox, and Safari.

Integrating published RemoteApp programs into the Start menu

If you want to integrate access to available RemoteApp programs with the **Start** menu or to use file-type associations (also called *document invocation*) for starting RemoteApp programs, you start RemoteApp programs on a client that has RemoteApp and Remote Desktop Connections installed. If you need to add connections for multiple users simultaneously, you can add RemoteApp and Remote Desktop Connection manually by using Control Panel or Group Policy. To add RemoteApp and Remote Desktop Connections, you need to provide either the RD Web feed URL or your email address. The default URL for an RD Web feed is `https://FQDNofRDWebAccessServer/Rdweb/webfeed.aspx`.

If you want to add RemoteApp and Remote Desktop Connections, the RD Web Access portal must use an SSL certificate that the client trusts. By default, RD Web Access uses a self-issued certificate, which means that you must configure it with a different certificate before you can add RemoteApp and Remote Desktop Connections on clients.

What is RemoteApp and Remote Desktop Connection

The RemoteApp and Remote Desktop Connection feature enables you to add available RDS resources on the client computer's **Start** menu. Users can start a RemoteApp program or remote desktop session from the **Start** menu without opening the RD Web Access portal first. A list of available RDS resources is periodically automatically updated, but you can also manually refresh it. You can add a connection from any Windows 8.1 or later device, regardless of whether it's a domain member.

Important: Before you can add RemoteApp and Remote Desktop Connection, RD Web Access must be configured with a trusted SSL certificate.

You can add RemoteApp and Remote Desktop either manually or by using Group Policy. If you use Group Policy, you can add a connection simultaneously to multiple users, but be aware that you must specify the default connection URL, and the client computers must be running Windows 8.1 or later. Users can

also add a connection manually via an email address or with a connection URL. In that case, users must also enter the credentials that are used for accessing the RD Web feed. After the connection is added, you can access its properties, such as its name, connection URL, when the connection was created, and when it was most recently updated. You can also update a connection manually.

If you want to add RemoteApp and Remote Desktop by using email addresses, you must add a text (TXT) resource record to the Domain Name System (DNS) server. The TXT resource record is used to map an email suffix to the RemoteApp and Desktop Connections URL address. Its name must use “_msradc” as its name, and the text box must contain the RemoteApp and Desktop Connections URL address.

You can use the Work Resources (RADIC) section in the App **Start** menu to review RDS resources that were added by RemoteApp and Desktop Connections. You can pin RDS resources to **Start**, but you cannot pin them to the taskbar. When you use search, it will also find RDS resources that match the search criteria.

The RemoteApp and Remote Desktop Connection feature offers several benefits:

- You can start RemoteApp programs from the **Start** menu, just like locally installed apps.
- Only RDS resources for which you have permissions are added.
- It adds all available RDS resources, including collections (Remote Desktop), virtual desktops, and RemoteApp programs.
- The list of available RDS resources refreshes automatically.
- File type association (document invocation) works for RemoteApp programs added by RemoteApp and Remote Desktop.
- Search works with RDS resources, just like with locally installed apps.
- You can add RemoteApp and Remote Desktop Connection regardless of a client computer's domain membership.
- You can add RemoteApp and Remote Desktop Connection to many users simultaneously by using Group Policy.

Note: RemoteApp and Remote Desktop Connection use a scheduled task to update connections. The default frequency is once per day at 12:00 A.M., but you can customize the Update connections task in the Task Scheduler Library in **Microsoft\Windows\RemoteApp and Desktop ConnectionsUpdate\EmailAddressOrConnectionURL**.

Overview of Remote Desktop Gateway

Remote Desktop Services (RDS) can provide a standardized environment for users who connect to an internal network. However, many users need to access their environments when they're not connected to an internal network, for example, when working from home or traveling on business. Users can connect to RDS resources from a public network after they established a virtual private network (VPN) connection to the organization's network. But with an RD Gateway server in place, users can securely connect to an organization's RDS resources even without a VPN. Remote Desktop Gateway (RD Gateway) is an optional RDS role service that you can install it on any server, regardless of its domain membership.

Why is remote access important for RDS

Most of today's computer and device users are extremely mobile. They typically have multiple devices that enable them to access their work environment from almost anywhere. Devices are often not connected to an organization's network, sometimes because the organization doesn't own them, or because they have a third-party operating system (OS) that isn't compatible with becoming a domain member.

However, users are familiar with their personal devices and they want to use them for work. They want to have a consistent user experience with the same applications and data available on all their devices, from anywhere, regardless of whether they're connected to the organization's network.

RDS provides a consistent user experience for users who connect to RDS session-based desktop deployments or virtual machine (VM)-based desktop deployments. RDS uses Remote Desktop Protocol (RDP), and it's available to users who are connected to an organization's internal network. However, users need to access their work environment at all times, even when they're not connected to the internal network. Therefore, organizations have to extend RDS availability to a public network, such as the internet, and have a solution to meet the following goals:

- Provide secure connectivity from a public network
- Use a standard protocol to provide remote access
- Not require firewall reconfiguration
- Control who can connect to the internal network
- Control which internal resources can be accessed, and by whom
- Control features that are available to users on a public network
- Monitor and manage established sessions
- Provide a highly available solution
- Require additional authentication, such as multi-factor authentication (MFA)

Methods for securing remote access to RDS

Users connect to RDS by using an RDP client. RDP includes built-in encryption, but most organizations don't want their RDS deployment to be available from a public network, so they block RDP traffic on an external firewall. In the past, users who were outside of an organization's network had to first establish a VPN connection, and then use Remote Desktop Connection (RDC) for establishing an RDP connection over the VPN to an RDS deployment.

A VPN provides an additional layer of security by encrypting network traffic between the client and the VPN server. A Microsoft VPN client is already included in the Windows operating system (OS), but a VPN connection must be configured before you can use it. Establishing and disconnecting a VPN connection is a manual process unless you're using Always On VPN or DirectAccess. The VPN client supports several VPN protocols, and some of them require additional firewall configurations. You should also be aware that some mobile devices don't include VPN clients.

RD Gateway enables authorized remote users to connect to resources on an internal network from any internet-connected device that can run the RDC client. RD Gateway uses RDP over HTTPS to establish a secure, encrypted connection between remote users on the internet and internal network resources. Internal RDS resources typically are on an internal network behind a firewall, and can also be behind a network address translation (NAT) device. RD Gateway tunnels all RDP traffic over HTTPS to provide a secure, encrypted connection. All traffic between a user's client computer and RD Gateway is encrypted while in transit over the internet.

When a user establishes connections from a public network to the destination RDS host on an internal network, data is sent through an external firewall to RD Gateway. RD Gateway decrypts HTTPS and contacts the domain controller to authenticate the user. RD Gateway also contacts the server that's running Network Policy Server (NPS) to verify if the user can cross the RD Gateway and contact the RDS host. If the user successfully validates and the connection is allowed, RD Gateway passes the decrypted

RDP traffic to the destination RDS host and establishes a security-enhanced connection between the user sending the data and the destination RDP host.

RD Gateway eliminates the need to configure VPN connections. This enables remote users to connect to an organization's network through the internet while providing a comprehensive security-configuration model.

RD Gateway adds the following benefits to an RDS deployment:

- Enables you to control access to specific internal network resources.
- Provides a secure and flexible connection to internal Remote Desktop resources from a public network. You can control which internal resources can be accessed and who can access them.
- Enables remote users to connect to internal network resources that are hosted behind firewalls on an internal network and across NAT devices.
- Enables you to configure authorization policies to define conditions for remote users to connect to internal network resources by using Remote Desktop Gateway Manager (RD Gateway Manager).
- Enables you to configure RD Gateway servers and RDC clients to use Network Access Protection (NAP) to enhance security.
- Provides tools to help you monitor RD Gateway connection status, health, and events. By using RD Gateway Manager, you can specify events that you want to monitor for auditing purposes such as unsuccessful connection attempts to the RD Gateway server.
- Integrates with additional security providers such as Microsoft Azure Multi-Factor Authentication or third-party authentication providers.

Remote Desktop Server Manager manages RDS deployments. RD Gateway might or might not be added to an RDS deployment, but in both cases, it's managed by using RD Gateway Manager.

Network configuration for RD Gateway

RD Gateway provides secure access to RDS resources on internal networks for clients connected to a public network such as the internet. RD Gateway must have connectivity to both an internal and public network to be able to proxy and inspect RDP traffic.

Location of RD Gateway

When planning your RD Gateway deployment, the RD Gateway network placement is one of the biggest considerations. In the simplest deployment, RD Gateway has a dedicated server on an internal network with two network interfaces: one interface is exposed to the internet, and the other interface is connected to the internal local area network (LAN). You can configure Domain Name System (DNS) resource records to allow name resolution of the server's fully qualified domain name (FQDN) from the internet and from the LAN. Though simple to set up, this deployment does not follow best security practices, and is not recommended.

A more secure deployment is to put RD Gateway in a perimeter network. In this design, RD Gateway listens for HTTPS traffic on the internet-facing interface, and for RDP traffic on the LAN. The drawbacks to this design are that RD Gateway can't communicate with Active Directory Domain Services (AD DS), and it can't be a domain member. This means that users will have to supply two sets of credentials: one for the local RD Gateway server, and one for the domain. A variation on this design is to have the RD Gateway server as a domain member to enable single sign-on (SSO), but this will require a large combination of ports to be open on the internal firewall to allow domain communications. The drawbacks to this second

configuration are the design complexity and diminished security from having multiple ports open to the internal LAN.

The most acceptable option is to place a reverse proxy server in the perimeter network to manage internet communications. In this design, the reverse proxy server inspects inbound internet traffic and then sends it to the internal RD Gateway server. This enables an RD Gateway that's a domain member in the internal LAN to authenticate and further proxy user sessions by using internal domain credentials. A reverse proxy server can also perform Secure Sockets Layer (SSL) bridging, which allows an encrypted request that comes in to the proxy to be decrypted, inspected, re-encrypted, and forwarded to the destination server. A reverse proxy server can be Web Application Proxy or another product that supports RD Gateway.

Installing and configuring RD Gateway

RD Gateway is an optional RDS role service. You can install it on a server that's a domain member and make it part of the RDS deployment, or you can install it on a separate server, regardless of its domain membership. If an RD Gateway server is not a domain member, you can't use AD DS groups to control who can use it. If an RD Gateway is not part of an RDS deployment, you must manually update RDS deployment properties so that links to RDS resources on the RD Web Access portal include information about RD Gateway.

Installing RD Gateway is a straightforward process; you only need to specify an SSL certificate common name. All client communication with RD Gateway is encrypted, and by default, RD Gateway uses a self-signed SSL certificate. Because clients don't trust a self-signed certificate, you should replace it with a certificate from a trusted certification authority (CA). You do this in Server Manager, where you can edit RDS deployment properties and configure certificates, including RD Gateway certificates.

You manage RD Gateway in RD Gateway Manager. This is different from most RDS role services, which you manage in Server Manager. RD Gateway can be installed and used separately, and not necessarily as part of an RDS deployment.

You can manage the following settings in RD Gateway Manager:

- **Maximum connections.** You can control how many simultaneous client connections can go through an RD Gateway server. You can allow as many connections as the server hardware supports or only a certain number of connections, or you can disable any new connection.
- **SSL certificate.** If RD Gateway is part of an RDS deployment, you should use RDS deployment properties in Server Manager to configure its certificate. If it's not part of an RDS deployment, you can either create a self-signed certificate, import a certificate, or select one of the existing server certificates for RD Gateway.
- **Transport settings.** By default, RD Gateway receives incoming HTTPS traffic over port 443 and connects to RDS resources over User Datagram Protocol (UDP) port 3391. You can modify IP addresses and ports for the HTTP and UDP transports that RD Gateway uses. If you modify the settings, RD Gateway's listener rules in Windows Defender Firewall are modified, all active connections to RD Gateway disconnect, and the RD Gateway service restarts.
- **RD CAP store.** RD Gateway uses Remote Desktop Connection authorization policies (RD CAPs) to control who can connect to RD Gateway. By default, RD CAPs are stored locally, but you can specify to store them on a central server that runs NPS.
- **Messaging.** You can configure RD Gateway with system messages that will be sent to all clients that use RD Gateway. You can also enable a sign in message that clients must accept before they can establish new sessions through RD Gateway. Finally, you can require that only clients that support RD Gateway messages can use an RD Gateway server.

- **SSL bridging.** If clients connect directly to RD Gateway, you don't need to enable SSL bridging. However, if clients connect to other devices and those devices connect to RD Gateway, you should enable SSL bridging. In this case, you must enable SSL termination on those devices, and select one of two SSL bridging modes on RD Gateway: HTTPS-HTTPS bridging, or HTTPS-HTTP bridging. If you enable SSL bridging, the Internet Information Services (IIS) server application pool is recycled, all active connections to RD Gateway disconnect, and the RD Gateway service restarts.
- **Auditing.** You can select RD Gateway events that should be audited. Seven different event types are available, and by default, all of them are audited.
- **Server farm.** In an environment with multiple RD Gateway servers, you should add them to a server farm. This ensures that client collections distribute among all available servers in the farm.

You can also use RD Gateway Manager to manage RD Gateway authorization policies, monitor active connections through RD Gateway, manage local computer groups, and import or export policy and RD Gateway configuration settings.

RDS licensing

If you want to use Remote Desktop Services (RDS), you must properly license all clients that connect to RDS. Every client must first have a license for the locally installed operating system (OS). Every client must also be licensed for using Windows Server Remote Desktop Services CAL (Client Access License), assuming they are domain members. Finally, you must obtain an RDS CAL, which allows a client to connect to a session-based desktop deployment of RDS. If you want to connect to a virtual machine (VM)-based desktop deployment of RDS to use a virtual desktop from a device that is not covered by a Microsoft Software Assurance (SA) agreement, you'll need to license those devices with Windows Virtual Desktop Access (Windows VDA) to access a virtual desktop. Be aware, though, that you must license applications that you use on RDS deployments separately, and they aren't included in an RDS CAL.

Whenever a client attempts to connect to an RDS deployment, the server that accepts the connection determines if an RDS CAL is needed. If one is required, then the server requests the RDS CAL on behalf of the client that is attempting the connection. If an appropriate RDS CAL is available, it issues to the client, enabling the client to connect to RDS.

Note: After installing RDS, you have an initial grace period of 120 days before you must install the valid CAL. This grace period begins after the Remote Desktop Session Host (RD Session Host) accepts the first client connection. If that grace period expires and you have not installed valid licenses, clients will not be able to sign in to the RD Session Host.

Remote Desktop Licensing (RD Licensing) manages the RDS CALs that are required for each device or user to connect to an RD Session Host server. You use RD Licensing to install, issue, and track the availability of RDS CALs on an RD Licensing server. At least one RD Licensing server must be deployed in the environment. The role service can install on any server, but for large deployments, the role service should not be installed on an RD Session Host server.

Tip: RDS supports two concurrent connections to administer a server remotely. You don't need an RD Licensing server and RDS CALs for these connections.

RD Licensing modes

RD Licensing modes determine the type of RDS CALs that an RD Session Host server requests from an RD Licensing server on behalf of a client that is connecting to an RD Session Host server. There are two licensing modes:

- **Per User.** This licensing mode gives one user the right to access any RD Session Host server in an RDS deployment from an unlimited number of client computers or devices. You should use RDS Per User CALs when the same user connects to RDS from many devices.
- **Per Device.** This licensing mode gives any user the right to connect to any RD Session Host server in an RDS deployment from a specific device. When a client connects to an RD Session Host server for the first time, a temporary license is issued by default. When the client computer or device connects to an RD Session Host server for the second time, if the license server is activated and enough RDS Per Device CALs are available, the license server issues the client computer or device a permanent RDS Per Device CAL. You should consider RDS Per Device CALs when multiple users use the same device for connecting to RDS, for example, a point-of-sale device that is used by different clerks.

The RD Licensing mode can be set in Server Manager by configuring the RDS deployment properties. In addition, you can set the RD Licensing mode in Group Policy by configuring the **Set the Remote Desktop licensing mode** policy setting or by using Windows PowerShell.

You should consider an RDS External Connector License if multiple external users who aren't employees of your organization need to access RDS. An RDS External Connector License allows an unlimited number of non-employees to connect to a specific RD Session Host. If you have multiple RD Session Host servers, you need multiple External Connector Licenses, in addition to any required Windows Server External Connector Licenses.

Additional reading: For additional information on how to license your RDS deployment, refer to **License your RDS deployment with client access licenses (CALs)**¹.

Windows VDA

Windows VDA is a device-based subscription that gives you the right to connect to a virtual desktop from a client device that is not covered by an SA agreement with Microsoft. If the client device is running the Windows OS and SA doesn't cover it, Windows VDA is not needed for that device.

RDS in Azure

As you have learned so far in this module, Remote Desktop Service (RDS) is a cost effective and high manageability solution for hosting your applications and Windows desktops. Previously, the RDS environment was installed in an organization's on-premises datacenter. Using Microsoft Azure, you can now also run your RDS environment in the cloud. To run RDS in Azure, you can either install RDS on virtual machines (VMs) in Azure, or run Windows Virtual Desktop (WVD).

Create an RDS deployment on Azure IaaS

If you choose this option, you might be able to migrate your existing servers running your RDS components to Azure. Perhaps your RDS servers are running as Microsoft Hyper-V VMs. Another option is to use

¹ <https://aka.ms/rds-client-access-license>

an Azure Marketplace offering which will install a complete RDS environment for you in Azure. The basic RDS environment will consist of six Azure virtual machines with the following components:

- 1 VM with Remote Desktop Connection Broker (RD Connection Broker) and Remote Desktop Licensing (RD Licensing)
- 1 VM with Remote Desktop Gateway (RD Gateway) and Remote Desktop Web Access (RD Web Access)
- 1 VM with a domain controller
- 1 VM with a file server (for user profile disks)
- 2 VMs with a Remote Desktop Session Host (RD Session Host)

Using the Azure Marketplace offering is an easier and faster way of getting an RDS environment up and running in Azure. Besides creating all the VMs that will run the RDS components, everything else is configured for you in Azure as well, including resource groups, virtual networks, and load balancers. However, because the offering creates a test domain, the choices you have at deployment time are somewhat limited, so this option is best suited for creating an RDS test environment. Note that you can resize virtual machines and make other changes to the environment after the deployment.

You can use the following high-level steps to create an RDS deployment using the Azure Marketplace:

1. Sign in to the Azure portal.
2. In the **Azure Marketplace** blade, search for **RDS**.
3. Select **Remote Desktop Services (RDS) Deployment**, and then select **Create**.
4. After you have been guided through the process, connect to your Azure RDS environment from a client to test the functionality.

You can also choose to use an Azure Resource Manager (ARM) quickstart template to deploy your RDS environment. Using a quickstart template gives more control over the deployment and is the recommended approach if you already have an existing Active Directory. You can then join the Azure VMs to your on-premises Active Directory. By using the Azure quickstart templates for RDS you can customize all aspects of the deployment, including:

- The Active Directory to join
- Virtual network configuration
- Using custom images for your RDS VMs

Additional reading: For additional information on Azure quickstart templates for RDS, go to **Azure Quickstart Templates**².

You could also create your own ARM template to automate RDS deployment, or configure all required components manually, such as VMs, virtual networks, and load balancers.

Introduction to WVD

WVD is a platform-as-a-service (PaaS) offering running in Azure. It enables you to create and run Windows 10 virtual desktops or RemoteApp programs in Azure. Because this is a PaaS offering you don't have to install the RDS roles such as RD Web, RD Connection Broker, RD Gateway, and RD Licensing. All you need to do is configure the environment the way you want it and Azure takes care of the rest.

² <https://aka.ms/azure-qs-templates>

To start using WVD, you must have:

- An Azure subscription.
- An on-premises Active Directory that synchronizes with Azure Active Directory (Azure AD) using Azure Active Directory Connect (Azure AD Connect). Instead, you can use Azure Active Directory Domain Services (Azure AD Domain Services).
- A virtual network in Azure that's connected to your on-premises network via either a site-to-site virtual private network (VPN), or ExpressRoute.

In addition, VMs that are either running the RemoteApp programs or are being used as Virtual Desktop Infrastructure (VDI) machines must be joined to an Active Directory domain.

Another important factor to consider when working with WVD is cost. You must pay for the Azure infrastructure, and you must have the appropriate license.

The following licenses enable you to use WVD:

- Microsoft 365 E3 or E5
- Microsoft 365 A3 or A5
- Microsoft 365 A5
- Microsoft 365 F3
- Microsoft 365 Business Premium
- Windows 10 Enterprise E3 or E5
- Windows 10 Enterprise A3 or A5

All of these licenses include the cost of the operating system (OS) and the WVD PaaS management service in Azure. This means that you won't need to pay for server licenses or RDS CALs to run WVD.

The VM resources in Azure that users connect to and use incur a cost as well. You must pay for the VMs, user profile storage, and egress bandwidth. To help estimate these costs you use the Azure Pricing Calculator.

Additional reading: For additional information on Azure Pricing Calculator, refer to **Pricing calculator**³.

Users can either use personal desktops that run Windows 10 Enterprise, or *pooled desktops*, which are VMs share by more than one user and running Windows 10 Enterprise multi-session (formerly known as *Windows 10 Enterprise for Virtual Desktops*).

Windows 10 enterprise multi-session is a special version of Windows 10 that allows more than one concurrent RDP connection. One of the advantages to using WVD is the ability to use a client version of the OS to provide session-based desktops to users. When using RDS session-based desktops, the user experience is the same as a Windows 10 desktop even though they're connected to and running applications on a server OS. Some applications might not function properly when running on a server OS. By using WVD, your users will always get an up-to-date Windows 10 client platform with support for most applications.

You can access WVD either through any HTML5-supported browser or by using the Windows Desktop client. This client must be downloaded and installed and is available for the following platforms:

- Windows Desktop
- Web
- macOS

³ https://aka.ms/pricing_calculator

- iOS
- Android

Important: You cannot use the built-in Remote Desktop Connection (MSTSC) in the OS to access WVD resources such as desktops or RemoteApp programs.

Test your knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

You're deploying a session-based Remote Desktop Services (RDS) deployment on your company's network. Users will only be connecting from your on-premises network. Which roles would you need to deploy? Choose all that apply.

- Remote Desktop Connection Broker (RD Connection Broker)
- Remote Desktop Gateway (RD Gateway)
- Remote Desktop Web Access (RD Web Access)
- Remote Desktop Virtualization Host (RD Virtualization Host)
- Remote Desktop Session Host (RD Session Host)

Question 2

How is RDS different from the Remote Desktop feature?

Question 3

Is RD Gateway required if you want to enable Internet clients to connect to your internal RDS resources?

Configuring a session-based desktop deployment

Lesson overview

Performing a session-based desktop deployment of Remote Desktop Services (RDS) is a straightforward process. It's important to remember that you deploy RDS differently than any other Windows Server roles. Instead of the role-based or feature-based installation that you use with other roles, you must select RDS installation in Server Manager to deploy RDS. A wizard then guides you through the process. You can only install RDS role services on servers that Server Manager is aware of—if servers are not already added to Server Manager, you should add them before you deploy RDS.

Collections are an essential element of each RDS deployment because they can provide high availability and load balancing for the servers in a collection. In this lesson, you'll learn more about collections, how to create them, and how to configure their properties.

An RDS infrastructure must be constantly available, and you'll be introduced to the various methods that you can use to provide high availability for it. You can make most role services if you add multiple servers with those role services installed, to an RDS deployment or to a collection.

You'll also learn about RemoteApp programs in RDS deployments. Instead of installing applications locally on each client, you install applications on Remote Desktop Session Host (RD Session Host) servers, and then publish them as RemoteApp programs. This enables users to run the applications even when they're not installed locally on client computers. A user can start RemoteApp programs from both a Remote Desktop Web Access (RD Web Access) portal, and from the **Start** menu.

Lesson objectives

After completing this lesson, you'll be able to:

- Describe the session-based desktop deployment process.
- Install RDS.
- Describe a collection.
- Explain how to configure session collection settings.
- Create and configure a session collection.
- Understand high availability options for RDS.
- Describe RemoteApp.

Overview of the session-based desktop deployment process

Remote Desktop Services (RDS) includes multiple role services. If you use Server Manager for RDS deployment, and if you use role-based or feature-based installations, you can install individual RDS role services. However, if you install an RDS role service in this way, you can't manage it. If you want to manage RDS, a deployment must have at least three role services: Remote Desktop Connection Broker (RD Connection Broker), Remote Desktop Web Access (RD Web Access), and either Remote Desktop Session Host (RD Session Host) or Remote Desktop Virtualization Host (RD Virtualization Host).

Note: Individual RDS role services can't be managed if they're not part of an RDS deployment.

The preferred method for installing RDS is to use Server Manager and select RDS installation. This way, you install all required RDS role services at once, and you can then manage the RDS deployment. You should use Server Manager to add all the servers you plan to install RDS role services on. If you don't add them, you won't be able to select them in the install process.

The **Add Roles and Features Wizard** helps you through the install process. It will prompt you to define an RDS desktop deployment scenario, which can be either Virtual Machine based (VM-based) or session-based. Depending on what you select, you are able to install RD Virtualization Host or RD Session Host; RD Connection Broker and RD Web Access are always installed. After the RDS installation finishes, you can add additional role services to the RDS deployment and start to configure the deployment.

Note: The Remote Desktop Gateway (RD Gateway), RD Web Access and RD Session Host role services are not supported on the Server Core edition of Windows Server. However, the RD Virtualization Host, RD Connection Broker, and Remote Desktop Licensing (RD Licensing) role services are supported.

After you plan your RDS deployment, you must complete a number of tasks to configure a session-based desktop deployment scenario. The process is described in the following high-level steps:

1. Add servers to Server Manager
 - To install role services on remote servers, you must first add the servers to Server Manager. You can install RDS role services only on the servers that Server Manager is aware of, which means that they are in Server Manager's management scope.
2. Install Remote Desktop Services
 - Use the **Add Roles and Features Wizard** to select the RDS installation option. When you use this option, the wizard installs all RDS role services required to manage an RDS deployment.
3. Select either RDS session-based desktop QuickStart or RDS session-based desktop standard deployment.
 - With standard deployment, you can deploy RDS across multiple servers and select which servers have specific role services installed. QuickStart installs all of the required role services on a single server and performs basic initial configuration by creating a collection and then publishing RemoteApp programs.
4. Choose either a session-based desktop deployment or a VM-based desktop deployment.
 - Session-based desktop deployment adds an RD Session Host to a deployment. If you want to provide users with virtual desktops, you need to select VM-based desktop deployment.
5. Choose the servers that have installed RDS role services
 - From the pool of managed servers, select servers that have the RD Connection Broker, RD Web Access, and the RD Session Host role services. Each role service must be installed on at least one server. Multiple role services can install on the same server, and you can install the same role service on multiple servers.

During deployment, the servers on which you installed the RD Session Host role restart. After the installation, you can perform initial RDS deployment configuration. You can also add additional servers to the deployment. At the very least, you should add RD Licensing because you can't connect to an RD Session Host without valid RDS client access licenses (RDS CALs) after the initial grace period of 120 days expires.

You should also consider installing multiple instances of the RDS role services for high availability. You can install session-based desktop deployments of RDS by using the **New-SessionDeployment** cmdlet in the Windows PowerShell command-line interface. This cmdlet performs all of the required configurations, including the system restart.

Demonstration: Install RDS

In this demonstration you will learn how to install Remote Desktop Services (RDS).

Preparation Steps

For this demonstration, you will use the following virtual machines (VMs):

- **WS-011T00A-SEA-DC1**
- **WS-011T00A-SEA-RDS1**
- **WS-011T00A-SEA-CL1**

Sign in to **WS-011T00A-SEA-RDS1** by using the following credentials:

User Name: Contoso\Administrator

Password: Pa55w.rd

Sign in to **WS-011T00A-SEA-CL1** by using the following credentials:

- User name: **jane**
- Password: **Pa55w.rd**
- Domain: **Contoso**

After completing the demonstration, leave all the virtual machines running as they will be used in a later demonstration.

Demonstration steps

Install RDS using Server Manager

1. On **SEA-RDS1**, select **Start**, and then select the **Server Manager** tile.
2. In Server Manager, select **Manage**, and then select **Add Roles and Features**.
3. In the **Add Roles and Features Wizard**, on the **Before you begin** page, select **Next**.
4. On the **Select installation type** page, select **Remote Desktop Services installation**, and then select **Next**.
5. On the **Select deployment type** page, verify that **Standard deployment** is selected, and then select **Next**.

Note: Even though you could have selected the **Quick Start** deployment option and have all three required RDS role services installed on **SEA-RDS1**, you selected the **Standard deployment** option to demonstrate how to select different servers for the RDS role services. Furthermore, the **Quick Start** deployment option will create a collection named **QuickSessionCollection** and publish the RemoteApp programs **Calculator**, **Paint**, and **WordPad**.

6. On the Select deployment scenario page, select **Session-based desktop deployment**, and then select **Next**.
7. On the **Review role services** page, review the description of the role services, and then select **Next**.
8. On the **Specify RD Connection Broker server** page, in the **Server Pool** section, select **SEA-RDS1 . Contoso.com**. Add the computer to the **Selected** section by selecting the Right arrow key, and then select **Next**.

9. On the **Specify RD Web Access server** page, in the **Server Pool** section, select `SEA-RDS1.Contoso.com`. Add the computer to the **Selected** section by selecting the Right arrow, and then select **Next**.
10. On the **Specify RD Session Host servers** page, in the **Server Pool** section, select `SEA-RDS1.Contoso.com`. Add the computers to the **Selected** section by selecting the Right arrow, and then select **Next**.
11. On the **Confirm selections** page, select **Cancel**.

Install RDS using Windows PowerShell

Note: You'll now learn how to install RDS using Windows PowerShell. The previous steps were included to demonstrate how to install RDS using Server Manager.

1. Switch to **SEA-DC1**, and at the Windows PowerShell command prompt enter the following command, and then select Enter:

```
powershell
```
2. Enter the following command, and then select Enter:

```
$SVR="SEA-RDS1.contoso.com"
```
3. Enter the following command, and then select Enter:

```
New-RDSessionDeployment -ConnectionBroker $SVR -WebAccessServer $SVR -SessionHost $SVR
```
4. Wait for the installation to complete and for **SEA-RDS1** to restart automatically. This will take approximately 5 minutes.
5. Switch to **SEA-RDS1**, and sign in as **Contoso\Administrator** with the password **Pa55w.rd**
6. Select **Start**, and then select the **Server Manager** tile.
7. Wait for **Server Manager** to refresh.
8. In **Server Manager**, in the navigation pane, select **Remote Desktop Services**, and then minimize **Server Manager**.

After completing the demonstration, leave all the virtual machines running as you will be using them in their current state in a later demonstration.

What is a collection?

Remote Desktop Services (RDS) deployments support two collection types: Session collections, and virtual desktop collections.

Collections are used as groupings of either Remote Desktop Session Host (RD Session Host) servers or virtual machines (VMs) that are used as virtual desktops. Collection members should be configured identically. For example, you should install the same applications on all RD Session Host servers that are members of the same collection.

Collections simplify the administration process by enabling you to manage all collection members as a unit instead of managing each individually. For example, after you configure a collection with session settings, those settings automatically apply to all the servers in the collection.

If you add an additional server to a collection, session settings also automatically apply to the added server. You can create a collection if you use Server Manager, or you can use the Windows PowerShell cmdlets **New-RDSessionCollection** and **New-RDVirtualDesktopCollection**. You can configure additional collection settings after the initial collection properties are defined and the collections are created.

When a session collection has multiple RD Session Host servers, a collection can provide high availability. Multiple servers in a collection are provisioned to offer the same resources. For example, if one server in a collection fails, Remote Desktop Connection Broker (RD Connection Broker) no longer directs session requests to that server. Instead, RD Connection Broker distributes session requests among the remaining servers in the collection.

Session collections

Before any remote applications or remote desktops can be accessed, you must create a session collection. You can begin creating collections after RDS is installed. Session collections enable you to provide separate configurations for remote desktop connections or for groups of RemoteApp programs; the same collection can't provide both. An RD Session Host can be a member of one collection only, and you can add it or remove it from a collection at any time.

The **Create Collection Wizard** guides you through the steps for creating a session collection. When running the wizard, you must provide the following information:

- The friendly name of the collection
- The RD Session Hosts that will be added to the collection
- Security groups that will be associated with the collection
- Whether users that connect to servers in the session collection will use user profile disks

Keep in mind that Group members will be allowed to access the remote desktop or run RemoteApp programs on servers in the collection. By default, Domain Users is associated with a new session collection.

After you have created a session collection, you can modify the properties that are specific to that collection. Session collections can provide remote desktops or have published RemoteApp programs, but not both.

You can also create a session collection by using the Windows PowerShell cmdlet **New-RDSessionCollection**. For example, to create a session collection named **RemoteApps** on a server named `SEA-RDS1.contoso.com`, you would run the following command:

```
New-RDSessionCollection -CollectionName RemoteApps -SessionHost SEA-RDS1.  
Adatum.com -CollectionDescription "Marketing remote apps" -ConnectionBroker  
SEA-RDS1.contoso.com
```

Virtual desktop collections

The virtual desktop collection contains VMs that are used as users' virtual desktops. VMs for a collection can be created automatically or manually, depending on whether a collection is managed or unmanaged. Before you can create a managed virtual desktop collection, you must create a VM that will be the template for the collection, and generalize it by using the System Preparation (Sysprep) tool.

VMs in virtual desktop collections run Windows 10. There are two types of virtual desktop collections: pooled, and personal. The virtual desktop collection type determines whether a user can connect to a different virtual desktop from the collection each time, or if they are assigned a personal desktop and always connects to the same VM in the collection.

You can use the **Create Collection Wizard** to guide you through the process of creating a virtual desktop collection. However, you will need to provide more information when you create a virtual desktop collection than when you create a session collection. It can also take considerable time and resources to

create managed virtual desktop collections because VMs are created during the process. Creating an unmanaged virtual desktop collection is faster, as VMs must already exist to add them to a collection.

Note: You can also create a virtual desktop collection by using the **New-RDVirtualDesktopCollection** cmdlet.

Configure session collection settings

When you create a session collection, only a limited set of configuration options are available. This is true regardless of whether you are creating a collection by using Server Manager or Windows PowerShell. When a session collection is created, you can modify its initial settings, configure additional settings, and perform additional tasks. For example, you can publish RemoteApp programs, add, or remove Remote Desktop Session Hosts (RD Session Hosts) from a collection, or disconnect existing connections to servers in a collection.

You can modify the following session collection properties:

- **General** page
 - On the **General** page, you can edit the collection name and description, and you can determine whether to present a collection in Remote Desktop Web Access (RD Web Access). The default setting is to display collections in RD Web Access.
- **User Groups**
 - You can specify which Active Directory Domain Services (AD DS) security groups are allowed to connect to RD Session Hosts in a collection and access any RemoteApp programs in the collection. Be aware that for a running RemoteApp program, users must also have access to the RemoteApp program.
- **Session**
 - On the **Session** page, you can configure the following session settings:
 - How long after a session is disconnected to end it
 - Session time limit for active sessions
 - Idle session time limits
 - Whether to disconnect or end a session when a session limit is reached or the connection is broken
 - The other temporary folders should be used per session, and whether they should be deleted on exit
- **Security**
 - From here, you can specify:
 - The security settings between a client and a server. The default setting is to negotiate security settings to use the most secure layer that a client can support.
 - An encryption level. The default encryption level is to use a client-compatible encryption level.
 - Whether to allow connections only from clients that are running Remote Desktop with Network Level Authentication.

- **Load Balancing**
 - If you have multiple RD Session Hosts in a collection, you can specify how many sessions can be created on each RD Session Host and prioritize session creation among servers in the collection by using the relative weight value.
- **Client Settings**
 - You can specify devices and resources on a client device, such as the **Clipboard** and printers, when a user connects to a session-based desktop deployment. You can also limit the maximum number of redirected monitors per user session
- **User Profile Disks**
 - You can configure the use of user profile disks when a client connects to servers in a collection. You can also configure where user profile disks are stored, and limit their size.
 - You can also specify which files and folders from a user profile to exclude from a user profile disk.

Note: If you want to configure session collection settings by using Windows PowerShell, you can use the **Set-RDSessionCollectionConfiguration** cmdlet.

Demonstration: Create and configure a session collection

In this demonstration you will learn how to:

- Create and configure a session collection.
- Connect to Remote Desktop Services (RDS).

Preparation Steps

The required virtual machines (VMs) **WS-011T00A-SEA-DC1**, **WS-011T00A-SEA-RDS1**, and **WS-011T00A-SEA-CL1** should still be running after the previous demonstration.

After you complete this demonstration, revert all running VMs.

Demonstration steps

Create and configure a session collection using Server Manager

Use the following steps to create a session collection using Server Manager:

1. On **SEA-RDS1**, open Server Manager, and then select **Remote Desktop Services**.
2. On the **Remote Service Overview** page, select **Collections**.
3. In the details pane under **COLLECTIONS**, select **TASKS**, and then select **Create Session Collection**.
4. In the **Create Collection Wizard**, on the **Before you begin** page, select **Next**.
5. On the **Name the collection** page, in the **Name** field, enter **Demo**, and then select **Next**.
6. On the **Specify RD Session Host servers** page, in the **Server Pool** section, select **SEA-RDS1 . Contoso.com**.
7. Add the computer to the **Selected** section by selecting the Right arrow, and then selecting **Next**.

8. On the **Specify user groups** page, select **Next**.
9. On the **Specify user profile disks** page, clear the check box next to **Enable user profile disks**, and then select **Next**.
10. On the **Confirm selections** page, select **Cancel**, and then when prompted, select **Yes**.
11. Minimize **Server Manager**.

Create and configure a session collection using Windows PowerShell

Use the following steps to create a session collection using Windows PowerShell:

1. ON **SEA-RDS1**, right-click or access the context menu for **Start**, and then select **Windows PowerShell (Admin)**.
2. In Windows PowerShell, at the command prompt, enter the following command, and then select Enter:


```
New-RDSessionCollection -CollectionName Demo -SessionHost SEA-RDS1.Contoso.com -CollectionDescription "This Collection is for Demo purposes" -ConnectionBroker SEA-RDS1.Contoso.com
```
3. Wait for the command to complete, which will take approximately 1 minute.
4. Maximize **Server Manager**, and then select **Overview**.
5. Refresh **Server Manager** by selecting **F5**.
6. In **Server Manager**, in the navigation pane, select **Collections**, and verify that a collection named **Demo** is listed in the details pane.
7. In the navigation pane, select the **Demo** collection.
8. Next to ****PROPERTIES**, select **TASKS**, and then select **Edit Properties**.
9. On the **Session Collection** page, select the various settings and notice how the collection is configured.
10. Select **Client Settings**, and verify that both **Audio and video playback** and **Audio recording** are enabled.
11. In the **Demo Properties** dialog box, select **Cancel**.
12. Minimize **Server Manager**.
13. In Windows PowerShell, enter the following command, and then select Enter:


```
Get-RDSessionCollectionConfiguration -CollectionName Demo -Client | Format-List
```
14. Examine the output, and notice that next to **ClientDeviceRedirectionOptions** the following are listed:
 - **AudioVideoPlayBack**
 - **AudioRecording**
 - **PlugAndPlayDevice**
 - **SmartCard**
 - **Clipboard**
 - **LPTPort**

- **Drive**

15. Enter the following command, and then select Enter:

```
Set-RDSessionCollectionConfiguration -CollectionName Demo -ClientDeviceRedirectionOptions PlugAndPlayDevice, SmartCard,Clipboard,LPTPort,Drive
```

16. Enter the following command, and then select Enter:

```
Get-RDSessionCollectionConfiguration -CollectionName Demo -Client | Format-List
```

17. Examine the output and notice that next to **ClientDeviceRedirectionOptions**, only the following entries are now listed:

- **PlugAndPlayDevice**
- **SmartCard**
- **Clipboard**
- **LPTPort**
- **Drive**

Connect to Remote Desktop Session Host (RD Session Host) from a client

1. On **SEA-CL1**, on the taskbar, select the **Microsoft Edge** icon.
 2. In **Microsoft Edge**, in the address bar, enter `https://SEA-RDS1.Contoso.com/rdweb`, and then select Enter.
 3. On the **This site is not secure** page, select **Details**, and then select **Go on to the webpage**.
- Note:** You are getting this warning because Remote Desktop Web (RD Web) is using a self-signed certificate that is not trusted by the client. In a real production deployment, you would use trusted certificates.
4. On the **RD Web Access** page, in the **Domain\user name** field, enter `contoso\jane`. In the **Password** field, enter `Pa55w.rd`, and then select **Sign in**.
 5. If prompted by **Microsoft Edge** to save the password, select **Never**.
 6. On the **RD Web Access** page, under **Current folder: /**, select **Demo**, and then when prompted, select **Open**.
 7. In the **Remote Desktop Connection** dialog box, select **Connect**.

Note: You are receiving the **Unknown publisher** warning because you have not yet configured certificates for RDS.

8. In the **Windows Security** dialog box, in the **Password** field, enter `Pa55w.rd`. You are now connected to the RD Session host.
9. In the RDP connection, right-click or access the context menu for **Start**, select **Shut down or sign out**, and then select **Sign out**.
10. On the **RD Web Access** page, select **Sign out**.
11. Close **Microsoft Edge**.

After completing the demonstration, revert all VMs.

High availability options for RDS

In an ideal computer environment, servers would always be available and free of failures. Bandwidth and other resources would be infinite, and you wouldn't need to worry about high availability. In reality, no matter how reliable hardware is, components fail from time to time. Although rare, power outages or natural disasters such as earthquakes or hurricanes are always a possibility. And system restarts make server downtime unavoidable.

You need to take these forms of downtime into consideration if you plan to provide uninterrupted (or *highly available*) services. *High availability* means that systems and services are up and running, regardless of service outages. The goal of high availability is to make systems and services constantly available and to eliminate potential single points of failure. Different services provide high availability in different ways. For example, Active Directory Domain Services (AD DS) achieves high availability by using multiple domain controllers. If one domain controller fails or needs to restart, clients automatically connect to another domain controller. Failover clustering provides high availability for file servers. The role automatically fails over to another node if a node with the file server role fails.

High availability is often expressed numerically as the percentage of time that a service is available. For example, a requirement for 99.9 percent availability allows 8.75 hours of downtime per year, or approximately 40 minutes of downtime every four weeks. However, with 99.999 percent uptime, the allowed service downtime reduces to only 5 minutes per year. If your service runs on a single system, these high availability rates are virtually unachievable because a single restart most likely uses up those 5 minutes. In addition, many actions such as upgrading hardware or applying updates require a system restart.

To make an RDS deployment highly available, you first must ensure that the hardware on which RDS role services are running is as reliable as possible. You should also store application data and user state data on highly available storage to make them available if one of the servers fails. You should provide redundancy for all components, including power supplies and network paths. If users cannot access RDS because of network failure, there is no benefit to running RDS.

What Is Network Load Balancing?

Network Load Balancing (NLB) is an effective and scalable way to achieve high availability for stateless services such as a Web server. The term *stateless* refers to workloads that respond to each request independently from previous requests, and without keeping client state. For example, when a client requests a webpage, a Web server gathers all of the necessary information from the request and then returns a generated webpage to the client. When the client requests another webpage, it might request the webpage from the same Web server or from any other identically configured Web server in an NLB farm. This is because all the information that the Web server needs is in the request.

Using NLB to achieve high availability provides the following benefits:

- NLB enhances the availability and scalability of other Internet server applications, such as File Transfer Protocol (FTP), firewall, proxy, virtual private network (VPN), Remote Desktop Connection Broker (RD Connection Broker), Remote Desktop Web Access (RD Web Access), and other mission-critical servers. Both the Standard and Datacenter editions of Windows Server include the NLB feature.
- You can include up to 32 servers in an NLB farm, and you can dynamically add or remove a server from an NLB farm. For a load-balanced service, the load redistributes automatically among the servers that are still operating when a server fails or goes offline. If the failure is unexpected, only active connections to the failed server are lost. When you repair the server, it can rejoin the NLB farm transparently and regain its share of requests.
- RDS can use NLB for load-balancing requests and for providing high availability for most RDS role services. When you use NLB in unicast mode to distribute load among servers, you must enable media

access control (MAC) spoofing. This is because the network adapter doesn't use its own MAC address, but the MAC address of the unicast NLB. This is not required in multicast mode, as servers communicate with clients by using their own MAC address.

High availability for RD Session Host

A Remote Desktop Session Host (RD Session Host) server is required in each session-based desktop deployment of RDS. RD Session Host can provide remote desktops and RemoteApp programs to your users. When an RDS deployment has a single RD Session Host server, that server becomes a single point of failure. When it fails, the failure will affect all users who are connected to it and who run RemoteApp programs on that server. You must consider the possibility of failure or the lack of RD Session Host server availability in your disaster recovery plan.

You can take several steps to improve RD Session Host availability. You should always use reliable and redundant hardware from respected vendors to minimize the probability of hardware failure. You should also make sure that the network is reliable and that there are multiple network paths to an RD Session Host server. However, you should be aware that failures are unavoidable, and no single server can always be available without downtime. Therefore, you should make sure that an RDS deployment has more than one RD Session Host server and that multiple RD Session Host servers are in each collection. Servers in the same collection must be configured similarly, as client connections can be directed to any of the servers in a collection. For example, if RemoteApp programs publish in a collection, those applications must install on all servers in the collection. Because servers in a collection have the same configuration, it doesn't matter which RD Session Host server in a collection the clients connect to.

Even with multiple RD Session Host servers in the same collection, a collection is accessible from a single connection point, or *fully qualified domain name* (FQDN), which should point to RD Connection Broker. In a request to an RD Connection Broker, a client passes the information about the collection to which it wants to connect. RD Connection Broker accepts the request and directs the client to an available RD Session Host in the collection. If any of the servers in the collection are not available, clients will not be directed to those servers.

Note: Information on which collection the clients want to connect to is included in the .rdp file that downloads from RD Web Access when you select the link to a collection. There is no user interface in the Remote Desktop Connection (RDC) client to enter this information, and because of that, you should always connect to a collection by using RD Web Access and not by using RDC directly.

When a collection includes multiple RD Session Host servers, it can provide better performance compared to a collection with a single server. In such an environment, client connections distribute among the servers in the collection, and they use resources on all servers, not just from a single server. For example, if 50 users run the same RemoteApp program in an environment with a single RD Session Host server, all 50 instances of the application run on the same server. If a collection has two RD Session Host servers, each of them runs 25 application instances.

To understand how a collection provides high availability for multiple RD Session Host servers in the same collection, consider the following example. If a collection has a single server and that server fails, clients will not be able to connect to any RDS resource in that collection. If two servers are in the collection, RD Connection Broker decides which RD Session Host server in the collection is most suitable and sends that information to the client. The client establishes a connection with that RD Session Host server, for example, Server 2. If Server 2 fails, it can no longer accept connections. When clients initiate requests to the collection, RD Connection Broker is aware that Server 2 is not available and directs the client to connect to Server 1, which is still available and is in the same collection. The client establishes a connection with Server 1, and because Server 1 has the same configuration as Server 2 (as both servers are in the same collection), the user experience is the same. In this way, a client can access a required resource even

when one or more RD Session Host servers in a collection are not available. A client can access a resource providing there are servers in the collection that can accept a connection.

High availability for RD Connection Broker

When an RDS deployment has multiple RD Session Host servers in the same collection, RD Connection Broker provides high availability for RDS resources in that collection. RD Connection Broker receives client requests and directs clients to the available RD Session Hosts in the collection even if some of the RD Session Host servers in the collection are unavailable. However, this configuration relies on the availability of the RD Connection Broker server, which presents a single point of failure.

To avoid single point of failure, you can configure RD Connection Broker servers in a highly available active/active mode configuration. With this configuration, you can include two or more RD Connection Broker servers in an RDS deployment and configure them to use an external database. This provides fault tolerance and load balancing for clients that connect to RDS, and eliminates RD Connection Broker as a single point of failure, thereby increases scalability. When you configure RD Connection Broker high availability, more clients can redirect at the same time, which reduces the time that is required for clients to connect, especially in large RDS deployments.

Configuring RD Connection Broker for high availability

If you want RD Connection Broker to be highly available, you must perform the following high-level steps:

1. Prepare the environment.
2. Configure RD Connection Broker for high availability.
3. Add additional RD Connection Broker servers to the RDS deployment.

To prepare the environment for RD Connection Broker high availability, several prerequisites must be met:

- You must have a Microsoft SQL Server or an Azure SQL Database running that is accessible by all RD Connection Broker servers. The database must run SQL Server 2014 or later, and the RD Connection Broker servers must have permission to create a database on the server. SQL Server is used to host the shared RD Connection Broker database.
- You must create a folder to store files for an RD Connection Broker database. This folder can be a local folder on a computer that is running SQL Server or a shared folder that is accessible by using a Universal Naming Convention (UNC) path.
- The SQL Native Client must be installed on all RD Connection Broker servers.
- All RD Connection Broker servers must have static IP addresses.
- You must configure a round robin Domain Name System (DNS) record for the RD Connection Broker servers.

Note: When you configure RD Connection Broker for high availability, its database moves from a local Windows Internal Database (WID) to a computer that is running SQL Server. Even when an RDS deployment has multiple RD Connection Broker servers, SQL Server can still be a single point of failure. You should make sure that SQL Server is highly available by running it in a failover cluster or using an Azure SQL database.

To prepare an environment for RD Connection Broker high availability, you need to perform the following high-level steps:

1. Create a security group in AD DS and add all of the RD Connection Broker computer accounts as group members. You will also need to restart the RD Connection Broker servers to update their group membership security token.
2. On the computer that runs SQL Server, create a SQL sign-in account that maps to the AD DS group, and place it in a server role that has permissions to create databases.
3. Install the SQL Native Client on all RD Connection Broker servers.
4. Create DNS round robin resource records and map them to each RD Connection Broker server, as in the following example:

```
rds.contoso.com 172.16.10.30 rds.contoso.com 172.16.10.31 rds.contoso.com
172.16.10.32
```

To configure RD Connection Broker for high availability, you should perform following steps:

1. In Server Manager, on the **Remote Desktop Services** page, in the **Overview** section, right-click or access the context menu for the **RD Connection Broker** node, and then select **Configure High Availability**.
2. Enter configuration information for the RD Connection Broker servers:
 - **Database connection string.** This string contains the necessary information to connect to the computer that is running SQL Server and to access the RD Connection Broker connections database. For example, the following sample string would connect to the RDCB_DB that is hosted on the computer that is running a SQL Server named SEA-SQL.contoso.com:


```
DRIVER=SQL Server Native Client 11.0;SERVER=SEA-SQL.contoso.com.;Trust-
ed_Connection=Yes;APP=Remote Desktop Services Connection Broker;Data-
base=RDCB_DB
```
 - **Folder to store database files.** This is a local or UNC path to the folder that you created for storing RD Connection Broker database files.
 - **DNS round robin name.** This is the FQDN for the DNS name that you assigned to a server farm—for example, rds.contoso.com.

After you've configured RD Connection Broker for high availability, you should add additional RD Connection Broker servers to the RDS deployment. In Server Manager, on the **Remote Desktop Services** page, in the **Overview** section, right-click or access the context menu for the **RD Connection Broker** node, select **Add RD Connection Broker Server**, and then complete the wizard.

Note: In an RDS deployment with multiple RD Connection Brokers servers, you should configure all RD Connection Broker servers with the digital certificate whose common name is the same as the FQDN that clients will use to connect to the RDS deployment.

You can also configure RD Connection Broker for high availability if you use the Windows PowerShell **Set-RDConnectionBrokerHighAvailability** cmdlet.

Additional reading: For additional information on RD Connection Broker high availability, refer to **Add the RD Connection Broker server to the deployment and configure high availability**⁴.

⁴ <https://aka.ms/rds-connection-broker-cluster>

High availability for other RDS sessions infrastructure

If you want an RDS deployment to be highly available, all RDS role services must be highly available. How to achieve high availability for the RD Session Host and RD Connection Broker role services was discussed earlier. In this topic, you will learn how to achieve high availability for other RDS role services.

High availability for RD Web Access

RD Web Access provides a portal with links to available remote desktops, RemoteApp programs, and virtual desktops. It provides information on which RDS resources are available and where you need to connect to access them. That information is also available as an RDWeb feed, which can integrate with the client **Start** menu.

RD Web Access uses a Web server as its platform. You can achieve RD Web Access high availability in the same way as any other Web server—by adding multiple RD Web Access servers to an RDS deployment. Next, you should provide a single entry point to all RD Web Access servers by adding servers to the NLB cluster, or you can use DNS round robin.

Note: RD Web Access uses a self-signed Secure Sockets Layer (SSL) certificate by default. In your RDS deployment, you should replace it with an SSL certificate from a trusted CA. When you add multiple RD Web Access to an RDS deployment and they are all accessible by using the same FQDN, you should make sure that the SSL certificate is also configured with the same FQDN as its common name.

High availability for RD Licensing

Remote Desktop Licensing (RD Licensing) issues licenses to RDS users and devices. Client devices cannot connect to RD Session Host without a valid RDS license if Per Device licensing mode is configured. After a client obtains an RDS license, they can connect to an RD Session Host server even if an RD Licensing server is not available. RD Licensing must be available if devices that do not yet have valid RDS client access licenses (CALs) want to connect to an RD Session Host. You can achieve high availability for RD Licensing if you add additional RD Licensing servers to an RDS deployment. You should also activate the RD Licensing server and add RDS licenses to a RD Licensing server.

High availability for RD Virtualization Host

A Remote Desktop Virtualization Host (RD Virtualization Host) runs virtual machines (VMs) that are used as virtual desktops. Virtual desktops can be either personal or pooled. To achieve high availability, you need to add additional RD Virtualization Host servers that are hosting pooled virtual desktops to the same collection. Then, if one RD Virtualization Host server fails, the RD Connection Broker redirects clients to the remaining RD Virtualization Host servers that host the same pooled virtual desktops.

Note: VMs that are used for virtual desktops must be highly available to achieve high availability for personal virtual desktops. VMs must run on a Microsoft Hyper-V failover cluster to achieve high availability.

High availability for RD Gateway

Remote Desktop Gateway (RD Gateway) provides access to RDS resources for external users, typically over the internet. It encapsulates Remote Desktop Protocol (RDP) traffic into HTTPS packets and sends them securely over a public network. You can control who can access RDS resources and to which RDS resources users can connect on an RD Gateway. You can also enforce organizational policy—for example, which local resources can redirect to RD Session Host servers.

The RD Gateway is an entry point to an RDS deployment for external users. If an RD Gateway fails, external users will not be able to connect to internal RDS resources. You can achieve high availability for an RD Gateway if you add multiple RD Gateway servers to an RDS deployment. You should use the same FQDN by adding them to an NLB cluster or by using DNS round robin to make the RD Gateways accessible.

Note: RD Gateway uses an SSL certificate for securing network communication. When you add multiple RD Gateway servers to an RDS deployment and they all use the same FQDN to access the network, make sure that the SSL certificate is also configured with the same FQDN as its common name.

Overview of RemoteApp

With RemoteApp programs, you can use Remote Desktop Services (RDS) to make programs that are accessed remotely seem as if they are running on a user's local computer. RemoteApp program windows display and are integrated with a client's desktop instead of being presented as a part of a Remote Desktop Session Host (RD Session Host) server's desktop. A RemoteApp program runs in its own resizable window, which a user can move between multiple monitors, and it has its own icon on the taskbar.

If a remote application uses a notification area icon, this icon appears in the client's notification area. RDS redirects dialog boxes and other windows to the local desktop. If a user runs more than one RemoteApp program on the same RD Session Host server, the RemoteApp programs share the same session. Users can access RemoteApp programs through a browser when they use Remote Desktop Web Access (RD Web Access), or from the **Start** menu when they use RemoteApp and Remote Desktop Connection.

RemoteApp programs are especially useful for:

- Remote users. Users often need to access applications from remote locations, for example, when working from home or travelling. RemoteApp programs allow these users to access applications over an internet connection. Using Remote Desktop Gateway (RD Gateway) with RemoteApp programs helps secure remote access to applications. Additionally, you can allow users to access remote applications through an **RD Web Access** page, or integrate applications on their **Start** menu.
- Line-of-business application deployments. Organizations often need to run consistent line-of-business (LOB) applications on computers and devices that run different versions of Windows and non-Windows operating systems (OSs). Instead of deploying LOB applications locally, you can install applications on an RD Session Host server and make them as available as RemoteApp programs.
- Roaming users. In some organizations, a user might work on several different computers. If users work on a computer where an application is not installed, they can access the application remotely through RDS.
- Branch offices. Branch office environments often have limited local IT support and limited network bandwidth. By using RemoteApp programs, you can centralize application management and improve remote application performance in limited bandwidth scenarios.

RemoteApp programs can integrate with a local client desktop, and operate just like applications that install locally. Users might not even be aware that RemoteApp programs are running remotely. They also might not be aware that the programs are running side by side with locally installed applications.

RemoteApp programs include the following features:

- Start a RemoteApp program without prompts. When a user selects a RemoteApp program link, the program can start without any prompts or user interaction. In the background, the client establishes a Remote Desktop Protocol (RDP) connection, signs in, starts the remote program, and displays its window.

- Has its own window. A RemoteApp program displays in its own window on a client. You can move, resize, minimize, maximize, or close the window the same way as any other application window. A RemoteApp window can display its content while you move or resize the window.
- Multiple ways to start RemoteApp programs. You can start a RemoteApp program from an **RD Web Access** page, from the **Start** menu, or by double-clicking or selecting the spacebar and then entering a file with an associated file name extension.
- Live thumbnails and application switching. A RemoteApp program icon displays on the taskbar even if the program is minimized. If multiple instances of a RemoteApp program are running, multiple, tabbed program icons display on the taskbar. When you move the pointer to the taskbar icon, a live thumbnail of the program window displays. You can use the standard Alt+Tab key combination to switch between running programs, including RemoteApp programs.
- Similar icons. The RemoteApp program icons on the taskbar are similar to locally installed application icons, but they include a Remote Desktop symbol. Any change in status for a RemoteApp program is represented with an icon overlay. For example, Microsoft Outlook uses a letter overlay to notify that new email has been received.

Test your knowledge

Use the following questions to check what you've learned in this lesson.

Suggested answer

Question 1

Can you use Windows Internal Database (WID) with a highly available Remote Desktop Connection Broker (RD Connection Broker) configuration?

Suggested answer

Question 2

Why would you use a RemoteApp program instead of a remote desktop session?

Suggested answer

Question 3

You want to install the RD Connection Broker role service on a server named "SEA-RDS2", but only a server named "SEA-RDS1" displays in Server Manager when you need to specify the RD Connection Broker server. What should you do to add "SEA-RDS2" as a possible selection?

Overview of personal and pooled virtual desktops

Lesson overview

Virtual Desktop Infrastructure (VDI) is the Microsoft technology that enables users to access desktops running in a datacenter. Virtual Desktop Infrastructure (VDI) is based on virtual machines (VMs) that run on a Microsoft Hyper-V server. These VMs can either be assigned to one user (a personal virtual desktop), or shared between users as a pooled virtual desktop. Unlike Remote Desktop Services (RDS) session-based desktop deployments where users share resources on a single server, VM-based desktop deployments enable you to allocate specific resources to each user, because each user is allocated their own VM. To ensure that you select and implement an appropriate solution, you need to understand the characteristics of personal and pooled virtual desktops.

One key difference between personal and pooled virtual desktops is state retention. Pooled virtual desktops don't retain state information between sessions, but personal virtual desktops do. A *virtual desktop template* is the base VM that is copied to create personal and pooled virtual desktops. A poorly configured desktop template will result in poorly performing personal and pooled virtual desktops. To configure a desktop template properly, you need to understand the process for creating a desktop template, configuration considerations such as which operating system (OS) to select, and how to configure application updates. To have a successful VDI, you'll also need to understand how to provide high availability for personal and pooled virtual desktops.

Lesson objectives

After completing this lesson, you will be able to:

- Describe VM-based desktop deployments of VDI.
- Describe how pooled virtual desktops work.
- Describe how personal virtual desktops work.
- Explain the differences between VDI options.
- Describe how to provide high availability for pooled virtual desktops.
- Describe how to provide high availability for personal virtual desktops.
- Explain how to optimize and prepare a virtual desktop template.

Overview of VM-based desktop deployments of Virtual Desktop Infrastructure

In a development or test environment, you can meet your virtual machine (VM) requirements by manually creating a few VMs and hosting them on a server running Microsoft Hyper-V or by using Client Hyper-V on a desktop computer. This type of solution is not suitable for VM-based desktop deployments of Virtual Desktop Infrastructure (VDI), where you need to create many VMs to support users. You also need a method for connecting users to an appropriate VM.

When you implement a VM-based desktop deployment of VDI, you implement many of the same server role services as when you implement a session-based desktop deployment of VDI. The main difference is that in a VM-based desktop deployment of VDI, users connect to a VM that is hosted on a Remote Desktop Virtualization Host (RD Virtualization Host) server instead of a session that is hosted on a

Remote Desktop Session Host (RD Session Host) server. VM resources are dedicated to that VMs user, whereas RD Session Host resources are shared among multiple users.

A VM-based desktop deployment of VDI uses the following server role services:

- Remote Desktop Connection Broker (RD Connection Broker). Clients connect to an RD Connection Broker server and are directed to an appropriate VM to which they have been granted access.
- Remote Desktop Web Access (RD Web Access). RD Web Access provides an .rdp file that contains configuration information that is necessary for a client to connect to an RD Connection Broker.
- RD Virtualization Host. RD Virtualization Host servers host VMs for VDI. They also have the Hyper-V server role service installed.
- Remote Desktop Gateway (RD Gateway). For organizations that provide external access to VMs, RD Gateway controls access to them.

Overview of pooled virtual desktops

Pooled virtual desktops are a type of virtual machine (VM)-based desktop deployment of Virtual Desktop Infrastructure (VDI), where users connect to a pool of identically configured VMs. Each VM has exactly the same resources and operating system (OS) configuration. If applications install on a VM, then all pooled virtual desktops have those same applications.

VMs for pooled virtual desktops don't assign to specific users; any user can connect to any pooled virtual desktop that is available. The RD Connection Broker server is responsible for selecting a pooled virtual desktop for a specific user during the connection process. When the user signs out from the pooled virtual desktops. When a user signs out from a pooled virtual desktop, the pooled virtual desktop reverts to the state it was in before the user signed in. User state information is not retained, and any information that is stored on a pooled virtual desktop is removed.

A *virtual desktop image* is a VM that you have installed and configured. A virtual desktop image that you prepare and create defines the starting state for a pooled virtual desktop. Pooled virtual desktops have differencing virtual hard disks that use a virtual desktop image as a base. All changes to pooled virtual desktop hard disks are stored on the differencing virtual hard disk. When a user signs out, the differencing virtual hard disk is removed.

Overview of personal virtual desktops

Personal virtual desktops are a type of virtual machine (VM)-based desktop deployment of Virtual Desktop Infrastructure (VDI), where users connect to a VM that is specifically assigned to them. Unlike pooled virtual desktops, personal virtual desktops are not shared between multiple users. Each personal virtual desktop is dedicated to one specific user.

One of the primary reasons to implement personal virtual desktops is to support user customization of virtual desktops. A personal virtual desktop retains all changes to it when a user signs out. This enables users to customize personal virtual desktops with additional applications and user state information. If you want to provide users with the ability to install their own applications, then you should give the users administrative permissions on their personal virtual desktops.

You can create personal virtual desktops in two ways:

- Base on a virtual desktop image. When you create personal virtual desktops based on a virtual desktop image, all personal virtual desktops have an identical starting configuration. This method allows you to create many personal virtual desktops quickly for a deployment, in much the same way as using images for operating system (OS) deployments on physical computers.

- Use an existing VM. When you create a personal virtual desktop from a specific VM, the VM converts to a personal virtual desktop. This can be useful if you have existing VMs that you want to convert to personal virtual desktops.

Compare VDI options

Choosing the correct type of Virtual Desktop Infrastructure (VDI) for your users is essential because it ensures that you meet their needs in a cost-effective way. The four main considerations for VDI options are personalization, application compatibility, ease of management, and cost effectiveness.

Personalization

The best option for personalization is personal virtual desktops. With personal virtual desktops, users can be assigned permissions to customize their personal virtual desktop completely, including applications. This can be useful when users have unique configuration needs for their virtual desktops, such as installing their own applications or performing operating system (OS) customization.

Note: Pooled virtual desktops and session-based desktop deployments of VDI do provide personalization options through other methods. You can implement Windows user state virtualization to make user states persistent. To provide access to unique applications, you can use RemoteApp programs.

Application compatibility

For the best application compatibility, use personal and pooled virtual desktops. The majority of applications that will install on a desktop computer can also be used with both personal and pooled virtual desktops, and can install on session-based desktop deployments. However, but some end-user applications do not run properly on a Remote Desktop Session Host (RD Session Host) server because the application uses a server OS.

Ease of management

The ease of management for VDI solutions depends on the amount of standardization. Fewer images generally result in easier to manage solutions. A session-based desktop deployment of VDI is the easiest to manage because you only need to update applications and the OS on the RD Session Host servers. Pooled virtual desktops also are easy to manage because you update only the desktop virtual template. Personal virtual desktops, however, require you to provide updates to each individual virtual machine (VM).

Cost effectiveness

The cost to implement a VDI solution is based on the resources that are required to support a specific number of users. A session-based desktop deployment of VDI is the most cost effective because it has much higher user density per server than VM-based desktop deployments of VDI. Pooled virtual desktops are more cost-effective than personal virtual desktops because by using a base image and a differencing virtual hard disk that clears when users sign out, they use less storage. Personal virtual desktops can be customized, and the disk sizes for personal virtual desktops grow over time as applications and updates are installed.

High availability for personal and pooled desktops

The process for providing highly available pooled virtual desktops is similar to providing highly available session-based desktop deployments. Each server role should be redundant, and there must be multiple Remote Desktop Virtualization Host (RD Virtualization Host) servers. The following table details how to make each server role highly available.

Table 2: Server role high availability methods

Server role	High availability method
Remote Desktop Connection Broker (RD Connection Broker)	Domain Name System (DNS) round robin and Microsoft SQL Server or Microsoft Azure SQL Database configured to store RD Connection Broker configuration
Remote Desktop Web Access (RD Web Access)	Load balancing
RD Virtualization Host	Multiple RD Virtualization Hosts

When you configure high availability for pooled virtual desktops, individual virtual machines (VMs) are not highly available, but the entire infrastructure is. When an RD Connection Broker server or RD Web Access fails, it doesn't affect users who are currently signed in to pooled virtual desktops. However, if an RD Virtualization Host fails, then all currently running pooled virtual desktops on that server fail and users are disconnected.

When users reconnect, the RD Connection Broker server connects them to a pooled virtual desktop that is running on a different RD Virtualization Host server. All state information from the previous pooled virtual desktop is lost unless some type of user state virtualization has been implemented. Functionally, this is similar to the process when a user signs out from a pooled virtual desktop and then signs in again.

High availability for personal virtual desktops

Unlike pooled virtual desktops, which are interchangeable, personal virtual desktops are unique. To make individual VMs highly available for personal virtual desktops, you need to configure RD Virtualization Host servers as nodes in a failover cluster. You also need to configure shared storage that is accessible to all nodes in the failover cluster.

Failover clustering

Failover clustering is a feature in both the Standard and Datacenter editions of Windows Server. You can use failover clustering to enable personal virtual desktops to move between RD Virtualization Host servers in cases of unplanned downtime. If you have planned downtime, then you can use the Live Migration feature to move the personal virtual desktop between nodes. When you use Live Migration, a personal virtual desktop continues to run while it moves. A user is unaware that a live migration has even occurred and continues to work without interruption.

If a personal virtual desktop has an unplanned move to another node (for example, because of the failure of an RD Virtualization Host server), then the user is disconnected from the personal virtual desktop and must wait for it to restart on another node. Any unsaved work in progress at the time of the failure will be lost. This process is similar to what happens when a standard desktop computer loses power and restarts.

Shared storage

You can make a personal virtual desktop available to all nodes in a failover cluster by placing it on storage that all of the nodes can access. This allows all of the nodes in a failover cluster to use the virtual disk and configuration files that are required to start the personal virtual desktop. Traditionally, shared storage for a failover cluster was:

- Shared serial-attached small computer system interface (SCSI)
- Internet SCSI (iSCSI) storage area network (SAN)
- Fibre Channel SAN

In these traditional configurations, shared storage is normally configured as a Cluster Shared Volume (CSV). You can store multiple VMs on a single CSV because multiple nodes can access a CSV simultaneously. A node locks individual files when they are in use to ensure that files are not corrupted by two nodes accessing a file at the same time. Another option is to use Storage Spaces Direct to store your virtual desktops.

You can also use a file share as storage for Microsoft Hyper-V VMs. This is less complex to implement than traditional shared storage, and possible because of performance improvements in the Server Message Block (SMB) 3.0 protocol. You can make file shares highly available in Windows Server by implementing the Scale-Out File Server feature.

Additional reading: For additional information on Scale-Out File Server, refer to **Scale-Out File Server for application data overview**⁵.

Networking

You typically configure a failover cluster with multiple networks. Each node in a failover cluster has access to all of the networks. Some networks that you might configure include:

- **Management network.** Administrators use this network to connect to failover cluster nodes to perform management actions.
- **VM network.** Clients use this network to connect to personal virtual desktops.
- **Heartbeat network.** Nodes use this network to communicate with each other and to identify when other nodes have failed.

Prepare a virtual desktop template

A *virtual desktop template* is a virtual machine (VM) that functions as a starting point for personal or pooled virtual desktops. The process of creating a desktop template is important because a poorly configured template results in poorly performing virtual desktops for users. It might also waste server-side resources on a Remote Desktop Virtualization Host (RD Virtualization Host) and reduce scalability, which increases overall costs.

Before you begin creating and configuring a desktop template, you need to determine information such as:

- Will you use a 32-bit or 64-bit operating system (OS)?
- How much random access memory (RAM) do virtual desktops require?
- How many virtual central processing units (CPUs) are required?

⁵ <https://aka.ms/sofs-overview>

- Will the desktop template include applications, or will another method deliver them?
- Which application configuration options can reduce resource utilization?
- Which Windows services are unnecessary and should be disabled?

To prepare a desktop template, perform the following high-level steps:

1. Create a VM on a server running Microsoft Hyper-V.
2. Install the selected OS on the VM.
3. Install selected applications on the VM.
4. Optimize application configuration for virtual desktops.
5. Optimize OS configuration for virtual desktops.
6. Run the System Preparation Tool (Sysprep) to prepare the OS.

After configuring the desktop template, you can create a Remote Desktop session collection for the personal or pooled virtual desktops. The **Create Collection Wizard** asks you to identify the virtual desktop template and copies the virtual desktop template to create virtual desktops.

Selecting an OS for personal and pooled virtual desktops

For both personal and pooled virtual desktops, you need to select the Windows client OS that will install on VMs for the users. The OS version and edition that you select will be based on the features that you require. The Windows client OS you should consider using is Windows 10 Enterprise.

64-bit vs. 32-bit OS versions

The main differences between 32-bit and 64-bit versions of Windows client OSs are:

- **Memory support.** A 32-bit OS supports a maximum of 4 gigabytes (GB) of RAM. If you require more than 4 GB of RAM for application performance, then you should use a 64-bit Windows client OS.
- **Resource utilization.** Memory usage for Windows 10 x86 is 1 GB, and 2 GB for Windows 10 x64. Storage utilization for both 64-bit and 32-bit versions of Windows 10 is approximately 15 GB. When you create hundreds or thousands of personal and pooled virtual desktops, a savings of approximately 2 GB per desktop can be significant. You can save considerably amounts of storage by enabling Data Deduplication on the volume where the VMs are stored.
- **Application compatibility.** A 64-bit OS can run both 32-bit and 64-bit applications. In reality, you should only consider using 32-bit if you have older 16-bit applications, because 64-bit OSs do not support 16-bit applications.

Unless there is a specific reason to use the 32-bit version of Windows 10, you should use a the 64-bit version.

Features

You should use the Windows 10 client OSs with your personal and pooled virtual desktops. In addition, only the Enterprise edition of Windows 10 supports personal and pooled virtual desktops. The Enterprise editions of Windows 10 supports the following virtualization features:

- **RemoteApp.** Windows 10 Enterprise can host and provide applications to other computers by using RemoteApp. For example, if a user has a standard desktop computer with an OS and installed applications, a personal or pooled virtual desktop that uses Windows 10 Enterprise can provide an installed

application to the standard desktop by using RemoteApp. This can be useful when an application cannot install on a Remote Desktop Session Host (RD Session Host) or is required by only a few users.

- **Microsoft Multi-Touch.** Windows 10 Enterprise supports touch-enabled devices. This is important if you use touch-enabled devices to access personal and pooled virtual desktops.
- **USB Redirection.** Windows 10 Enterprise supports the redirection of USB devices from a local client to a personal or pooled virtual desktop. This allows personal or pooled virtual desktops to use various local USB devices, such as printers, scanners, and audio devices.
- **Discrete Device Assignment (DDA).** The DDA feature allows you to pass PCIe devices into a VM. This enables the VM to use a graphics card or NVMe storage devices from the virtualization host in Windows 10 Enterprise using the normal drivers.
- **User profile disk.** You can use a user profile disk with Windows 10 Enterprise to implement user state virtualization. This is useful for pooled virtual desktops, where changes are not retained between sessions.

Activation considerations for desktop templates

When you create a virtual desktop template with Windows 10, you need to consider how you will activate Windows 10 after creating personal and pooled virtual desktops. The virtual desktop template activates following the personal or pooled virtual desktop deployments. Entering product keys manually and performing activation after deploying personal and pooled virtual desktops isn't practical. Therefore, you should automate the activation process.

When you install a volume license version of Windows 10, it includes a generic volume license key (VLK). This license key instructs Windows 10 to search for an automated method of activation for volume licensing. For personal and pooled virtual desktops, you can use Active Directory–based activation or Key Management Service (KMS).

Active Directory–based activation

You can use Active Directory–based activation to store all necessary activation information for personal and pooled virtual desktops. The information is stored in Active Directory Domain Services (AD DS). When personal or pooled virtual desktops that have not been activated join the domain, they can access the activation objects in AD DS and activate. Personal and pooled virtual desktops join the domain as part of the creation process.

To configure Active Directory–based activation for an Active Directory forest, complete the following high-level steps:

1. Install the Volume Activation Services server role.
2. Open **Volume Activation Tools**, and select **Active Directory–based activation** in the **Volume Activation Tools Installation Wizard**.
3. Enter the KMS host key—the VLK that you have purchased—and then activate the key. (Typically, this is done online, but you can also activate it by phone.)
4. Close the **Volume Activation Tools Installation Wizard**.

KMS

When you install KMS on a computer, a service installs that can activate software. Similar to Active Directory–based activation, you enter a KMS host key into KMS to make the licenses available for clients

to activate. Client computers use Domain Name System (DNS) to identify the location of the KMS service. The KMS service creates a service (SRV) resource record in DNS to advertise its location. The KMS service listens on port 1688.

Comparing Active Directory–based activation and KMS

Active Directory–based activation is automatically highly available if you have multiple domain controllers. To make KMS highly available, you need to install two KMS servers. By default, clients will load balance between the two available KMS servers, and if one is unavailable, they will use the other.

One critical consideration for KMS is the minimum activation threshold. KMS will not activate client computers until a minimum of 25 computers have contacted KMS. This can be important for small Virtual Desktop Infrastructure (VDI) deployments and pilot projects where you will only be deploying a small number of personal and pooled virtual desktops. In contrast, Active Directory–based activation doesn't have minimum activation thresholds. However, for Active Directory–based activation, clients must be domain members and must reactivate every 180 days. Providing clients are connected to the domain network, reactivation occurs automatically.

Application considerations for desktop templates

When you plan personal and pooled virtual desktop deployments, you need to consider how applications will install and update on VMs. Personal and pooled virtual desktops have unique characteristics that require you to manage applications and updates in a different way for each type of virtual desktop.

Pooled virtual desktops

Pooled virtual desktops have unique application installation and update issues because they don't retain changes between sign-ins. Application and OS updates that are delivered by a solution such as Windows Server Update Services (WSUS) or Microsoft Endpoint Configuration Manager would be lost at each sign-out. Consequently, you can't use those tools. If applications include their own automatic update functions, you should disable them.

To deploy application and OS updates to a pooled virtual desktop, you modify the pooled virtual desktop and update it. Consequently, if you include many applications directly in the virtual desktop template for pooled virtual desktops, you'll need to update the pooled virtual desktops many times. This can become a significant management burden if done manually and should be avoided. As an alternative, consider using Microsoft Endpoint Configuration Manager to automatically build and configure images for pooled virtual desktops.

Personal virtual desktops

Application and OS updates for personal virtual desktops are managed similarly to standard desktop computers because they retain state the same way. Updates that install on a personal virtual desktop are retained between sign-ins. This means that you can use standard application updating methods such as WSUS, Microsoft Endpoint Configuration Manager, and Windows Update for Business for personal virtual desktops. You can also use application deployment tools such as Microsoft Endpoint Configuration Manager to deploy applications to personal virtual desktops. You are more likely to include applications in a virtual desktop image for personal virtual desktops than pooled virtual desktops. This is because you can deploy updates to applications without redeploying personal virtual desktops.

Antivirus considerations for desktop templates

When you implement personal and pooled virtual desktops, you need to consider the impact of antivirus software on system performance. Like desktop computers, personal and pooled virtual desktops can become infected with malware that can steal personal information or control the infected system. You should always use antivirus software for personal and pooled virtual desktops. Even though signing out from a pooled virtual desktop reverts its state and removes malware, malware is functional and can propagate while a pooled virtual desktop runs.

Windows Defender Antivirus software is part of Windows 10 and is an excellent choice when running Windows 10 in an VDI environment. Windows Defender Antivirus contains performance optimizations and features designed especially for Windows 10 running in VMs.

Performance

Antivirus software can adversely affect personal and pooled virtual desktop performance. Storage input/output (I/O) increases by using antivirus software because each time a file is scanned, it uses the storage subsystem. On a desktop computer, this effect is minimal because only a single computer is using the storage. However, each little bit of I/O per desktop adds up to a large amount of I/O for hundreds or thousands of personal and pooled virtual desktops.

If you are using third-party antivirus software, you should check with the software vendor to determine whether they have specific recommendations for implementing their product with personal and pooled virtual desktops. Here are some general recommendations for configuring antivirus software for personal and pooled virtual desktops:

- Disable scheduled scans on pooled virtual desktops. Scheduled scans read the entire disk searching for malware. This causes a high amount of I/O, which isn't necessary on pooled virtual desktops because signing out effectively removes any malware that installed.
- Randomize scheduled scans on personal virtual desktops. If your organization runs scheduled scans on desktop computers, you should also run scheduled scans on personal virtual desktops. However, because personal virtual desktops share the same storage infrastructure, you need to ensure that scheduled scans on personal virtual desktops are randomized and don't run at the same time. Running schedule scans simultaneously results in a large burst of I/O on the storage infrastructure. If the scheduled scans are randomized, then the load on the storage infrastructure is evened out.
- Randomize virus signature updates. Updating virus definition updates for antivirus software causes a brief burst of I/O. Just as with scheduled scans on personal desktops, this should be randomized to minimize the impact on the storage infrastructure.
- Don't perform scans on virus signature updates. Antivirus software has the option to perform entire system scans after updating virus signatures. The purpose of this scan is to identify recently installed malware that wasn't detected by the previous set of virus signatures. However, this places a large load on the storage infrastructure similar to running a scheduled scan.
- Prevent scanning of virtual desktop template files. If possible, in your antivirus software disable scanning of any files from the VM template. The files in this template are known to be malware free from your original build, and preventing their scan reduces overall resource utilization.

Additional reading: For additional information on running Windows Defender Antivirus in an VDI environment, refer to **Deployment guide for Microsoft Defender Antivirus in a virtual desktop infrastructure (VDI) environment**⁶.

⁶ <https://aka.ms/vdi-windows-defender-antivirus>

Optimizing OS services for desktop templates

OS services consume system resources. They can consume processing, memory, storage, and network resources. On an individual desktop computer, running unnecessary services has minimal impact because each service consumes a small amount of resources. For personal and pooled virtual desktops, even small amounts of unnecessary resource utilization per desktop adds up to a large amount of resource utilization on an overall infrastructure. Some Windows 10 services are not required in most VDI deployments, and you can disable them. Some of the services you should consider disabling are:

- Background Intelligent Transfer Service (BITS). This service downloads data as a background process. Most VDI deployments do not require this because they have fast network connectivity.
- Block Level Backup Engine Services. This service is used to back up data on a computer. This is not required for most VDI deployments because VMs typically are not backed up.
- Bluetooth Support Service. Personal and pooled virtual desktops do not support Bluetooth.
- Diagnostic Policy Service. This service is used for problem detection and resolution, which is not necessary for most VDI deployments. If necessary, you can manually enable it for troubleshooting.
- Shell Hardware Detection. This service provides notifications for hardware events that trigger AutoPlay. This isn't necessary for most VDI deployments because events that trigger AutoPlay, such as inserting a DVD, are not typically performed.
- Volume Shadow Copy Service (VSS). This service manages volume shadow copies for backup and restore points. This is not necessary for most VDI deployments because backups are not necessary.
- Windows Search. This service indexes local files such as cached Microsoft Outlook mailboxes for faster searching. You can disable this for VDI deployments that don't store data locally.

Other Windows 10 optimizations

In addition to disabling unnecessary OS services, consider making the following additional modifications to Windows 10 to optimize personal and pooled virtual desktop performance:

- Minimize event log sizes. Doing this minimizes storage utilization. You can increase event log size when necessary for troubleshooting.
- Disable unneeded scheduled tasks. This minimizes overall resource utilization. For example, the disk defragmentation task that runs once per week isn't useful in a virtualized environment.
- Minimize visual effects. This reduces virtual GPU (vGPU) utilization.
- Disable features that increase storage utilization. You can disable a number of features to reduce storage utilization, including:
 - Disable the New Technology File System (NTFS) file system's last access time stamp.
 - Disable System Restore.
 - Disable hibernation.

Additional reading: For additional information on optimizing Windows 10 for VDI, refer to **Optimizing Windows 10, version 1909, for a Virtual Desktop Infrastructure (VDI) role**⁷.

Additional reading: For additional information on the Windows 10 optimizing script for VDI, refer to **Windows 10 (Build 1803) VDI Optimization Script Primer**⁸.

⁷ https://aka.ms/rds_vdi-recommendations-1909

⁸ <https://aka.ms/vdi-optimization-script-primer>

Test your knowledge

Use the following questions to review what you've learned in this lesson.

Question 1

Which two types of virtual desktops can you create (or use) in a Remote Desktop Service (RDS) deployment? Choose all that apply.

- Pooled virtual desktops
- Virtual machine (VM) desktops
- Server desktops
- Personal virtual desktops
- Windows virtual desktops

Question 2

What is the reason for disabling unnecessary services in a virtual desktop template?

?

Question 3

You would like to provide your users with the option to customize their virtual desktop, and have those customizations be persistent between sign-ins. Which type of virtual desktop should you use?

Question 4

What feature would you use to make individual personal virtual desktop VMs highly available?

- NLB
- SQL clustering
- Failover clustering
- VM replication

Module 09 lab and review

Lab: Implementing RDS in Windows Server

Scenario

You have been asked to configure a basic Remote Desktop Services (RDS) environment as the starting point for the new infrastructure that will host the sales application. You would like to deploy RDS services, perform initial configuration, and demonstrate to the delivery team how to connect to an RDS deployment.

You are evaluating whether to use user profile disks for storing user profiles and making the disks available on all servers in the collection. A coworker reminded you that users often store unnecessary files in their profiles, and you need to explore how to exclude such data from the profile and set a limit on the profile size.

As the sales application will publish on the RD Web Access site, you also have to learn how to configure and access RemoteApp programs from the Remote Desktop Web Access (RD Web Access) portal.

You been tasked with creating a proof of concept (POC) for a virtual machine (VM)—based session deployment of Virtual Desktop Infrastructure (VDI). You will create a virtual desktop template on a preexisting Microsoft Hyper-V VM manually with a few optimizations.

Objectives

After completing this lab, you'll be able to:

- Implement RDS
- Configure session collection settings and use RDS
- Configure virtual desktop template

Estimated Time: 90 minutes

Module review

Use the following questions to review what you've learned in this module.

Question 1

Which Remote Desktop Service (RDS) role service tracks user sessions across multiple Remote Desktop Session Host (RD Session Host) servers and virtual desktops?

- RD Session Host
- Remote Desktop Virtualization Host (RD Virtualization Host)
- Remote Desktop Connection Broker (RD Connection Broker)
- Remote Desktop Web Access (RD Web Access)
- Remote Desktop Gateway (RD Gateway)

Question 2

Can you connect to RDS only from a Windows-based computer?

Question 3

In which tool can you publish RemoteApp programs on a Remote Desktop Session Host (RD Session Host) server?

Question 4

You are creating a new virtual desktop template for a group of users. You have created and configured the virtual machine (VM). You've also optimized the VM appropriately for use as a virtual desktop. What is the last step in preparing a virtual desktop template?

Question 5

Which port must you allow on your firewall to enable external clients to use RD Gateway to connect to internal RDS resources?

Answers

Question 1

You're deploying a session-based Remote Desktop Services (RDS) deployment on your company's network. Users will only be connecting from your on-premises network. Which roles would you need to deploy? Choose all that apply.

- Remote Desktop Connection Broker (RD Connection Broker)
- Remote Desktop Gateway (RD Gateway)
- Remote Desktop Web Access (RD Web Access)
- Remote Desktop Virtualization Host (RD Virtualization Host)
- Remote Desktop Session Host (RD Session Host)

Explanation

You would need to deploy RD Connection Broker, RD Web Access, and RD Session Host. The RD Gateway role is only necessary if you have external users that need to connect to RDS, and the RD Virtualization Host is used for VM-based deployments.

Question 2

How is RDS different from the Remote Desktop feature?

You can enable the Microsoft Remote Desktop feature on a Windows client and on a server operating system (OS). RDS is a server role, and you can add it only to the Windows Server OS. Remote Desktop on a Windows client allows only a single session; on Windows Server, it allows two sessions. Conversely, RDS supports as many connections as you have hardware resources and RDS client access licenses (CALs). RDS provides many additional features, such as RemoteApp programs, RD Web Access, RD Gateway, and VM-based sessions (Virtual Desktop Infrastructure). These features are not available when you enable only the Remote Desktop feature. An enhanced client experience, advanced device redirection, and media redirection is only available with RDS.

Question 3

Is RD Gateway required if you want to enable Internet clients to connect to your internal RDS resources?

No, RD Gateway is not required. You can configure an external firewall to allow RDP connections to internal RDP resources. However, this is a not very secure solution, and you should avoid using it. Because RD Gateway provides an additional layer of security, we strongly recommend implementing it when you need to enable Internet clients to connect to your internal RDS resources.

Question 1

Can you use Windows Internal Database (WID) with a highly available Remote Desktop Connection Broker (RD Connection Broker) configuration?

WID is used when you have a single RD Connection Broker server in your Remote Desktop Services (RDS) deployment. However, when you configure RD Connection Broker for high availability, the RD Connection Broker database must be stored on a computer that is running SQL Server or in an Azure SQL database.

Question 2

Why would you use a RemoteApp program instead of a remote desktop session?

A remote desktop provides you with the full desktop of a remote server, while a RemoteApp program offers only an application window. RemoteApp programs integrate with local desktops and provide the same user experience as locally installed applications, while a remote desktop adds an additional desktop, which can sometimes be confusing.

Question 3

You want to install the RD Connection Broker role service on a server named "SEA-RDS2", but only a server named "SEA-RDS1" displays in Server Manager when you need to specify the RD Connection Broker server. What should you do to add "SEA-RDS2" as a possible selection?

Server Manager can install RDS role services only on the servers of which it is aware. You should first add "SEA-RDS2" to Server Manager, to the "All Servers" node, and then start the RDS deployment process again.

Question 1

Which two types of virtual desktops can you create (or use) in a Remote Desktop Service (RDS) deployment? Choose all that apply.

- Pooled virtual desktops
- Virtual machine (VM) desktops
- Server desktops
- Personal virtual desktops
- Windows virtual desktops

Explanation

You can use both Personal and Pooled virtual desktops. They are based on VMs that run on a server running Microsoft Hyper-V. All the other answers are incorrect.

Question 2

What is the reason for disabling unnecessary services in a virtual desktop template?

Disabling unnecessary services reduces the resources that are used by each personal or pooled virtual desktop that is created from the virtual desktop template.

Question 3

You would like to provide your users with the option to customize their virtual desktop, and have those customizations be persistent between sign-ins. Which type of virtual desktop should you use?

The best option for personalization is personal virtual desktops. With personal virtual desktops, you can assign users permissions to customize their own personal virtual desktop, including applications.

Question 4

What feature would you use to make individual personal virtual desktop VMs highly available?

- NLB
- SQL clustering
- Failover clustering
- VM replication

Explanation

Failover clustering is the correct answer. To make individual VMs highly available for personal virtual desktops, you need to configure Remote Desktop Virtualization Host (RD Virtualization) Host servers as nodes in a failover cluster. All other answers are incorrect.

Question 1

Which Remote Desktop Service (RDS) role service tracks user sessions across multiple Remote Desktop Session Host (RD Session Host) servers and virtual desktops?

- RD Session Host
- Remote Desktop Virtualization Host (RD Virtualization Host)
- Remote Desktop Connection Broker (RD Connection Broker)
- Remote Desktop Web Access (RD Web Access)
- Remote Desktop Gateway (RD Gateway)

Explanation

RD Connection Broker is the correct answers. Its role service manages connections to RemoteApp programs and virtual desktops, and it directs client connection requests to an appropriate endpoint. It also provides session reconnection and session load balancing. All the other answers are incorrect.

Question 2

Can you connect to RDS only from a Windows-based computer?

No. You can connect to RDS from any device that has a Remote Desktop Protocol (RDP) client, regardless of whether it's running the Windows operating system or any other operating system (OS), or whether the device is a domain member.

Question 3

In which tool can you publish RemoteApp programs on a Remote Desktop Session Host (RD Session Host) server?

You cannot publish RemoteApp programs on an individual RD Session Host server. You can only publish them per session collection, which means that they will publish for all RD Session Host servers in that collection. You can publish RemoteApp programs by using Server Manager or Windows PowerShell.

Question 4

You are creating a new virtual desktop template for a group of users. You have created and configured the virtual machine (VM). You've also optimized the VM appropriately for use as a virtual desktop. What is the last step in preparing a virtual desktop template?

The last step in preparing a virtual desktop template is to run Sysprep and shut down the VM.

Question 5

Which port must you allow on your firewall to enable external clients to use RD Gateway to connect to internal RDS resources?

Clients connect to RD Gateway by using the HTTPS protocol, which uses Transmission Control Protocol (TCP) port 443 by default.



Module 10 Remote Access and web services in Windows Server

Overview of RAS in Windows Server

Lesson overview

Many organizations have mobile users or users working from home that require remote access to files and apps. The Remote Access server role in Windows Server provides multiple role services that you can use to facilitate remote access. You can use DirectAccess and virtual private network (VPN) to provide remote access to data and services on internal networks. Depending on your requirements, you should consider whether to deploy a public key infrastructure (PKI) to distribute certificates to users and computers for authentication. If you have web-based apps, you can implement Web Application Proxy for secure access.

Lesson objectives

After completing this lesson, you'll be able to:

- Describe remote access features in Windows Server.
- Describe considerations for remote app access.
- Install and manage the Remote Access role.
- Manage remote access in Windows Server.
- Describe the considerations for when to deploy a PKI.
- Describe the purpose of the Web Application Proxy.
- Explain authentication options for Web Application Proxy.
- Publish web apps with Web Application Proxy.
- Identify remote access options for an organization.

Remote Access features in Windows Server

The Remote Access server role in Windows Server provides multiple remote access options. Each option represents a unique technology that organizations can use to access internal resources from offices in remote site locations or from the internet. The technology that they use depends on their different business scenarios.

DirectAccess

DirectAccess enables remote users to securely access corporate resources such as email servers, shared folders, and internal websites, without connecting to a virtual private network (VPN). DirectAccess also provides increased productivity for a mobile workforce by offering the same connectivity experience both inside and outside the office.

The main benefit of using DirectAccess is that it provides seamless connectivity back to internal resources. Users don't need to initiate a connection; the connection is created automatically even before the user signs in. This always-on functionality allows you to manage remote computers similar to how you manage computers that are on the internal network.

Note: Only Windows 10 Enterprise and Education editions support DirectAccess. Other editions of Windows 10 do not support DirectAccess.

VPN

VPN connections enable users who are working offsite (for example, from home, a customer site, or a public wireless access point) to access apps and data on an organization's private network 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 a computer, the VPN client, and an organization's server. The exact infrastructure of the shared or public network is irrelevant because it appears to the user as if the data is sent over a dedicated private link.

Routing

Windows Server can function as a router or network address translation (NAT) device between two internal networks, or between the internet and the internal network. Routing works with routing tables and supports routing protocols such as Routing Information Protocol (RIP) version 2, Internet Group Management Protocol (IGMP), and Dynamic Host Configuration Protocol (DHCP) Relay Agent. Although you can use Windows Server for these routing tasks, it's uncommon to do so, because most organizations have specialized hardware devices to perform these tasks.

Web Application Proxy

Web Application Proxy provides reverse proxy functionality for users who must access their organization's internal web applications from the internet. Web Application Proxy preauthenticates users by using the following options:

- Active Directory Federation Services (AD FS) technology, where Web Application Proxy acts as an AD FS proxy.
- Pass-through authentication, where the published application, not Web Application Proxy, performs authentication.

Overview of remote application access

Remote application access is an important part of supporting mobile users and users in remote offices. How you provide remote access to apps varies depending on the architecture of the app. However, for all apps, you need to ensure that remote access to the app is secure.

Remote access to data files

When you use a virtual private network (VPN) or DirectAccess to access data files like Microsoft Word documents or Microsoft Excel spreadsheets, it takes longer to open and close files than when you're in the office, but performance is typically acceptable. The slower performance is primarily because of slower network speeds at remote locations and over the internet. If performance is too slow, you always have the option to copy it locally, work on the file, and then copy the file back to the server on premises.

Remote access to desktop apps

For apps that use shared data storage, such as a database, using a VPN or DirectAccess often causes slow performance. Most developers do not optimize their apps to run over slower connections with high latency. Thus, the apps have many communication calls with the backend data storage. The additional latency on each call adds up to very slow performance.

To support apps with shared data storage it is common to implement Remote Desktop Services (RDS). When you implement RDS, the app is installed on a Remote Desktop Session Host (RD Session Host) located on the internal network that is shared by multiple users. The app remains close to the data, so network latency doesn't cause performance issues. Users connect to the RD Session Host by using the Remote Desktop client which uses Remote Desktop Protocol (RDP). RDP is optimized to send screen drawing information over slower, and higher latency, network connections.

To secure remote access to RD Session Hosts, you can implement a Remote Desktop Gateway (RD Gateway). The RD Gateway is a reverse proxy that you host in a perimeter network specifically for the RDP protocol. The RDP packets are encapsulated in HTTPS packets between Remote Desktop clients and the RD Gateway. This encrypts the data and facilitates transport because HTTPS is not blocked on public networks. You can also secure Access to RD Session Hosts by using a VPN or RemoteAccess.

Remote access to web-based apps

Web-based apps have good performance on slower and higher latency networks. This is because the application logic is stored on a web server that is close to the application data. Only a limited amount of data to be displayed on screen is sent to the web browser. This means that web-based apps are well suited to use by mobile users and remote offices.

The HTTPS protocol which encrypts communication is typically used for web-based apps. This ensures that data is not intercepted while in transit, but most companies also require the web-based app to be isolated from the internet by a reverse proxy. Remote users communicate with a reverse proxy in a perimeter network and the reverse proxy communicates with the web-based app on the internal network. Web Application Proxy functions as a reverse proxy for web-based apps.

If all users are using a VPN or DirectAccess, then you don't need to implement a reverse proxy for the web-based app, because both a VPN and DirectAccess provide a secure way to gain access to the internal network.

Manage remote access in Windows Server

After you install the Remote Access role on a server that is running Windows Server, you can manage the role by using the **Microsoft Management Console (MMC)**, and by using Windows PowerShell. You can use the **MMC** for your day-to-day tasks of managing remote access, and you can use Windows PowerShell for managing multiple servers and for scripting or automating management tasks.

There are two **MMCs** for managing the Remote Access server role: the **Remote Access Management console**, and the **Routing and Remote Access console**. You can access these consoles from the **Tools** menu in **Server Manager**.

Remote Access Management console

The **Remote Access Management console** allows you to manage DirectAccess, virtual private networks (VPN), and Web Application Proxy. When you open this console for the first time, you'll use a wizard-based interface to configure remote access settings according to your business requirements. After you configure the initial remote access settings, you can manage your remote access solution with the following options in the console:

- **Configuration.** You can edit the remote access settings by using wizards and by using the graphical representation of the current network configuration in the console.
- **Dashboard.** You can monitor the overall status of servers and clients that are part of your remote access solution.
- **Operations status.** You can access detailed information on the status of the servers that are part of your remote access solution.
- **Remote Client Status.** You can access detailed information on the status of the clients that are connecting to your remote access solution.
- **Reporting.** You can generate historical reports on different parameters, such as remote access usage, access details, connection details, and server load statistics.

Routing and Remote Access console

You can use the **Routing and Remote Access console** to configure a server running Windows Server as a network address translation (NAT) device, as a router for both IPv4 and IPv6 protocols, as a DHCP proxy, and as a virtual private network (VPN) server. After you complete the configuration, you can manage the remote access solution by using the following options in the console:

- **Server Status.** You can monitor the status of the Remote Access server, the ports in use, and how long the server has been operational (that is, the server uptime).
- **Remote Access Client, Ports, Remote Access Logging.** You can monitor the client status, port status, and detailed logging information about clients that are connected to the Remote Access server.
- **IPv4.** You can configure the IPv4 settings such as NAT, IPv4 routing with static routes, and these routing protocols: RIP version 2, IGMP, and the DHCP Relay Agent.
- **IPv6.** You can configure IPv6 settings, such as IPv6 routing with static routes and the DHCP Relay Agent routing protocol.

Windows PowerShell commands

You can use Windows PowerShell commands in Windows Server to configure remote access and create scripts to automate the configuration and management procedures. Some examples of Windows PowerShell commands for remote access include:

- **Set-DAServer.** Sets the properties specific to the DirectAccess server.
- **Get-DAServer.** Displays the properties of the DirectAccess server.
- **Set-RemoteAccess.** Modifies the configuration that is common to both DirectAccess and VPN, such as Secure Sockets Layer (SSL) certificate, internal interface, and internet interface.
- **Get-RemoteAccess.** Displays the configuration of DirectAccess and VPN (both remote access VPN and site-to-site VPN).

Additional reading: For more information about remote access cmdlets, refer to **RemoteAccess**¹.

When to deploy a public key infrastructure for remote access

You can use digital certificates to verify and authenticate the identity of each party involved in an electronic transaction. Digital certificates also help establish trust between computers and the corresponding applications that are hosted on application servers. Remote access uses certificates to verify the identity of servers and provide encryption. You can also use certificates to verify the identity of users or computers signing in for remote access. You can use Windows Server to build a public key infrastructure (PKI) to issue certificates.

Methods for obtaining certificates

In most cases, you obtain certificates from a certification authority (CA). The most important consideration for a CA is trust. If a certificate is issued by a trusted CA, then that certificate is trusted and can be used for authentication. If a CA is not trusted, then the certificates issued by that CA can't be used for authentication.

To obtain certificates, you can:

- Create your own private CA by using Windows Server. The certificates issued by a private CA are automatically trusted by domain joined Windows clients and servers. However, the certificates issued by an internal CA are not automatically trusted by any devices not joined to the domain.
- Purchase certificates from a public CA. The certificates issued by a public CA are trusted automatically by almost all devices whether they are domain joined or not. Windows does not include tools to automatically deploy certificates from a public CA to users or computers.
- Generate self-signed certificates within some applications. By default, these certificates 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 configured with the **Getting Started Wizard**, which provides easy setup and configuration.

Considerations when planning PKI

To determine whether you should implement an internal PKI for remote access, you need to plan how you will use certificates. If you are using certificates only on a few servers, then the cost of using a public

¹ <https://aka.ms/remotearch-win10-ps>

CA is low. Certificates from a public CA are also beneficial if you expect devices that are not joined to the domain to access the servers.

A private CA is beneficial primarily for remote access when you are issuing certificates to client devices and individual users for authentication. For example, it is common to require a valid computer certificate to allow VPN access as a second level of authentication beyond a username and password. If you are issuing certificates to many computers, then the automatic enrollment provided by a private CA is important. There is also a significant cost savings because you don't need to pay for certificates issued by a private CA.

The following table summarizes the advantages and disadvantages of certificates issued by private and public CAs.

Table 1: Advantages and Disadvantages of certificates by issuer

CA type	Advantages	Disadvantages
Private CA	A private CA provides greater control over certificate management and has a lower cost when compared to a public CA. There is no cost per certificate. You also have the option to use customized templates and automatic enrollment.	By default, certificates by private CAs are not trusted by external clients (web browsers and operating systems), and require greater administration.
Public CA	A certificate issued by a public CA is trusted by many external clients (web browsers and operating systems), and requires minimal administration.	The cost is higher when compared to a private CA. Cost is based per certificate. Certificate procurement is also slower.

What is Web Application Proxy?




Figure 1: Web Application Proxy controlling communication

Web Application Proxy is a Remote Access role service. This role service functions as a reverse web proxy, and provides users located on the internet with access to internal corporate web applications or Remote Desktop Gateway servers. Web Application Proxy can use Active Directory Federation Services (AD FS) to preauthenticate internet users and acts as an AD FS proxy for publishing claims-aware applications. A claims-aware application can use any information about a user such as group membership, email address, department, or company as part of user authorization.

Before you install Web Application Proxy, you must deploy AD FS as a prerequisite. Web Application proxy uses AD FS for authentication services. One feature provided by AD FS is single sign-on (SSO) functionality, which means that if users enter their credentials for accessing a corporate web application

once, they will not be asked to enter their credentials again for subsequent access to the corporate web application. You can also use AD FS to authenticate users at Web Application Proxy before users communicate with the application.

Placing the Web Application Proxy server in the perimeter network between two firewall devices is a typical configuration. The AD FS server and applications that are published are located on the corporate network, and together with domain controllers and other internal servers, are protected by the second firewall. This scenario provides secure access to corporate applications for users located on the internet, and at the same time protects the corporate IT infrastructure from security threats on the internet.

Authentication options for Web Application Proxy

When you configure an application in Web Application Proxy, you need to select the type of preauthentication. You can choose either Active Directory Federation Services (AD FS) preauthentication or pass-through preauthentication. AD FS preauthentication provides more features and benefits, but pass-through preauthentication is compatible with all web apps.

AD FS preauthentication

AD FS preauthentication uses AD FS for web applications that use claims-based authentication. When a user initiates a connection to the corporate web application, the first entry point the user connects to is the Web Application Proxy. Web Application Proxy preauthenticates the user in the AD FS server. If the authentication is successful, Web Application Proxy establishes a connection to the web server in the corporate network where the application is hosted.

By using AD FS preauthentication, you ensure that only authorized users can send data packets to the web application. This prevents hackers from taking advantage of web-app flaws before authentication. AD FS preauthentication significantly reduces the attack surface for a web app.

Pass-through preauthentication

Pass-through preauthentication doesn't use AD FS for authentication, nor does Web Application Proxy preauthenticate the user. Instead, the user is connected to the web application through Web Application Proxy. The web application proxy rebuilds the data packets as they are delivered to the web app, which provides protection from flaws such as malformed packets. However, the data portion of the packet passes to the web app. The web app is responsible for authenticating users.

AD FS preauthentication benefits

AD FS preauthentication provides the following benefits over pass-through preauthentication:

- **Workplace join.** *Workplace join* allows devices that are not members of the Active Directory domain, such as smartphones, tablets, or non-company laptops, to be added to a workplace. After these non-domain devices are added to the workplace, you can use the workplace join status as part of AD FS preauthentication.
- **Single sign-on (SSO).** *SSO* allows users that are preauthenticated by AD FS to enter their credentials only once. If users subsequently access other applications that use AD FS for authentication, they won't be prompted again for their credentials.

- Multifactor authentication. *Multifactor authentication* allows you to configure multiple types of credentials to strengthen security. For example, you can configure the system so that users enter their username and password together with a smart card.
- Multifactor access control. *Multifactor access control* is used in organizations that want to strengthen their security when publishing web applications by implementing authorization claim rules. The rules are configured so that they issue either a permit, or a deny claim, which determines whether a user or a group is allowed or denied access to a web application that is using AD FS preauthentication.

Publish applications with Web Application Proxy

After the Web Application Proxy role service is installed, you configure it by using the **Web Application Proxy Configuration Wizard** from the **Remote Access Management console**. When the **Web Application Proxy Configuration Wizard** completes, it creates the **Web Application Proxy console**, which you can use for further management and configuration of Web Application Proxy.

The **Web Application Proxy Configuration Wizard** requires that you enter the following information during the initial configuration process:

- AD FS name. To locate this name, open the **AD FS Management** console, and, under **Edit Federation Service Properties**, find the value in the **Federation Service** name box.
- Credentials of local administrator account for AD FS.
- AD FS Proxy Certificate. This is a certificate that Web Application Proxy will use for AD FS proxy functionality.

Note: The AD FS proxy certificate must contain the AD FS name in the subject field of the certificate, because the **Web Application Proxy Configure Wizard** requires it. Additionally, the subject alternative names field of the certificate should include the AD FS name.

After completing the **Web Application Proxy Configuration Wizard**, you can publish your web app by using **Web Application Proxy console** or Windows PowerShell cmdlets. The Windows PowerShell cmdlets for managing published apps are:

- **Add-WebApplicationProxyApplication**
- **Get-WebApplicationProxyApplication**
- **Set-WebApplicationProxyApplication**

When you publish your web app, you must provide the following information:

- The type of preauthentication, for example, pass-through
- The application to publish
- The external URL of the application, for example, `https://lon-svr1.adatum.com`
- A certificate whose subject name covers the external URL, for example, `lon-svr1.adatum.com`
- The URL of the backend server, which is entered automatically when you enter the external URL

Additional reading: For additional guidance from Microsoft for publishing complex applications, refer to **Publishing Applications with SharePoint, Exchange and RDG²**.

² <http://aka.ms/Qopw7d>

Discussion: Remote access options usage scenarios

In this discussion, you'll consider the following scenario and discuss the questions that follow.

Scenario

Remote access technologies provide various solutions that allow secure access to an organization's infrastructure from different locations. While organizations usually own and protect local area networks (LANs) entirely 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.

Questions

Consider the following questions in your discussion:

- Do you allow users to connect to your network resources remotely? If so, how?
- What are your business requirements for using remote access?

Test your knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

Which remote access feature in Windows Server automatically connects clients to the internal network when outside the office?

- Network Policy Server (NPS)
- Web Application Proxy
- Router
- DirectAccess

Question 2

Which information is required when publishing a web app with Web Application Proxy? (Choose three.)

- External URL of the applications
- A certificate from a private certification authority
- URL of the backend server
- Type of preauthentication

Implementing VPNs

Lesson overview

virtual private networks (VPNs) provide secure access to the internal data and applications that organizations provide for clients and devices that are using the internet. If you want to implement and support a VPN environment within your organization, you must understand how to select a suitable tunneling protocol, configure VPN authentication, and configure the server role to support your chosen configuration.

One of the advantages of using VPN compared to other remote access technologies are that VPN supports different kinds of devices. These devices include mobile devices, tablet devices, computers that are not domain members, workgroup computers, and computers that are running nonenterprise versions of Windows 10 and Windows 8.1 operating systems.

Lesson objectives

After completing this lesson, you will be able to:

- Describe the various VPN scenarios
- Understand site-to-site VPN
- Describe the options for VPN tunneling protocols
- Describe the VPN authentication options
- Describe the **VPN Reconnect** feature
- Describe how to configure a VPN by using the **Getting Started Wizard**

VPN scenarios

Similar to previous Windows Server versions, Windows Server 2019 supports two types of virtual private network (VPN) connections:

- Remote access VPN connection
- Site-to-site

Remote access VPN connections

Remote access VPN connections enable users who work offsite, such as at home, at a customer site, or from a public wireless-access point, to access a server on your organization's private network by using the infrastructure that a public network provides, such as the internet. 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 displays as though you sent the data over a dedicated private link.

Site-to-site VPN connections

Site-to-site VPN connections, or router-to-router VPN connections, enable your organization to establish routed connections between separate offices or with other organizations over a public network, while helping to maintain secure communications.

Properties of VPN connections

VPN connections that use the Point-to-Point Tunneling Protocol (PPTP), Layer 2 Tunneling Protocol with Internet Protocol Security (L2TP/IPsec), Secure Socket Tunneling Protocol (SSTP) and Internet Key Exchange version 2 (IKEv2) have the following properties:

- Encapsulation. VPN technology encapsulates private data with a header that contains routing information, which allows the data to traverse the transit network.
- Authentication. There are three types of authentication for VPN connections, including:
 - User-level authentication by using Point-to-Point Protocol (PPP) authentication. To establish the VPN connection, the VPN server authenticates the VPN client that is attempting to connect by using a PPP user-level authentication method. It then verifies that the VPN client has the appropriate authorization. If you use mutual authentication, the VPN client also authenticates the VPN server.
 - Computer-level authentication by using Internet Key Exchange (IKE). To establish an IPsec security association, the VPN client and the VPN server use the IKE protocol to exchange computer certificates or a pre-shared key. In either case, the VPN client and server authenticate each other at the computer level. We recommend computer-certificate authentication because it is a much stronger authentication method than a pre-shared key. Please note, however, that computer-level authentication occurs only for L2TP/IPsec connections.
 - Data-origin authentication and data integrity. To verify that the data sent on a VPN connection originated at the connection's other end and was not modified in transit, the data contains a cryptographic checksum that is based on an encryption key known only to the sender and the receiver. Note that data-origin authentication and data integrity are available only for L2TP/IPsec connections.
- Data encryption. To ensure data confidentiality as it traverses the shared or public transit network, the sender encrypts the data, and the receiver decrypts it. The encryption and decryption processes depend on the sender and the receiver both using a common encryption key.

Packets that are intercepted in the transit network are unintelligible to anyone who does not have the common encryption key. The encryption key's length is an important security parameter. Therefore, it is important to use the largest possible key size to ensure strong data encryption and confidentiality. However, stronger encryption consumes more central processing unit (CPU) resources. Therefore, organizations should plan for hardware resources if they plan to require stronger encryption.

Site-to-site VPN

About site-to-site VPNs

A site-to-site VPN (virtual private network) connection connects two portions of a private network. The VPN server provides a routed connection to the network to which the VPN server attaches. The calling router, which is the VPN client, authenticates itself to the answering router, which is the VPN server. For mutual authentication, the answering router authenticates itself to the calling router. In a site-to-site VPN connection, the packets sent from either router across the VPN connection do not typically originate at the routers.

When you create a demand-dial interface, you specify the same information as you would when creating a VPN profile. Furthermore, you must specify the credentials used to connect to the answering router.

The name of the answering router's demand-dial interface must match the name of the user account that the calling router specifies.

When you configure a site-to-site VPN, you can create a one-way connection or a two-way connection. If you configure a one-way connection, one VPN server always initiates the connection, and one VPN server always answers. If you configure a two-way connection, either of your VPN routers can initiate the connection, and either can function as the calling or answering router.

You can restrict a calling router from initiating unnecessary connections by using demand-dial filtering or dial-out hours. You can use demand-dial filtering to configure the type of traffic that can initiate a connection, or you can specify the traffic that can't initiate a connection. You do this by right-clicking or accessing the context menu at the demand-dial interface in **Routing and Remote Access**, and then selecting **Set IP Demand-dial Filters**. You also can configure times during which a calling router can, or can't, initiate a connection. You do this by right-clicking or accessing the context menu at the demand-dial interface and then selecting **Dial-out Hours**.

A routed VPN connection across the internet operates logically as a dedicated wide area network (WAN) link. When networks connect over the internet, a router forwards packets to another router across a VPN connection. To the routers, the VPN connection operates as a data-link layer link.

Types of site-to-site VPNs

You can create three types of site-to-site VPNs in Windows Server, including:

- PPTP (Point-to-Point Tunneling Protocol), which uses the Microsoft Point-to-Point Encryption (MPPE) for encryption and the PPP (Point-to-Point Protocol) protocol for authentication.
- L2TP (Layer 2 Tunneling Protocol), which uses certificates for encryption, integrity, and data authentication, and PPP for authentication.
- IKEv2 (Internet Key Exchange version 2), which uses Advanced Encryption Standard (AES) 256, AES 192, AES 128, and Triple Data Encryption Standard (3DES) for encryption.

Additionally, a site-to-site VPN connection can be persistent or on-demand:

- **On-demand VPN Connection.** When traffic is being forwarded to the remote location, a site-to-site VPN connection occurs. When the transfer completes, the connection closes shortly thereafter, contingent on the configuration for your remote access policy. You also can configure the calling router (VPN client) to close the connection after a specified idle timeout interval. You can configure this in the properties of the demand-dial interface.
- **A persistent VPN Connection.** A persistent site-to-site VPN has a constant connection. Additionally, if the connection inadvertently closes or drops, it immediately establishes again. To configure the connection as persistent, on the **Properties** page of the demand-dial interface, on the **Options** tab, select **Persistent connection**. You can also configure this on the answering router by clearing the **Idle Timeout** and **Session Timeout** boxes on the network policy's **Constraints** tab.

You must create a demand-dial interface on the calling router. This interface is a VPN profile that connects to the answering router.

Options for VPN tunneling protocols

PPTP (Point-to-Point Tunneling Protocol), L2TP (Layer 2 Tunneling Protocol), and SSTP (Secure Socket Tunneling Protocol) depend heavily on the features that you specified originally for PPP (Point-to-Point Protocol), which sends data across dial-up or dedicated point-to-point connections. For IP, PPP

encapsulates IP packets within PPP frames, and then transmits the encapsulated PPP packets across a point-to-point link. PPP was originally the protocol used between a dial-up client and a network access server.

PPTP

You can use PPTP for remote access and site-to-site VPN (virtual private network) connections. When you use the internet as the VPN public network, the PPTP server is a PPTP-enabled VPN server that has one interface on the internet and one on your intranet.

PPTP enables you to encrypt and encapsulate multiprotocol traffic in an IP header that it then sends across an IP network or a public IP network, such as the internet:

- **Encapsulation.** PPTP encapsulates PPP frames in IP datagrams for network transmission. PPTP uses a TCP (Transmission Control Protocol) connection for tunnel management and a modified version of Generic Route Encapsulation (GRE) to encapsulate PPP frames for tunneled data. You can encrypt and compress payloads of the encapsulated PPP frames.
- **Encryption.** You can encrypt the PPP frame with MPPE (Microsoft Point-to-Point Encryption) by using encryption keys that are generated from the Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAPv2) or Extensible Authentication Protocol-Transport Layer Security (EAP-TLS) authentication process. VPN clients must use the MS-CHAPv2 or EAP-TLS authentication protocol to ensure encryption of payloads of PPP frames. PPTP uses the underlying PPP encryption and encapsulates a previously encrypted PPP frame.

L2TP

L2TP enables you to encrypt multiprotocol traffic that is sent over any medium that supports point-to-point datagram delivery, such as IP or asynchronous transfer mode (ATM). L2TP is a combination of PPTP and Layer 2 Forwarding (L2F). L2TP represents the best features of PPTP and L2F.

Unlike PPTP, the Microsoft implementation of L2TP does not use MPPE to encrypt PPP datagrams. L2TP relies on IPsec in transport mode for encryption services. The combination of L2TP and IPsec is L2TP/IPsec (Layer 2 Tunneling Protocol with Internet Protocol Security).

To use L2TP/IPsec, both the VPN client and server must support L2TP and IPsec. Windows Server 2012, Windows 8.1 and newer Windows operating systems remote access clients include client support for L2TP. Windows Server 2012 or later operating systems all have VPN server support for L2TP.

The encapsulation and encryption methods for L2TP are as follows:

- **Encapsulation.** Encapsulation for L2TP/IPsec packets consists of two layers, L2TP encapsulation and IPsec encapsulation. L2TP encapsulates and encrypts data as follows:
 - **First layer.** The first layer is the L2TP encapsulation. A PPP frame (an IP datagram) is wrapped with an L2TP header and a User Datagram Protocol (UDP) header.
 - **Second layer.** The second layer is the IPsec encapsulation. The resulting L2TP message is wrapped with an IPsec Encapsulating Security Payload (ESP) header and trailer, an IPsec Authentication trailer that provides message integrity and authentication, and a final IP header. The IP header contains the source and destination IP address that corresponds to the VPN client and server.

- Encryption. The L2TP message is encrypted with AES (Advanced Encryption Standard) or 3DES (Triple Data Encryption Standard) by using encryption keys that the IKE negotiation process generates.

SSTP

SSTP is a tunneling protocol that uses the HTTPS protocol over TCP port 443 to pass traffic through firewalls and web proxies, which otherwise might block PPTP and L2TP/IPsec traffic. SSTP provides a mechanism to encapsulate PPP traffic over the SSL channel of the HTTPS protocol. The use of PPP allows support for strong authentication methods, such as EAP-TLS. SSL provides transport-level security with enhanced key negotiation, encryption, and integrity checking.

When a client tries to establish an SSTP-based VPN connection, SSTP first establishes a bidirectional HTTPS

layer with the SSTP server. The protocol packets flow over this HTTP layer as the data payload by using the

following encapsulation and encryption methods:

- Encapsulation. SSTP encapsulates PPP frames in IP datagrams for transmission over the network. SSTP uses a TCP connection (over port 443) for tunnel management and as PPP data frames.
- Encryption. SSTP encrypts the message with the SSL channel of the HTTPS protocol.

IKEv2

IKEv2 (Internet Key Exchange version 2) 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 the following strong authentication and encryption methods:

- Encapsulation. IKEv2 encapsulates datagrams by using IPsec ESP or Authentication Header (AH) for transmission over the network.
- Encryption. IKEv2 encrypts the message with one of the following protocols by using encryption keys that it generates during the IKEv2 negotiation process: AES 256, AES 192, AES 128, and 3DES encryption algorithms.

IKEv2 is supported only on computers that are running Windows Server 2019, Windows Server 2016, Windows 10, and Windows 8.1 operating systems. IKEv2 is the default VPN tunneling protocol in Windows 10.

Note: You should not use PPTP because of security vulnerabilities. L2TP is an old VPN protocol. Instead, use IKEv2 whenever possible because it is more secure and offers advantages over L2TP.

VPN 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. The

Remote Access server role supports the methods that the following sections describe.

PAP

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. Windows Server includes PAP to support older client operating systems that support no other authentication method.

CHAP

The Challenge Handshake Authentication Protocol (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. However, because CHAP requires that you use a reversibly encrypted password, you should consider using another authentication protocol, such as MS-CHAPv2.

MS-CHAPv2

MS-CHAPv2 (Microsoft Challenge Handshake Authentication Protocol version 2) is a one-way, encrypted password, mutual-authentication process that works as follows:

1. The authenticator, which is the Remote Access server or computer that is running Network Policy Server (NPS), sends a challenge to the remote access client. The challenge consists of a session identifier and an arbitrary challenge string.
2. The remote access client sends a response that contains a one-way encryption of the received challenge string, the peer challenge string, the session identifier, and the user password.
3. The authenticator checks the response from the client, and then sends back a response that contains an indication of the connection attempt's success or failure and an authenticated response based on the sent challenge string, the peer challenge string, the client's encrypted response, and the user password.
4. The remote access client verifies the authentication response and, if correct, uses the connection. If the authentication response is not correct, the remote access client terminates the connection.

EAP

If you use EAP (Extensible Authentication Protocol), an arbitrary authentication mechanism authenticates a remote access connection. The remote access client and the authenticator, which is either the Remote Access server or the Remote Authentication Dial-In User Service (RADIUS) server, negotiate the exact authentication scheme they will use. Routing and Remote Access includes support for EAP-TLS (Extensible Authentication Protocol-Transport Layer Security) by default. You can plug in other EAP modules to the server that is running Routing and Remote Access to provide other EAP methods.

Note: RADIUS is an industry-standard authentication protocol that many vendors use to support the exchange of authentication information between elements of a remote-access solution. NPS is the Microsoft implementation of a RADIUS server. You will learn more about this in Lesson 2: Implementing Network Policy Server.

Note: We strongly recommend that you disable the PAP and CHAP authentication protocols, because they are insecure when compared to the MS-CHAPv2 and EAP authentication protocols.

Other authentication options

You can enable additional options when you select an authentication method, including:

- **Unauthenticated Access.** This is not an authentication method, but it is the lack of an authentication method. Unauthenticated access allows remote systems to connect without authentication. You should never enable this option in a production environment, as it leaves your network at

risk. However, this option can be useful for troubleshooting authentication issues in a test environment.

- **Allow machine certificate authentication for IKEv2.** Select this option if you want to use VPN Reconnect.

What is VPN Reconnect?

In dynamic business scenarios, users must be able to access data securely at any time, from anywhere, and be able to access it continuously without interruption. For example, users might want to securely access data that is on the company's server, from a branch office, or while they are traveling.

Therefore, to meet this requirement, you can configure the **VPN Reconnect** feature that is available in Windows Server 2012 or later, Windows 10, and Windows 8.1. This feature allows users to access organizational data by using a VPN (virtual private network) connection, which reconnects automatically if connectivity inadvertently disconnects. **VPN Reconnect** also enables roaming between different networks.

VPN Reconnect uses the IKEv2 (Internet Key Exchange version 2) technology to provide seamless and consistent VPN connectivity. Users who connect via a wireless mobile broadband will benefit most from this capability. Consider a scenario where you are using a laptop that is running Windows 10. When you travel to work by train, you connect to the internet with a wireless mobile broadband card and then establish a VPN connection to your organization's network. When the train passes through a tunnel, you lose the internet connection. After the train emerges from the tunnel, the wireless mobile broadband card automatically reconnects to the internet. With earlier versions of Windows client and server operating systems, the VPN did not reconnect automatically. Therefore, you would have had to repeat the multistep process of connecting to the VPN manually. This was time-consuming and frustrating, as it provided only intermittent connectivity.

However, with **VPN Reconnect**, clients that are running Windows Server 2019, Windows Server 2016, and Windows 10 reestablish active VPN connections automatically when the network reestablishes internet connectivity. Even though the reconnection might take several seconds, you do not need to reconnect manually or authenticate again to access internal network resources.

VPN Reconnect requirements

The system requirements for using the **VPN Reconnect** feature include:

- A computer that is running Windows Server 2012 or later as a VPN server.
- A computer that is running Windows 10, Windows 8.1, Windows Server 2012 or later, as a VPN client.
- A public key infrastructure (PKI), because **VPN Reconnect** requires a computer certificate for a remote connection. You can use certificates that an internal certification authority (CA) or a public CA issue.

Configure a VPN by using the Getting Started Wizard

The Getting Started Wizard

You can configure a VPN (virtual private network) solution by using the **Getting Started Wizard** in the **Remote Access Management Console**. You can use this wizard to configure both DirectAccess and VPN, or only DirectAccess or VPN. If you choose to configure VPN only, the **Routing and Remote Access**

console

display allows you to specify VPN configuration settings and deploy the VPN solution.

Before you deploy your organization's VPN solution, you must:

- Ensure that your VPN server has two network interfaces. You must determine which network interface will connect to the internet and which will connect to your private network. During configuration, you must choose which network interface connects to the internet. If you specify the incorrect network interface, your remote-access VPN server will not operate correctly.
- 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 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, the remote-access VPN server can automatically 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 a RADIUS (Remote Authentication Dial-In User Service) server or a remote-access VPN server that you configure to authenticate connection requests from VPN clients. 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 (Network Policy Server) can function as a RADIUS proxy or server.

- Remember that by default, the **Getting Started Wizard** configures Windows authentication for VPN clients.
- Ensure that the person who deploys your VPN solution has the necessary administrative group memberships to install server roles and configure necessary services. These tasks require membership to the local **Administrators** group.

Options for modifying VPN configurations

After you deploy and configure your VPN solution, your server is ready for use as a remote access server. However, you can perform additional tasks on your remote access server, including the ability to:

- Configure static packet filters. Add static packet filters to provide additional network protection.
- Configure services and ports. Choose the services on the private network that you want to make available for remote users.
- Adjust logging levels. Configure the level of event details that you want to log. You can decide which information you want to track in log files.
- Configure the number of VPN ports. Add or remove VPN ports. For example, you might want to increase **L2TP** and remove all PPTP (Point-to-Point Tunneling Protocol) and SSTP (Secure Socket Tunneling Protocol) connections. Configure the ports to support the number of users and the types of connections that you want to allow.
- Create a Connection Manager profile. Manage the client connection experience for users and simplify configuration and troubleshooting of client connections.
- Add Active Directory Certificate Services (AD CS). Configure and manage a CA on a server for use in a PKI (Public Key Infrastructure).

- Increase remote access security. Protect remote users and the private network by implementing methods such as enforcing the use of secure authentication methods and requiring higher levels of data encryption.
- Increase VPN security. Protect remote users and the private network by implementing methods such as requiring the use of secure tunneling protocols and configuring account lockout.
- Implement **VPN Reconnect**. Consider adding **VPN Reconnect** to reestablish VPN connections automatically if you lose your internet connections temporarily.

Demonstration: Configure VPN

In this demonstration you will learn how to:

- Install Routing and Remote Access Server virtual private networks (VPNs) using PowerShell.
- Configure and enable VPN configuration.
- Review the default VPN configuration.

Preparation Steps

For this demonstration, you will use the following virtual machines:

- **WS-011T00A-SEA-DC1**
- **WS-011T00A-SEA-ADM1**
- **WS-011T00A-SEA-SVR1**
- **WS-011T00A-SEA-CL1**

Sign in by using the following credentials:

- User Name: **Contoso\Administrator**
- Password: **Pa55w.rd**

Note: You do not need to sign in to **WS-011T00A-SEA-DC1**.

After completing the demonstration, leave all the virtual machines running. You will use them in a later demonstration.

Demonstration steps

Install Routing and Remote Access Server (VPN) using PowerShell

- On **SEA-ADM1**, open a Windows PowerShell command prompt, enter the following command, and then select Enter:

```
Install-WindowsFeature -name RemoteAccess,Routing -IncludeManagementTools
```

Wait for the command to complete, which should take approximately 1 minute.

Request certificate for SEA-ADM1

1. On **SEA-ADM1**, open a command prompt, enter the following command, and then select **Enter: mmc**

2. Add the Certificates snap-in for the computer account and local computer.
3. In the Certificates snap-in console tree, navigate to **Certificates (local)\Personal**, and then request a new certificate.
4. Under **Request Certificates**, configure the Contoso Web Server certificate with the following setting:
 - Subject name: Under **Common name**, enter `vpn.contoso.com`
 - Friendly name: **Contoso VPN**
5. In the Certificates snap-in, expand **Personal and select Certificates**, and then, in the **details** pane, verify that a new certificate with the name `vpn.contoso.com` is enrolled with Intended Purposes of Server Authentication.
6. Close the **Microsoft Management Console (MMC)**. When you receive a prompt to save the settings, select **No**.

Change the HTTPS bindings

1. Open the **Internet Information Services (IIS) Manager console**.
2. In **Internet Information Services (IIS) Manager**, in the console tree, navigate to **SEA-ADM1/Sites**, and then select **Default Web site**.
3. Configure site bindings by selecting **Contoso VPN as SSL Certificate**. When prompted, select **Yes**.
4. Close the **Internet Information Services (IIS) Manager console**.

Configure and enable VPN configuration

1. On **SEA-ADM1**, open **Routing and Remote Access**.
2. Right-click **SEA-ADM1 (local)** or access the context menu, and then select **Configure and Enable Routing and Remote Access**.
3. On the **Welcome to Routing and Remote Access Server Setup Wizard**, select **Next**.
4. On the **Configuration** page, select **Custom configuration**, and then select **Next**.
5. On the **Custom Configuration** page, select **VPN access and LAN routing**, and then select **Next**.
6. On the **Completing the Routing and Remote Access Server Setup Wizard** page, select **Finish**. When prompted, select **Start service**.

Review the default VPN configuration

1. Expand **SEA-ADM1 (local)**, right-click **Ports** or access the context menu, and then select **Properties**.
2. Verify that 128 ports exist for **Wan Miniport (SSTP)**, **Wan Miniport (IKEv2)**, and **Wan Miniport (L2TP)**. Change the number of ports to five, for each type of connection. Disable the use of **Wan Miniport (PPTP)**.
3. Close the **Ports Properties** dialog box, and when prompted, select **Yes**.
4. Right-click **SEA-ADM1 (local)** or access the context menu, and then select **Properties**.
5. On the **General** tab, verify that **IPv4 Remote access server** is selected.

6. On the **Security** tab, select the drop-down arrow next to **Certificate**, and then select `vpn-contoso.com`.
7. Select **Authentication Methods**, and then verify that **EAP** is selected as the authentication protocol.
8. On the **IPv4** tab, verify that the VPN server is configured to assign IPv4 addressing by using a Dynamic Host Configuration Protocol (DHCP).
9. To close the **SEA-ADM1 (local) Properties** dialog box, select **OK**, and then when you receive a prompt, select **Yes**.

After completing the demonstration, leave all the virtual machines running. You will use them in a later demonstration.

Test your knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

Which two types of VPN (virtual private network) connections does Windows Server 2019 support? Choose two.

- Remote access
- IKEv2 (Internet Key Exchange version 2)
- Point-to-site
- Site-to-site
- VPN reconnect

Question 2

Which types of site-to-site VPNs can you create in Windows Server? Choose three.

- PPTP (Point-to-Point Tunneling Protocol)
- IKEv2
- L2TP (Layer 2 Tunneling Protocol)
- SSTP (Secure Socket Tunneling Protocol)

Question 3

Which VPN tunneling protocol should you use for your mobile users?

- PPTP
- L2TP
- SSTP
- IKEv2

Question 4

You configure your VPN server to use IP addresses from a DHCP server on your private network to assign those addresses to remote clients. How many IP addresses is it going to lease at a time?

- 2
- 10
- 25
- 5

Implementing NPS

Lesson overview

NPS (Network Policy Server) is part of the **Network Policy and Access Services** server role in Windows Server. It enables you to create and enforce organization-wide network access policies for connection request authentication and authorization. You also can use NPS as a RADIUS proxy to forward connection requests to NPS or other RADIUS (Remote Authentication Dial-In User Service) servers that you configure in remote **RADIUS server** groups.

You can use NPS to centrally configure and manage network-access authentication and authorization.

Lesson objectives

After completing this lesson, you will be able to:

- Describe NPS.
- Plan an NPS deployment.
- Describe Connection request processing.
- Configure policies on NPS.
- Describe how to implement RADIUS with NPS.

Overview of NPS

NPS (Network Policy Server) enables you to create and enforce organization-wide network access policies for connection request authentication and authorization. You can also use NPS as a RADIUS proxy to forward connection requests to NPS or other RADIUS (Remote Authentication Dial-In User Service) servers that you configure in remote RADIUS server groups.

You can use NPS to implement network-access authentication, authorization, and accounting with any combination of the following functions:

- RADIUS server
- RADIUS proxy
- RADIUS accounting

RADIUS server

RADIUS is an industry-standard authentication protocol that vendors use to support the exchange of authentication information between elements of a remote-access solution. NPS is the Microsoft implementation of a RADIUS server. NPS enables the use of a heterogeneous set of wireless, switch, remote access, or VPN equipment. You can use NPS with the Routing and Remote Access service, which is available in the Windows Server operating system. In addition, you can use NPS with the **Remote Access** role in Windows Server.

NPS performs centralized connection authentication, authorization, and accounting for wireless, Remote Desktop (RD) Gateway servers, authenticating switches, virtual private networks (VPNs), and dial-up connections. When using NPS as a RADIUS server, you configure network access servers (NASs), such as wireless access points and VPN servers, which are also known as RADIUS clients in NPS. You also configure the network policies that NPS uses to authorize connection requests, and you can configure RADIUS

accounting so that NPS logs accounting information to log files on the local hard disk or in a **Microsoft SQL Server** database.

Important: You can't install NPS on Server Core editions of Windows Server.

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) capability. This means that the same set of user credentials enable network-access control, such as authenticating and authorizing access to a network, and access to 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, when 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 the NPS server will forward to other RADIUS servers and to which RADIUS servers you want to forward connection requests. You can also configure NPS to forward accounting data for logging by one or more computers in a remote RADIUS server group. With NPS, your organization can also outsource its remote-access infrastructure to a service provider, while retaining control over user authentication, authorization, and accounting. You can create NPS configurations for the following solutions:

- VPN servers.
- Wireless access points.
- Remote Desktop (RD) Gateway servers.
- Outsourced VPN, dial-up, or wireless access.
- Internet access.
- Authenticated access to extranet resources for business partners.

RADIUS accounting

You can configure NPS to perform RADIUS accounting for user authentication requests, Access-Accept messages, Access-Reject messages, accounting requests and responses, and periodic status updates. NPS enables you to log to a **Microsoft SQL Server** database in addition to, or instead of, logging to a local file.

Plan NPS deployment

When you implement NPS (Network Policy Server), it performs centralized connection authentication, authorization, and accounting for wireless, Remote Desktop (RD) Gateway servers, authenticating switches, virtual private networks (VPNs) and dial-up connections.

You also configure the network policies that NPS uses to authorize connection requests, and you can configure RADIUS (Remote Authentication Dial-In User Service) accounting so that NPS logs accounting information to log files on the local hard disk or in a **Microsoft SQL Server** database.

After the installation, you can use the **NPS Management** console, Netsh NPS commands, or Windows PowerShell to manage NPS.

Additional reading: For additional information on Netsh NPS commands refer to **Netsh Commands for Network Policy Server in Windows Server 2008**³.

Additional reading: For additional information on NPS PowerShell cmdlets, refer to **Network Policy Server (NPS) Cmdlets in Windows PowerShell**⁴.

NPS configuration

First, you need to decide to which domain the NPS should belong, in case you have multiple domains in your environment. By default, the NPS can authenticate all users in its own domain and all trusted domains. To grant the NPS permission to read the dial-in properties of user accounts, you must add the computer account of the NPS to the **RAS** (Remote Access Service) and **NPS** groups for each domain.

NPS clients

When using NPS as a RADIUS server, you configure network access servers (NASs), such as wireless access points supporting 802.1X, Remote Desktop (RD) Gateway servers, 802.1X authenticating switches, virtual private networks (VPNs), and dial-up connections, as RADIUS clients in NPS.

Note: Client computers that use VPN servers are not RADIUS clients, only NASs that support the RADIUS protocol are RADIUS clients.

You must configure your NASs as RADIUS clients by specifying the name or IP address of the NPS as the authenticating server. Likewise, you must configure the NPS server with the IP address of the RADIUS client.

NPS authentication methods

Authentication is the process of verifying the identity of a user or computer that is attempting to connect to a network. NPS must receive proof of identity from the user or computer in the form of credentials. NPS authenticates and authorizes a connection request before allowing or denying access when users attempt to connect to your network through NASs.

When you deploy NPS, you can specify the required type of authentication method for access to your network.

Some authentication methods implement the use of password-based credentials. The NAS passes these credentials to the NPS server, which verifies the credentials against the user accounts database.

Other authentication methods implement the use of certificate-based credentials for the user, the client computer, the NPS server, or a combination of the three. Certificate-based authentication methods provide strong security and we recommend them over password-based authentication methods.

Any authentication method has advantages and disadvantages in terms of security, usability, and breadth of support. However, password-based authentication methods don't provide strong security because malicious individuals can guess passwords. For that reason, we do not recommend them. Instead, consider using a certificate-based authentication

³ <https://aka.ms/Netsh-Commands-for-Network-Policy-Server>

⁴ <https://aka.ms/powershell-NPS>

method for all network access methods that support certificate use. This is especially true for wireless connections.

For these types of connections, consider using PEAP-MS-CHAP v2 or PEAP-TLS.

The configuration of the NAS determines the authentication method you require for the client computer and network policy on the NPS server. Consult your access server documentation to determine which authentication protocols are supported.

You can configure NPS to accept multiple authentication protocols. You can also configure your NASs, also called RADIUS clients, to attempt to negotiate a connection with client computers by requesting the use of the most secure protocol first, then the next most secure, and so on, to the least secure. For example, the Routing and Remote Access service tries to negotiate a connection by using the protocols in the following order:

1. EAP
2. MS-CHAP v2
3. MS-CHAP
4. Challenge Handshake Authentication Protocol (CHAP)
5. Shiva Password Authentication Protocol (SPAP)
6. Password Authentication Protocol (PAP)

When you choose EAP as the authentication method, the negotiation of the EAP type occurs between the access client and the NPS server.

Warning: You should not use PAP, SPAP, CHAP, or MS-CHAP in a production environment as they are considered highly insecure.

NPS network policies

Network Policies in NPS determine whether to allow or deny a connection request from a RADIUS client. This process also uses the dial-in properties of a user.

When NPS performs authorization of a connection request, it compares the request with each network policy, in the order that they are listed in the NPS snap-in. It starts with the first policy and then moves to the next item in the list.

Because of this, you should make sure that your most restrictive policies appear first in 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 you configure the dial-in properties of the user account to deny access, NPS will deny access to the user, even though the policy would grant them access.

When you create your network policies, you will specify the conditions that you must satisfy to apply the policy. Then, you can specify the settings to apply to the connection request when you apply the policy.

NPS accounting

You also need to consider how you should configure logging for NPS.

You can log user authentication requests and accounting requests to log files in text format or database format, or you can log to a stored procedure in a **Microsoft SQL Server** database.

Use request logging primarily for connection analysis and billing purposes, and as a security investigation tool, because it enables you to identify a hacker's activity.

To make the most effective use of NPS logging:

- Turn on logging initially for authentication and accounting records. Modify these selections after you determine what is appropriate for your environment.
- Ensure that you configure event logging with sufficient capacity to maintain your logs.
- Back up all log files on a regular basis because you can't recreate them when they are damaged or deleted.
- Use the RADIUS **Class** attribute to track usage and simplify identification of which department or user to charge for usage. Although the **Class** attribute, which generates automatically, is unique for each request, duplicate records might exist in cases where the reply to the access server is lost and the request resends. You might need to delete duplicate requests from your logs to track usage accurately.
- To provide failover and redundancy with **Microsoft SQL Server** logging, you could use an SQL server failover cluster. If you need failover and redundancy with log file-based logging, you could place them on a file server cluster.
- If RADIUS accounting fails because of a full hard-disk drive or other causes, NPS stops processing connection requests. This prevents users from accessing network resources.
- If you don't supply a full path statement in **Log File Directory**, the default path applies. For example, if you enter **NPSLogFile** in **Log File Directory**, you will locate the file at **%systemroot%\System32\NPSLogFile**.

Additional reading: For additional information on how to interpret logged data, refer to **Interpret NPS Database Format Log Files**⁵.

Configuring NPS event logging

To configure NPS event logging by using the Windows interface, perform the following tasks:

1. Open the Network Policy Server snap-in.
2. Right-click **NPS (Local)** or access the context menu, and then select **Properties**.
3. On the **General** tab, select each of the following options, as required, and then select **OK**:
 - **Rejected authentication requests**
 - **Successful authentication requests**

Note: To complete this procedure, you must be a member of the **Domain Admins** group or the **Enterprise Admins** group.

Using the event logs in **Event Viewer**, you can monitor NPS errors and other events that you configure NPS to record. NPS records connection request failure events in the System and Security event logs by default. Connection request failure events consist of requests that NPS rejects or discards. The **Event Viewer** records other NPS authentication events. Note that the **Event Viewer** security log might record events containing sensitive data.

⁵ <https://aka.ms/InterpretNPSDatabase>

Connection request failure events

Although NPS records connection request failure events by default, you can change the configuration according to your logging needs. NPS rejects or ignores connection requests for a variety of reasons, including the following:

- The RADIUS message is not formatted according to RFC 2865 "Remote Authentication Dial-in User Service (RADIUS)," and RFC 2866, "RADIUS Accounting."
- The RADIUS client is unknown.
- The RADIUS client has multiple IP addresses and has sent the request on an address other than the one that you define in NPS.
- The message authenticator, also known as a digital signature, that the client sent is invalid because the shared secret is invalid.
- NPS was unable to locate the username's domain.
- NPS was unable to connect to the username's domain.
- NPS was unable to access the user account in the domain.

When NPS rejects a connection request, the information in the event text includes the username, access server identifiers, the authentication type, the name of the matching network policy, and the reason for the rejection.

Connection request success events

Although NPS records connection request success events by default, you can change the configuration according to your logging needs. When NPS accepts a connection request, the information in the event text includes the username, access server identifiers, the authentication type, and the name of the first matching network policy.

Logging Schannel events

Secure channel (Schannel) is a security support provider (SSP) that supports a set of internet security protocols, such as SSL and TLS. These protocols provide identity authentication and secure, private communication through encryption.

Logging of client-certificate validation failures is a secure channel event and is not enabled on the NPS server by default. You can enable additional secure channel events by changing the following registry key value from 1 (**REG_DWORD type, data 0x00000001**) to 3:

```
(REG_DWORD type, data 0x00000003): HKEY_LOCAL_MACHINE\SYSTEM\CurrentControl-Set\Control\SecurityProviders\SCHANNEL\EventLogging
```

Overview of connection request processing

You can use *connection request processing* to decide where you want the authentication of connection requests to occur, either on the local computer or on a remote RADIUS (Remote Authentication Dial-In User Service) Server.

If you want the instance of NPS (Network Policy Server) running on the local computer to take care of the authentication for connection request, you do not need to do anything. The default connection request policy uses NPS as a RADIUS server and processes all authentication requests locally.

Alternatively, if you want to forward connection requests to other NPS or RADIUS servers, you must configure a remote RADIUS server group, and then add a new connection request policy that specifies conditions and settings that match the connection requests.

One of the advantages of using NPS is that you can administer all your connection policies centrally. Consider the following scenario:

- You have three VPN (virtual private network) servers running Windows Server 2019. These provide remote access to the employees in your company. **Active Directory** groups for the various departments in the company control remote access. If you are using the local instance of NPS on all three VPN servers and need to add a new group to provide them with access, you have to add this group to the policy on all three VPN servers. However, if you use NPS (RADIUS) and forward connection requests to a central NPS or RADIUS server, you only need to change it in one policy on one server.
- You can use the **New Connection Request Policy Wizard** to create a new remote RADIUS server group when you create a new connection request.
- If you want the NPS server to act as both a RADIUS server, by processing connection requests locally, and as a RADIUS proxy, by forwarding some connection requests to a remote RADIUS server group, then you should add a new policy, and then verify that the default connection request policy is the last policy processed.

If you forward connection requests to other NPS or RADIUS servers, RADIUS messages provides authentication, authorization, and accounting according to the following workflow:

1. The RADIUS client, which could be a VPN server, wireless access point, switch, or Remote Desktop (RD) Gateway server, receives a connection request from a client.
2. The RADIUS client creates an Access-Request message and sends it to the NPS.
3. The NPS checks the Access-Request message.
4. The NPS contacts a domain controller and verifies the user's credentials along with the dial-in properties of the user. The dial-in properties are only checked if the Access-Request message is from a VPN server.
5. The NPS authorizes the connection by observing the dial-in properties and the network policy configuration.
6. The NPS sends an Access-Accept message to the RADIUS client if it authenticates and authorizes the connection attempt. If the NPS doesn't authenticate or authorize the connection attempt, the NPS sends an Access-Reject message to the RADIUS client.
7. The RADIUS client allows the client to connect to the server, sends an Accounting-Request message to NPS, and logs details about the connection.
8. The NPS sends an Accounting-Response message to the server.

Ports for RADIUS and logging

By default, NPS observes RADIUS traffic on ports **1812**, **1813**, **1645**, and **1646** for both IPv6 and IPv4 for all installed network adapters.

Note: If you disable either IPv4 or IPv6 on a network adapter, NPS does not monitor RADIUS traffic for the uninstalled protocol.

The values of **1812** for authentication and **1813** for accounting are RADIUS standard ports defined in Request for Comments (RFC) 2865 "Remote Authentication Dial-in User Service (RADIUS)," and RFC 2866, "RADIUS Accounting." However, by default, many existing access servers, and often legacy access servers,

use port **1645** for authentication requests and port **1646** for accounting requests. When you are considering what port numbers to use, make sure that you configure NPS and the access server to use the same port numbers. If you don't use the RADIUS default port numbers, you must configure exceptions on the firewall for the local computer to enable RADIUS traffic on the new ports.

Configuring NPS UDP port information

You can use the following procedure to configure the User Datagram Protocol (UDP) ports that NPS uses for RADIUS authentication and accounting traffic.

Important: To complete this procedure, you must be a member of the **Domain Admins** group, the **Enterprise Admins** group, or the **Administrators** group on the local computer.

To configure the NPS UDP port information by using the Windows interface, follow these steps:

1. Open the **NPS** console.
2. Right-click **Network Policy Server** or access the context menu, and then select **Properties**.
3. Select the **Ports** tab, and then examine the settings for ports.
4. If your RADIUS authentication and RADIUS accounting UDP ports vary from the provided default values **1812** and **1645** for authentication and **1813** and **1646** for accounting, enter your port settings in **Authentication and Accounting**.

Tip: To use multiple port settings for authentication or accounting requests, separate the port numbers with commas.

Configure policies on NPS

NPS (Network Policy Server) supports connection request policies and network policies. These policies manage and control connection request attempts for remote access clients and determine which NPS servers manage and control connection attempts.

Connection request policies

Connection request policies allow you to choose whether the local NPS server processes connection requests or forwards them to another RADIUS server for processing.

- With connection request policies, you can use NPS as a RADIUS (Remote Authentication Dial-In User Service) server or as a RADIUS proxy, based on the following:
 - The time of day and day of the week
 - The realm name in the connection request
 - The connection type that you request
 - The RADIUS client's IP address
- When you install NPS, it creates a default connection request policy with the following conditions:
 - Authentication isn't configured
 - Accounting isn't configured to forward accounting information to a remote RADIUS server group
 - Attribute manipulation isn't 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 aren't configured
- The default connection request policy uses NPS as a RADIUS server

Network policies

Network policies allow you to designate which users you authorize to connect to your network and the circumstances under which they can or can't connect. A network policy is a set of conditions, constraints, and settings that enable you to designate who you will authorize to connect to the network, and the circumstances under which they can or can't 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 connection requests require a specific network connection method or type of network access server. Overview properties also enable you to specify whether to ignore the dial-in properties of user accounts in AD DS (Active Directory Domain Services). 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 doesn't match a constraint, NPS rejects the request automatically. Unlike the NPS response to unmatched conditions in the network, if a constraint doesn't match, NPS doesn't evaluate additional network policies, and denies the connection request.
- **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 match and the request is accepted.

When NPS authorizes a connection request, it compares the request with each network policy in the ordered list of policies, starting with the first policy and moving to the next item on 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 authorize the request. 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 doesn't 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.

Important 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.

Implement RADIUS with NPS

RADIUS (Remote Authentication Dial-In User Service) is an industry-standard authentication protocol that many vendors use to support the exchange of authentication information between elements of a remote-access solution. To centralize your organization's remote-authentication needs, you can configure NPS (Network Policy Server) as a RADIUS server or a RADIUS proxy. While configuring RADIUS clients and servers, you must consider several factors, such as the RADIUS servers that will authenticate connection requests from RADIUS clients and the ports that RADIUS traffic will use.

What is a RADIUS client?

RADIUS clients are usually NASs (Network Access Servers) such as wireless access points, 802.1X authenticating switches, and VPN (virtual private network) servers. A NAS is a device that provides access to a larger network. You can configure a NAS to function as a RADIUS client. RADIUS clients communicate with a RADIUS server for authentication, authorization, and accounting. By default, RADIUS devices communicate with each other over ports **1812** and **1813** or **1645** and **1646**. End-user computing devices such as wireless laptop computers, tablets, and other computing devices are not RADIUS clients. These types of devices are clients of the NAS devices. In addition to deploying NPS as a RADIUS server, a RADIUS proxy, or a NAP policy server, you must also configure RADIUS clients in NPS.

RADIUS client examples

Examples of NASs include the following:

- NASs that provide remote access connectivity to an organization's network or the internet. For example, a computer that is running the Windows Server operating system and the Remote Access Service (RAS) that provides either traditional dial-up or VPN remote access services to an organization's intranet.
- Wireless access points that provide physical-layer access to an organization's network by using wireless-based transmission and reception technologies.
- Switches that provide physical-layer access to an organization's network using traditional local area network (LAN) technologies, such as Ethernet.
- NPS-based RADIUS proxies that forward connection requests to RADIUS servers that are members of a remote RADIUS server group that you configure on the RADIUS proxy, or other RADIUS proxies.

What is a RADIUS proxy?

A RADIUS proxy routes RADIUS messages between RADIUS clients and RADIUS servers that perform user authentication, authorization, and accounting for the connection attempt. As a RADIUS proxy, NPS is a central switching or routing point through which RADIUS access and accounting messages flow. NPS records information in an accounting log about forwarded messages.

You can use NPS as a RADIUS proxy when:

- You are a service provider who offers outsourced dial, VPN, or wireless network-access services to multiple customers.

In this case, your NAS sends connection requests to the NPS RADIUS proxy. Based on the username's realm portion in the connection request, the NPS RADIUS proxy that your company maintains on your premises forwards the connection request to a RADIUS server. The customer maintains the RADIUS server and can authenticate and authorize the connection attempt.

- You want to provide authentication and authorization for user accounts that aren't:

- Members of the domain in which the NPS server is a member.
- Members of a domain that has a two-way trust with the NPS server's member domain.

This includes accounts in untrusted domains, one-way trusted domains, and other forests. Instead of configuring your access servers to send their connection requests to an NPS RADIUS server, you can configure them to send their connection requests to an NPS RADIUS proxy. The NPS RADIUS proxy uses the realm-name portion of the username, and then forwards the request to an NPS server in the correct domain or forest. NPS can authenticate connection attempts for user accounts in one domain or forest for a NAS in another domain or forest.

- You want to perform authentication and authorization by using a database that is not a Windows account database.

In this case, NPS forwards connection requests that match a specified realm name to a RADIUS server that has access to a different database of user accounts and authorization data. An example of another user database is a **Microsoft SQL Server** database.

- You want to process a large number of connection requests.

In this case, instead of configuring your RADIUS clients to balance their connection and accounting requests across multiple RADIUS servers, you can configure them to send their connection and accounting requests to an NPS RADIUS proxy.

The NPS RADIUS proxy dynamically balances the load of connection and accounting requests across multiple RADIUS servers, and it increases processing of large numbers of RADIUS clients and authentications each second.

- You want to provide RADIUS authentication and authorization for outsourced service providers and minimize intranet firewall configuration.

An intranet firewall is between your intranet and your perimeter network. A perimeter network is the network between your intranet and the Internet. By placing an NPS server on your perimeter network, the firewall between your perimeter network and intranet must allow traffic to flow between the NPS server and multiple domain controllers.

When replacing the NPS server with an NPS proxy, the firewall must allow only RADIUS traffic to flow between the NPS proxy and one or multiple NPS servers within your intranet.

Demonstration: Manage NPS

In this demonstration you will learn how to:

- Configure the Remote Access policies
- Create a VPN (virtual private network) profile on a Windows client
- Connect to the VPN server using a Windows client

Preparation steps

The required virtual machines (**WS-011T00A-SEA-DC1**, **WS-011T00A-SEA-ADM1**, and **WS-011T00A-SEA-CL1**) should be running after the previous demonstration.

After completing the demonstration, leave all the virtual machines running. You will use them in a later demonstration.

Demonstration steps

Configure the Remote Access policies

1. On **SEA-ADM1**, from **Server Manager**, open the **Network Policy Server** console.
2. In the **Network Policy Server** console, on the **navigation** pane, expand **Policies**, and then select **Network Policies**.
3. Create a new network policy by using the **New Network Policy Wizard** with the following settings:
 - Policy name: **Contoso IT VPN**
 - Type of network access server: **Remote Access Server(VPN-Dial up)**
 - Windows Groups: **IT**
 - Specify Access Permission: **Access granted**
 - Configure Authentication Methods:
 - Add Microsoft Secured password (**EAP-MSCHAP v2**)
 - Add Microsoft: Smart Card or other certificate
 - Clear the **Microsoft Encrypted Authentication (MS-CHAP)** check box
4. Complete the **New Network Policy Wizard** by accepting the default settings on the other pages.
5. Close all open windows.

Create a VPN profile on a Windows client

1. On **SEA-CL1**, right-click **Start** or access the context menu, and then select **Network Connections**.
2. In **Network & Internet**, select **VPN**, and then select **Add a VPN connection**.
3. In the **Add a VPN connection** wizard, use the following values and then select **Save**:
 - VPN provider: **Windows (built-in)**
 - Connection Name: **Contoso VPN**
 - Server name or address: `vpn.contoso.com`
 - VPN type: **Secure Socket Tunneling Protocol (SSTP)**
 - Type of sign-in info: **User name and password**
 - Remember my sign-in info: **Cleared**

Connect to the VPN using a Windows client

1. In **Network & Internet**, select **Contoso VPN**, and then select **Connect**.
2. In the **Sign in** dialog box, in the **User name** text box, enter `contoso\jane`, and in the **Password** text box, enter `Pa55w.rd`, and then select **OK**.
3. Verify your connection to the VPN server.

Verify connection on client and VPN server

1. On **SEA-CL1**, right-click **Start** or access the context menu, and then select **Windows PowerShell (Admin)**.

2. Enter the following command, and then select Enter:

```
Get-NetIPConfiguration
```

3. Examine the output and verify that **Contoso VPN** is listed next to **InterfaceAlias**. Also verify that the Contoso VPN has been issued an IP Address. This is the IP address for VPN connection assigned by RRAS.
4. Switch to **SEA-ADM1** and maximize the Routing and Remote Access snap-in.
5. In the Routing and Remote Access snap-in, select **Remote Access Clients (0)** and verify that **Contoso\jane** is listed under the **User Name** column. This indicates that the user is connected to the VPN Server.
6. Maximize **Server Manager** and the **Tools** menu, and then select **Remote Access Management**.
7. In the **Remote Access Management Console**, select **Remote Client Status** and verify that **CONTO-SO\jane** is listed under **Connected Clients**. Notice that the VPN protocol used displays under the **Protocol/Tunnel** field as **Sstp**.

Test your knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

Which of the following is a RADIUS (Remote Authentication Dial-In User Service) client?

- VPN (virtual private network) Server
- Wireless access point
- Windows 10 client
- Windows Server 2019 member server
- Remote Desktop (RD) Gateway server

Question 2

Which authentication protocols should you use in a production environment? Choose two.

- SPAP (Shiva Password Authentication Protocol)
- EAP
- PAP (Password Authentication Protocol)
- MS-CHAP
- CHAP
- MS-CHAP v2

Question 3

What kind of policies can you create on a Network Policy Server? Choose two.

- Connection Request Policies
- Group Policies
- Network Policies
- Configuration Policies

Implementing Always On VPN

Lesson overview

Always On VPN is the next generation VPN (virtual private network) solution for Windows 10 devices. It can provide very secure access to internal data and applications. To properly implement and support Always On VPN in your environment, you must understand how to select a tunnel mode, choose the VPN authentication protocol, and configure the server roles to support your chosen configuration.

One of the advantages of using Always On VPN compared to traditional VPN technologies is that it is fully automated. A VPN connection will automatically trigger based on network conditions. Furthermore, Always on VPN supports all Windows 10 editions as clients and might not require client domain membership depending on the tunnel mode.

Lesson objectives

After completing this lesson, you will be able to:

- Describe Always On VPN.
- Understand the prerequisites for deploying Always On VPN.
- Describe Always On VPN features and functionality.
- Explain why you would choose Always On VPN over Windows VPN.
- Understand how to deploy Always On VPN.

What is Always On VPN?

Always On VPN enables remote users running Windows 10 to securely access corporate resources such as email servers, shared folders, or internal websites, without manually connecting to a VPN (virtual private network). When a client is outside the company network, the Windows 10 VPN client automatically detects an untrusted network and connects securely to the VPN without any user intervention. When the client moves within the company's network, the Windows 10 VPN client detects a change and automatically disconnects from the VPN server.

Always On VPN also provides increased productivity for a mobile workforce by offering the same connectivity experience both inside and outside the office. You can consider Always On VPN to be the successor to DirectAccess, which works in a similar way with different technologies.

Additional reading: For detailed information about Always On VPN, refer to **Remote Access Always On VPN**⁶.

Prerequisites for Always On VPN deployment

To deploy and configure Always On VPN, your organization must support the following infrastructure components:

- Always On VPN Gateway (VPN server)
 - Microsoft remote access server, a virtual private network (VPN), is an excellent choice for the Always On VPN Gateway. This role supports the features needed for Always On VPN, namely LAN routing and IKEv2 VPN protocol. IKEv2 (Internet Key Exchange version 2) uses the IPsec Tunnel

⁶ <https://aka.ms/remote-access-always-on-vpn>

Mode protocol over UDP (User Datagram Protocol) 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. This means that Always On VPN automatically restores the connection without any user actions if you lose connectivity.

- Always On VPN clients
 - The client machine should run Windows 10 to support both user and tunnel mode. We recommend that you always use the latest version of the Windows 10 operating systems especially when working with Always On VPN. Each new release of Windows 10 adds new functionality and performance enhancements Always On VPN. If you are using tunnel mode, you must join the client machine to an AD DS (Active Directory Domain Services) domain.
- Network Policy Server (NPS)
 - NPS enables you to create and enforce organization-wide network access policies for connection request authentication and connection request authorization. NPS is Microsoft's implementation of a RADIUS server, but you can also use third-party RADIUS servers. Always ON VPN Gateway server configures as a RADIUS when using NPS.
- An Active Directory domain
 - You must have at least one Active Directory domain to host groups for your Always On VPN users, Always On VPN server, and NPS servers. Furthermore, PEAP (Protected Extensible Authentication Protocol) requires special user attributes stored in AD DS, which authenticates users and provides authorization for VPN connection requests.
- Group Policy
 - You should use Group Policy for the autoenrollment of certificates used by Always On VPN users and clients.
- Firewall configuration
 - You must configure all firewalls to allow the flow of VPN and RADIUS traffic, otherwise Always On VPN will not function correctly.
- Public key infrastructure (PKI)
 - Always On VPN requires certificates. You must implement certificates for authentication for every user that will participate in Always ON VPN communication. If you are using Tunnel mode, then every Windows 10 client requires a Computer Authentication certificate. The VPN server itself requires a Server Authentication Certificate and the NPS Server. You can use a public SSL certificate if you use the SSTP (Secure Socket Tunneling Protocol) protocol for Always On VPN. However, if you use the IKEv2 authentication protocol then you must use a certificate for your internal PKI.
- Domain Name System (DNS) server
 - The external and internal environments require DNS zones. The internal DNS zone could be a delegated subdomain of the external DNS name (such `corp.contoso.com`). If you use the same name externally and internally, support for split-brain DNS is also available.

Always On VPN features and functionalities

Always On VPN offers many features and enhancements when compared to traditional virtual private network (VPN) solutions. One of the most interesting features might be integration support for cloud services such as support for Azure Multi-Factor Authentication (MFA), Azure conditional access, Windows Hello for Business, and Windows Information Protection (WIP).

The following summary gives you an overview of the most important functionality that Always On VPN has to offer:

- **VPN Connection.** Always On VPN will automatically connect to the VPN server when a client moves from a trusted network to an untrusted network. This works with either a device tunnel or a user tunnel. In the past, it was not possible to trigger an automatic VPN when either the user or the device was authenticated.
 - The Always On VPN connection options include:
 - **Application triggering.** You can configure Win32 and UWP (Universal Windows Platform) applications to trigger the VPN connection when they start.
 - **Name-based triggering.** You can create rules that allow you to trigger an Always On VPN connection based on domain names. It does matter whether or not you join the Windows 10 client to a domain.
 - **Device tunnel.** Always On VPN supports two tunnel modes, user tunnel and device tunnel. User tunnel mode starts the VPN after the user logs on, where device tunnel mode starts the VPN connection before the user signs in. You can use device tunnel and user tunnel at the same time or separately. It's important to remember that device tunnel mode requires domain membership and only IKEv2 (Internet Key Exchange version 2) authentication protocol supports it.
- **Security.** Always On VPN supports new security options that enables it to restrict the VPN connection. You can control the applications, authentication methods, and type of traffic that can use the Always On VPN connection. An Always On VPN connection will be active for a longer duration than a traditional VPN connection, which the user starts when needed. It's important to secure your Always On VPN connection properly and securely. Always On VPN supports the use of the latest RSA (Rivest–Shamir–Adleman) cryptographic algorithms and natively supports Enrollment for Application Platform (EAP). The support of EAP enables you to use a variety of Microsoft or third-party authentication types (EAP type). This enables you to configure secure authentication based on username and password, user certificates, smart cards, Windows Hello for Business, and MFA. You can also use UWP plug-in authentication which enables third-party support for various custom options such as token or OTP (One Time Password) support.
- **Platform integration.** Always On VPN has support for many scenarios because of tight integration with the Windows operating system and support for third-party solutions.
- **Compatibility and configuration.** Always On VPN supports several ways of management and deployment, which is not always the case with other VPN client solutions.
 - **Supported platforms.** Always On VPN supports all Windows editions, Azure AD joined devices, workgroup devices, and domain joined devices.
 - **Deployment and management.** Always ON VPN supports several methods for creating and managing a VPN profile for Always On VPN. These methods include Microsoft Endpoint Configu-

ration Manager, Intune, Windows Configuration Designer, and PowerShell or any third-party mobile device management tool.

- **VPN Gateway compatibility.** Even though Microsoft remote access server (VPN) is an excellent choice for the Always On VPN Gateway, you can use third-party VPN gateways because of the support for the industry standard IKEv2 protocols. You can also use the UWP VPN plug-in to get support for solutions that are not related to Microsoft VPN server.
- **Networking.** Always On VPN works equally well with IPv4 and IPv6, but it's not dependent on IPv6 like DirectAccess. You can also create granular routing policies, which enables you to lock down or control access to individual applications. It's also possible to exclude certain applications so you can control traffic flow and routing. To use exclusion routes, you must run a split tunnel setup.

Additional reading: For detailed information about Always On VPN features and functionality, refer to **Always On VPN features and functionalities**⁷.

Why choose Always On VPN over Windows VPN?

VPN (virtual private network) connections enable users who are working offsite (for example, from home, a customer site, or a public wireless access point) to access a server on an organization's private network 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 a computer, the VPN client, and an organization's server. The exact infrastructure of the shared or public network is irrelevant because it displays as if you sent the data over a dedicated private link. You can access the same resource through a VPN as you can within the company office.

DirectAccess was designed to provide automatic always-on remote access to company resources from outside the company boundaries.

DirectAccess requires that any client must be a domain-joined computer that is running an Enterprise edition of the Windows 10 or Windows 8.1 operating systems.

DirectAccess has not received any new features since Windows Server 2012 and Microsoft recommends that companies currently considering DirectAccess should deploy Always On VPN instead. Even though DirectAccess hasn't been "officially" deprecated and Windows Server 2019 still supports it, there is no guarantee that Microsoft will continue to support it after Windows Server 2019. Always On VPN can replace DirectAccess, as it provides all the same features and improves on them.

The following table compares Always On VPN, traditional VPN, and DirectAccess features.

Feature	Always On VPN	Traditional VPN	DirectAccess
Domain join required	Not required for user tunnel, but required for device tunnel	Not required	Required for all clients
Client built-in	Built-in and third-party	Built-in and third-party	Only built-in
Manual or autoconnect	Always automatic	Manual	Always automatic
Firewall requirements	Depends on the protocols used, but normal VPN ports are blocked from some locations	Depends on the protocols used, but traditional VPN ports are blocked from some locations	Only port 443 with IP-HTTPS

⁷ <https://aka.ms/remote-access-vpn-map-da>

Feature	Always On VPN	Traditional VPN	DirectAccess
Manual disconnect	Yes	Yes	No
Operating system (OS) support	Only Windows 10	Support for every OS and device	Only Windows 10 Enterprise, Windows 8.1 Enterprise

The gap from traditional VPN to Always On VPN might not be as long as you think. Both traditional VPN and Always On VPN requires a backend infrastructure which includes the Remote Access Server and a RADIUS (Remote Authentication Dial-In User Service) server. Both VPN solutions can use the Remote Access Server in the Windows Server operating system and they can also use third-party VPN server solutions. Neither traditional VPN nor Always On VPN depends on the Remote Access Server in Windows Server. In theory, Always On VPN can use any VPN Server implementation because the Windows 10 client configures the features of an Always on VPN solution in the form of a special VPN profile.

You can use the VPN server for Always On VPN clients at the same time and with the same configuration as traditional VPN clients. Both VPN solutions can work without domain membership but if you configure Always On VPN to use a device tunnel, you will require domain membership. Traditional VPN also works with every OS and almost any device, but it might require a third-party VPN client. The Always On VPN client is built into the Windows 10 operating system and requires no installation or maintenance.

Even though Always On VPN provides the same automatic connect feature as DirectAccess, it doesn't offer the same functionality. If you are using the user tunnel with Always On VPN, the VPN connects when you sign in to the machine. On the other hand, device tunnel starts the VPN connection before you sign in. It's designed for manage-out scenarios and for inexperienced users that haven't previously logged to their machines. You should also remember that device tunnel requires IKEv2 (Internet Key Exchange version 2) and domain joined machines.

With traditional VPN, you can choose to disconnect, which means that you might not be able to manage the machine for Group Policy, Software updates, and so on. With Always On VPN, you can disconnect using the GUI if you are using user tunnel. But if you are using the device tunnel, there is no VPN profile in the GUI for you to disconnect.

When it comes to deciding whether to switch from using a traditional VPN solution to Always On VPN, the choice might not be an easy one. If you are using a third-party VPN solution and want something that doesn't require a lot of management, VPN client built-in in the operation system and part of the Windows license, then Always On VPN might be good choice. If you are already using traditional Windows VPN, you might already have all or most of the configuration on the VPN server needed for Always On VPN, and you can easily implement Always On VPN.

You can also choose to support both traditional VPN and Always On VPN, without choosing between either one of them. As previously described, you can use the VPN server for Always On VPN clients and traditional VPN clients at the same time and with the same configurations. This means that you can get the benefits of both. You can evaluate Always On VPN on your Windows 10 clients and continue to support traditional VPN connections from other devices.

Note: Before you begin a migration to Always On VPN from either a traditional VPN solution or DirectAccess, you should document your requirements for a remote access solution. Then you will know if Always On VPN is the correct solution for your environment.

Deploy Always On VPN

To properly implement and support Always On VPN in your environment, you must understand how to plan, configure, and scope your Always On VPN implementation. These steps will include actions such as

selecting a tunnel mode, choosing the VPN (virtual private network) authentication protocol, and configuring the server roles to support your chosen configuration.

In this topic, you will learn about the various steps to implement Always On VPN. The deployment of Always On VPN usually includes the following steps:

1. Always On VPN deployment planning
2. Always On VPN server infrastructure configuration
3. Remote Access Server configuration for Always On VPN
4. NPS (Network Policy Server) installation and configuration
5. Firewall and DNS configuration
6. Windows 10 Client configuration for Always On VPN

Always On VPN deployment planning

You will start your journey toward Always On VPN by planning your deployment.

- **VPN Authentication protocols.** You need to decide if you are going to use the built-in Windows VPN client or you are going to use a UWP VPN plug-in or a third party.
- **Routing.** You must choose whether your routing setup should be force tunneling or split tunneling.
- **Device or user tunnel.** Always On VPN supports two tunnel modes, user tunnel or device tunnel. It's important to remember that device tunnel mode can only use IKEv2 (Internet Key Exchange version 2), requires domain membership, and is only supported with Windows 10 enterprise clients.
- **Plan the Remote Access Server.** You can use the Remote Access Server is Windows Server operating system or you can use most third-party VPN implementations.
- **Plan for VPN tunneling protocol.** You must decide which VPN protocol you are going to use. You should use either IKEv2 or SSTP (Secure Socket Tunneling Protocol) because of the support for mobility and because they are secure. Even though you could use other protocols, you should avoid using other protocols that are "old" and insecure such as L2TP (Layer 2 Tunneling Protocol) and PPTP (Point-to-Point Tunneling Protocol).
- **Configure firewalls and routing.** You must open the ports on your firewall to allow VPN traffic. The ports you open depends on the VPN protocols that you use.

Always On VPN server in infrastructure configuration

- You must install and configure all the components required to support Always On VPN. These usually include enabling autoenrollment of certificates for both users and for computers if needed. You must create the certificate templates for the VPN server, NPS Server, and for users.

Remote Access Server configuration for Always On VPN

- To support Always On VPN, you install and configure a VPN server. Configure the allocation of IP addresses to remote clients, assign the VPN server certificate, and allow the VPN protocol you have decided to use. You should always disable the unused protocols on the VPN server.

NPS Server installation and configuration

You must install and configure the Network Policy Server to support Always ON VPN connections. This typically includes the following actions: Enroll the NPS certificate, configure Network Policies and add the VPN server as a RADIUS (Remote Authentication Dial-In User Service) client.

Firewall and DNS configuration

You must ensure that you update the DNS server with the name of the Always On VPN used in the certificate. This could differ from the actual host name of the machine or device. Your external firewalls must be configured to allow VPN traffic from the internet to reach the VPN server. The NPS server and VPN server must be able to communicate and you might need to allow NPS traffic in the Windows Defender Firewall.

Windows 10 Client configuration for Always On VPN

The most important aspect of the Always On VPN deployment, is the configuration of the Windows 10 VPN profile. Up until now, you have actually “only” configured a normal VPN implementation based on i.e. IKEv2 or SSTP. Windows 10 clients could use this but wouldn't be Always ON VPN. For Windows 10 to function as an Always On VPN client, you must configure a special VPN profile. You can deploy the Always On VPN profile using either Microsoft Endpoint Configuration Manager, PowerShell, or Intune.

Tip: The configuration of the Windows 10 Always On VPN profile uses the VPNv2 configuration service provider (CSP). A CSP is an interface that can manage settings for a given feature in Windows 10 and have some similarities with Group Policy client-side extensions. Mobile device management systems (MDMs) use CSPs to configure mobile devices.

Additional reading: For additional information on VPNv2, refer to **VPNv2 CSP**⁸.

ProfileXML

ProfileXML is the name of an URI (Uniform Resource Identifier) node within the VPNv2 CSP. The **ProfileXML** node is used to configure Windows 10 VPN client settings and the XML data contains all information needed to configure a Windows 10 Always On VPN profile.

The easiest way to create the ProfileXML file is to create a template VPN profile on a Windows 10 machine. Perform the following steps:

1. Make a note of the name of the NPS server by examining the NPS Certificate.
2. Create a VPN profile template in Windows 10 using the **Add a VPN connection wizard**. Remember to specify the external FQDN (Fully Qualified Domain Name) of your VPN server and under the Protected EAP properties, specify the name of the NPS server, you recorded in step 1.
3. Try to connect to the VPN server from the outside of your network using the newly created VPN profile template. This will ensure that you configured the VPN server correctly and it's working as expected. Furthermore, if you don't connect at least once, the VPN profile template won't have all the information required when we create the ProfileXML file later in the steps.
4. Use the **MakeProfile.ps1** script to create the **VPN_Profile.ps1** and the **VPN_Profile.xml** files. Before you run the script, you need to change some of the values such as the FQDN of your VPN server, the name of the Always On VPN profile, IP addresses of your internal DNS servers, and the name of your trusted network.

⁸ <https://aka.ms/mdm-vpnv2-csp>

Use the **VPN_Profile.ps1** script to create the Always On VPN Profile on your Windows 10 device. You can deploy this script using Microsoft Endpoint Configuration or run it manually on a Windows 10 device.

Use the **VPN_Profile.XML** file if you want to deploy your Always On VPN profile using Intune or a third-party MDM tool.

Additional reading: For detailed information about creating the ProfileXML and the MakeProfile.ps1 PowerShell script, refer to **Configure Windows 10 client Always On VPN connections**⁹.

Test your knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

Which of the following infrastructure components does Always On VPN require? Choose all that apply.

- Remote access server (virtual private network)
- Azure Virtual Network Gateway
- Group Policy
- Windows 10 clients
- Public Key Infrastructure
- Network Policy Server

Question 2

Which methods does Always ON VPN support to create and manage a VPN profile? Choose three.

- Group Policy
- PowerShell
- Intune
- Microsoft Endpoint Configuration Manager

Question 3

What is the name of the configuration item used to configure an Always On VPN profile?

- AlwaysOn.conf
- ProfileXML
- OMA-DM
- PEAP

⁹ <https://aka.ms/vpn-deploy-client-vpn-connections>

Implementing Web Server in Windows Server

Lesson overview

Microsoft Internet Information Services (IIS) version 10 is the Web Server included in the Windows Server 2019 operating system.

In this lesson, you will learn about the high-level architecture of IIS and about the new functionality included in IIS 10. You will also learn about the prerequisites for installing IIS, and how to perform a basic installation and configuration of IIS.

Lesson objectives

After completing this lesson, you will be able to:

- Describe IIS in Windows Server.
- Describe the new features in IIS.
- Describe the architecture of IIS.
- Describe network infrastructure requirements for a web server.
- Perform a basic installation and configuration of IIS.

IIS in Windows Server

IIS (Internet Information Service) is a Hypertext Transfer Protocol (HTTP) web server. The server accepts incoming HTTP requests, processes them, and sends a response back to the requestor.

HTTP is an application-level protocol that uses the Transmission Control Protocol (TCP) as its lower-level transport protocol. A Windows kernel-mode driver, `http.sys`, scans for incoming TCP requests on whatever ports to which IIS is configured. For example, in a typical internet web server, `http.sys` would scan for TCP **port 80**. The kernel-mode driver performs several basic security checks on incoming HTTP requests before passing the request to a user-mode worker process. The worker process fulfills the request. The response generated by the worker process is then sent back to the requestor. Usually, the requestor is a web browser.

The model of a kernel-mode driver and user-mode worker processes offers several advantages:

- The kernel-mode driver can run fast and can have more efficient access to the computer's network hardware.
- The user-mode worker process can be isolated so that a code problem does not affect other worker processes on the same computer.
- The kernel-mode driver can offer some basic protections, helping to protect user-mode code from several common kinds of attack. These include malformed HTTP requests.

You can enhance HTTP by the adding Secure Sockets Layer (SSL) and Transport Layer Security (TLS). When you add SSL to HTTP, the resulting protocol is HTTP Secure, or HTTPS. By using digital certificates, HTTPS can encrypt the connection between the server and client, protecting the contents of requests and responses from a third party. HTTPS can also provide authentication so that the requestor can make sure that responses are coming from the intended server.

Important: You should protect all your websites using HTTPS. This will protect your website's integrity and protect the communication between users and websites.

Moreover, some of the popular browsers used today to access content on websites may restrict certain features if HTTPS is not enabled on the website.

HTTP requests and responses are text-based. A typical internet web server will create responses whose text complies with the HTML specification, and a web browser will render that response into an observable page. However, HTTP responses are not limited to HTML and can contain any text-based information. Other common HTTP response types include graphics files, XML, JavaScript Object Notation (JSON), and other kinds of data. Not all response types are intended for direct display to users but are instead used to send data between computer software and processes.

Most web servers produce at least some of their content dynamically. That is, they accept information in an HTTP request and then run software code to produce the response. That code can come in the form of a compiled executable (.exe program), a script language such as PHP, a web development framework such as ASP.NET, and other forms. Web servers can also serve static content, such as HTML files that are stored on the server's disk and sent upon request to web browsers.

One aspect of HTTP often confuses inexperienced users. HTTP uses TCP, a connection-oriented protocol. Some inexperienced users believe that the server and client maintain the connection during a user's interaction with the server. That is not true. The connection between the client and the server lasts only for the time that is required for the client to send their request, or for the server to send its reply. After that, the connection is closed. Therefore, each request the client sends to the server is a new connection. If the server wants to maintain some information about the client, then you must write the software code to do this.

Re-identifying a client to the server is the main reason that cookies were created. A cookie is a small piece of information that is sent to the client by the server, which the client then sends to the server as a part of each new request. That information could enable the server to identify the client.

What's new in IIS 10.0?

IIS (Microsoft Internet Information Services) 10 is the version of the Web Server included in the Windows Server 2016 and later. With every new release of the operating system, new functionality has been added to IIS to support new standards and innovations in security.

IIS 10 introduced the following functionality:

- IIS in Nano Server container
- IIS in Containers
- Wildcard Host Headers
- Managing IIS
- HTTP/2
- IIS Thread Pool Ideal CPU Optimization for NUMA hardware

IIS on Nano Server

Nano Server was a new deployment option in Windows Server 2016 that has a smaller system resource footprint, starts up significantly faster, and requires fewer updates and restarts than by using the Server Core deployment option for Windows Server. Nano Server only deploys in a Windows container, and it functions primarily as an application host. PHP, Apache Tomcat, and ASP.NET Core can run on IIS in Nano Server.

IIS in 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. Containers provide an isolated operating environment that you can use to deliver 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.

Windows Server 2016 and later supports two 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, Hyper-V containers do not share the OS kernel with the container host. If you run more than one container in a Hyper-V virtual machine, then those containers share the virtual machine's kernel.

IIS 10.0 runs equally well in both Windows Server containers and in Hyper-V containers and offers support for either Nano Server or Server Core images.

Wildcard Host Headers

You must assign at least one IP address to your web server. Some web servers host multiple websites and might require multiple IP addresses so that you can access each site individually. Instead of using one IP address per website, you can assign a host header to each website. A host header is the DNS name used to access the website, such as `www.contoso.com`.

With Wildcard Host Headers, you can create host headers that will work for any subdomain. If you want all traffic for either the `north.contoso.com` or the `south.contoso.com` subdomains going to the same website, you would create a Wildcard Host Header for `.contoso.com`.

Managing IIS

With IIS support for either Server Core or container images running in Nano Server, IIS now provides a better management experience in the form of either IIS Administration PowerShell cmdlets or Windows Admin Center (WAC).

Note: The Windows Server 2016 release introduced IIS administration through Microsoft IIS Administration. Since then, IIS Web Manager has been deprecated and you should use the IIS extension for WAC.

The new IIS Administration PowerShell module will simplify your administration of IIS either directly from the command-line or through scripting.

Additional reading: For additional information on the IIS Administration PowerShell module, refer to **IISAdministration PowerShell Cmdlets**¹⁰.

¹⁰ <https://aka.ms/iisadministration-powershell-cmdlets>

HTTP/2

The release of Windows Server 2016 first included support for HTTP/S and releases prior to IIS 10.0 didn't include support.

HTTP/2 will reduce the connection load and latency on your web servers.

Additional reading: For additional information about HTTP/2, refer to **HTTP/2 on IIS**¹¹.

IIS thread pool ideal CPU optimization for NUMA hardware

IIS 10.0 now has support for ideal central processing unit (CPU) setting for non-uniform memory access (NUMA) hardware. This enables IIS to distribute IIS thread pool evenly across all CPUs in NUMA nodes. This feature is enabled by default when installing IIS 10.0 and it is recommended to leave it enabled.

Additional reading: For additional information about ideal CPU setting for NUMA hardware, refer to **IIS Thread Pool Ideal CPU Optimization for NUMA hardware**¹².

Overview of IIS architecture

IIS (Microsoft Internet Information Services) 7 and later uses a request-processing architecture which includes:

- The Windows Process Activation Service (WAS)
- A customizable web server engine, meaning you can add or remove modules depending on what functionality you need.
- Integrated request-processing pipelines from IIS and ASP.NET

IIS consists of various components that each perform functions for the web server and application such as reading configuration files and observing requests made by IIS.

These components are either services or protocol listeners and include the following:

- Protocol listeners
 - HTTP.sys
- Services
 - World Wide Web Publishing Service (WWW service)
 - Windows Process Activation Service (WAS)

HTTP.sys observes incoming TCP (Transmission Control Protocol) requests on the ports that configure IIS. The kernel-mode driver performs several basic security checks on incoming HTTP requests before passing the request to IIS for processing. When IIS processes the request, it transmits back to the requestor. Usually, the requestor is a web browser.

The World Wide Web Publishing Service is responsible for reading IIS configuration from IIS, which will update HTTP.sys if there are any changes in the configuration. The WWW service monitors IIS Performance information and manages performance counters.

The Windows Process Activation Service (WAS) is responsible for the management of worker processes and work pool configuration.

¹¹ <https://aka.ms/http2-on-iis>

¹² <https://aka.ms/ideal-cpu-numa-optimization>

Modules in IIS

Many of the IIS components installed by Server Manager or by Windows PowerShell are called modules. In IIS, a module is a binary component that is installed on the server, and that can provide functionality to all websites on the server. Modules can consist of native dynamic link library (DLL) files or .NET assemblies.

Modules provide some IIS functionality that you will learn about in the next topic. For example, modules provide the default document functionality, directory browsing feature, and static webpage capability.

You can install modules by using Windows PowerShell, Windows Admin Center, or Server Manager for modules that are included in the Windows operating system. You can install other modules by using the Web Platform Installer or by running the modules' own installer. After you install the modules, you must enable them to enable their functionality. You can enable them by using IIS Manager, Windows PowerShell, the AppCmd.exe command line tool, or Windows Admin Center with the Microsoft IIS Web Manager extension installed.

Application pools in IIS

When you use IIS Manager to create a new website, IIS Manager will automatically create a new application pool for the website and assign the new website to that new application pool. The new application pool will be given a default name that corresponds to your new website name. This naming convention intends to help you easily identify the relationship between websites and application pools.

Usually, it's desirable to have each website in its own application pool. However, you might want to rename the new application pool to comply with organizational naming standards. You might also want to review the configuration of the new application pool to make sure the configuration is correct for the web application that the website will host.

When you want the new website to be assigned to an existing application pool, you can reassign the website after you create it, and then delete the new application pool that was created by IIS Manager. You might also want to create an application pool manually before you create the new website. This enables you to completely configure the application pool before assigning the website to it.

Application isolation in IIS

Dynamic websites usually consist of two types of software code. The first type is the system code, which consists of IIS itself, and the parts of the Windows operating system that support IIS. For example, the http.sys kernel-mode driver is part of the system code. The other type of code is the user code, which consists of dynamic webpages that are created by a software developer, including ASP.NET pages and PHP pages.

The earliest versions of IIS ran both types of code in a single process. All the code ran in the same area of memory and was assigned a single priority by the operating system. The problem with that approach is that poorly written user code could make the web server run slowly or even crash. If a web server hosted multiple websites, then a single poorly written website could make them all run slowly or even become unavailable.

Application pools were created to help provide memory and central processing unit (CPU) isolation between different websites running on IIS. Each application pool is serviced by one or more worker processes. These worker processes can run on separate processes within the computer, and the operating system assigns each worker process to its own area of memory. Therefore, each worker process isolates the user code that it runs from other user code and from the system code. A poorly performing website might make its worker process run slowly. However, that will not necessarily impact other worker processes, and it will not necessarily impact IIS itself.

The overall purpose of an application pool is to provide configuration settings for those worker processes.

Websites are assigned to an application pool. Each application pool can manage multiple websites, or you can set aside an application pool to manage only a single website. You can configure each application pool to support different behavior for its worker processes, helping to support different web application requirements.

IIS also enables you to create applications in a website. You can assign these applications to a different application pool than the main website so that you can isolate sections of a website from one another.

Worker Processes in IIS

Application pools don't represent actual running processes on the computer. Instead, one or more worker processes service each application pool. These display in the Windows Task Manager as W3wp.exe.

Note: W3WP stands for World Wide Web (W3) Worker Process (WP). The abbreviation "W3" is often used to refer to objects that are associated with websites and web servers. This includes the W3 Consortium, the organization responsible for maintaining standards related to websites and web servers.

The kernel-mode http.sys driver receives incoming HTTP requests. Running this driver in kernel mode enables it to run very quickly, in part because it receives faster access to networking hardware and privileged access to server memory. Then, the request is sent to user-mode IIS code. That code examines the request and decides which website to send the request. The request is sent to one of the worker processes that manages that website, based on the website's application pool assignment.

Worker processes are completely self-contained, like any application process that's running on Windows. Application pools provide many configuration options for worker processes. Two of the main configuration settings relate to the number of worker processes and to worker process recycling.

Multiple worker processes

By default, each application pool has a single worker process. That process must run all user code associated with the websites that are assigned to the application pool. For small websites, a single worker process is usually sufficient. Worker processes are multithreaded so that they can manage more than one incoming request at a time.

However, application pools that manage larger websites, or that manage multiple small websites, might be unable to process all incoming requests quickly enough. In those cases, you can configure the application pool to use multiple worker processes.

Each worker process requires a small amount of additional overhead. Each one requires the OS to set aside a small amount of memory and a small amount of CPU time to run the worker process. There is a tradeoff between performance and having more worker processes. In some scenarios, you might obtain better performance by using a single worker process than by configuring multiple worker processes.

Newer server hardware usually supports non-uniform memory access (NUMA), a technology that can provide faster memory access to the server's CPUs. When it runs on NUMA-compatible servers, you should configure IIS to use the same number of worker processes as there are NUMA nodes in the server.

Worker Process Recycling in IIS

The user code run by a worker process might be poorly written. A common problem with poorly written websites is memory leakage. This means that the user code continually requests more memory from the OS but does not return that memory to the OS when the code finishes running. In that scenario, the worker process gradually consumes more memory until the server runs out.

Worker process recycling provides a workaround to the problem of poorly written user code. When a worker process is recycled, the process is terminated completely. That returns all its memory to the OS and shuts down all user code that was running in the process.

Overview of the Web Server server role

On Windows Server, IIS is installed as a server role named **Web Server (IIS)**. In addition to the basic role, numerous optional role services are available that enable specific IIS functionality. As a best practice, you should only install the role services needed for your particular application. Each role service installs executable code on the server. By installing only role services that you will use, you can help minimize vulnerabilities that a malicious hacker could exploit.

Some documentation refers to the IIS role services as IIS modules. You can find a list of available role services at **Modules in IIS 8.5**¹³.

Commonly used role services include:

- **Static content.** This gives IIS the ability to send static HTML pages as responses.
- **Default document.** This gives IIS the ability to send a default webpage when the HTTP request doesn't specify.
- **ASP.NET.** This gives IIS the ability to create dynamic responses by running ASP.NET code.
- **HTTP Logging.** This gives IIS the ability to log HTTP requests to a text log file or other location.

As you can observe, role services provide IIS with basic capabilities in addition to advanced features. If you were to install IIS and include no role services, it would not be a functional web server.

Install and configure Web Server

IIS (Microsoft Internet Information Services) is installed as a role in Windows Server and includes several optional role services. You can perform the installation of IIS using either PowerShell, Windows Admin Center, or Server Manager.

Hardware and software requirements

IIS 10 has specific hardware and software requirements. It's available in Windows Server 2019 and in Windows 10. Installing IIS on a client operating system is primarily intended for a software development environment.

We will focus on IIS installations on the server operating system, which is typical for a production environment.

You can install the IIS role on any hardware on which you can install Windows Server. But based on server load, you might require more resources (CPU, storage, network).

The primary consideration for IIS capacity planning is the code that the websites will run on the server.

DevOps or IT professionals must perform extensive testing in a lab environment to estimate the capacity of a given hardware configuration for a given website or web application.

It's important to perform a pilot implementation of any new website or application, so that you can make an estimate of how the computer will perform under the intended workload. Because almost all web servers run as virtual machines, the resources (processor and memory) for the operating system and IIS can increase or decrease more easily than in a physical computer.

¹³ <https://aka.ms/Modules-in-IIS-85>

Web servers have several infrastructure requirements. In addition to an operating network, these requirements provide services that make web servers more available and more secure.

Web servers must be connected to an IP-based network and must be assigned at least one IP address. Some web servers host multiple websites and might require multiple IP addresses so that each site can be accessed individually.

A Domain Name Service (DNS) is usually required. DNS enables a user to enter an easy-to-remember server name, such as `www.contoso.com`, instead of having to remember the server's numeric IP address. IP addresses can also change over time, and DNS lets users remember an unchanging name. Finally, DNS provides one method of balancing incoming requests across multiple identical web servers.

Internal and internet web servers usually require DNS. For example, when an internet user tries to access `www.contoso.com`, the user's computer must query a DNS server to find the IP address for the host "www" in the domain `contoso.com`. That process usually requires the cooperation of several DNS servers. For example, the user's local DNS server might have to query the top-level internet DNS server for the ".com" domain, and that server might return a reference to the DNS server handling domain names starting with "c," and so on.

When you work with IIS, it is especially important to understand the difference between DNS A records and canonical name (CNAME) records (or AAAA records, in an IPv6-enabled network). Typically, an A (or AAAA) record is supposed to provide the IP address for a computer's real host name, such as **SEA-SVR2**. A CNAME record provides a nickname, or an alias, which is an easier name to remember. For example, when an intranet user tries to access `http://intranet`, the user's computer will query DNS for the name "sharepoint." That might be a CNAME record that refers to the host named **SEA-ADM1**. The user's computer will then query the A (or AAAA) record for **SEA-ADM1** to find its IP address. The user's web browser address bar will still display `http://intranet`. A single host can be referred to by multiple CNAME aliases. A single alias can also refer to multiple hosts, such as in a web farm.

IIS security

Because web servers are exposed to end users, including anonymous users, they are a common target of attacks. These attacks might seek to take the web server offline, to reveal private information that is contained in the web server, or to deface the websites hosted by the server. Therefore, web servers typically require several kinds of protection.

- Web servers exposed only to an organization's internal users are frequently protected by a single firewall. This might be a software firewall (such as the Windows Defender Firewall) installed on the web server. These firewalls help ensure that users can only connect to designated TCP (Transmission Control Protocol) ports, such as the ports that IIS listens for incoming HTTP requests. Other ports, and the software that might be using them, are protected from any potentially malicious incoming connections.
- Internet web servers should always be protected by a firewall between the web server and the public internet. In most cases, this is a hardware firewall. It can offer better performance and protection than a locally installed software firewall. A hardware firewall can protect multiple web servers consistently. Another firewall might also isolate the web server might also be isolated from its owners' internal network. This secondary firewall creates a perimeter network that contains the web server. The perimeter network helps provide additional protection for the internal network, restricting how far into the network an internet-based malicious hacker can potentially penetrate.

Verify the installation of IIS

After completing the IIS installation, you should verify that IIS is working. By default, IIS creates a single website that observes HTTP requests on TCP port 80, using all the server's IP addresses.

If you are signed in to the server console, then you can quickly test the IIS installation by opening Microsoft Edge and browsing to `http://localhost`. From a remote computer, you can browse to the server's computer name. For example, if the server is named **SEA-SVR2**, then you could browse to `http://SEA-SVR2` to test the default website. The default installation of IIS includes a static HTML page in the default website so that administrators can verify the installation.

If the default webpage does not work, ensure that the IIS service is installed and started on the computer. You can do this by opening Windows PowerShell and running `Get-Service`.

In the resulting list, ensure that the IIS service is started. If you are attempting to connect from a remote computer, ensure that the web server's name resolves to an IP address by using DNS. Also ensure that Windows Firewall on the web server is configured to allow incoming connections. You can also try browsing to the server's IP address (for example, `http://172.16.10.13`) if DNS is unable to resolve the server's name to an IP address.

You can also use Windows PowerShell to check for IIS-related messages in the Windows event log. Run the following command:

```
Get-EventLog -LogName System -Source IIS
```

Review any warning or error messages that appear and take appropriate corrective actions.

IIS management

IIS Manager is included on all servers that have a graphical user interface (GUI) and that have IIS installed. IIS Manager is also available in the Remote Server Administration Tools (RSAT). The RSAT is available for different versions of the Microsoft Windows client operating system. As a best practice, you should install and run IIS Manager on the client computer and use it to connect to remote servers that run IIS. You could also use PowerShell to manage most aspects of IIS. PowerShell can be useful if you need to automate the administration of IIS. You could also manage IIS using the Microsoft IIS Web Manager extension for Windows Admin Center, which is currently in preview.

The **Start Page** displays after you select **Start Page** in the left tree view in IIS Manager. The tree view is also known as the **navigation** pane. The **Start Page** displays a list of servers that you connected to in the recent past and displays shortcuts for common tasks. Those tasks include connecting to the local computer, connecting to a server, and connecting to a specific site. The **Start Page** also displays links to online resources, such as the official Microsoft IIS.net website.

The **navigation** pane provides quick access to the main configuration containers in IIS. This includes sites and application pools. It also provides access to the global server configuration.

Examine the Web Server configuration page

If you select a server in the IIS Manager **navigation** pane, the right side of the window displays the global server configuration page. The global configuration page is a graphical representation of the data that is contained in three XML files:

- **Machine.config**. This file is located in `%windir%\Microsoft.NET\Framework\framework_version\CONFIG`.
- **Root Web.config for the .NET Framework**. This file is located in `%windir%\Microsoft.NET\Framework\framework_version\CONFIG`.
- **ApplicationHost.config**. This file is located in `%windir%\system32\inetsrv\config`.

You can modify the global server configuration either by using IIS Manager or by editing these files directly. Usually, administrators use IIS Manager because it provides a safer way to manage the server

configuration. When you directly edit the XML configuration files, it's easy to make a mistake and make the file unusable by IIS.

The server configuration provides the foundation for all site configurations. However, individual sites, directories, and applications can provide their own configurations, and those override the server configuration.

The server configuration is organized into several parts. Be aware that IIS is an extensible, modular service.

New role services or optional components might add more configuration items to the server configuration page. Important server configuration items include the following:

- **Authentication.** Enables and disables different authentication protocols.
- **Default Document.** Specifies the webpages that the server will deliver if a user sends a request to a website without requesting a specific webpage.
- **Error Pages.** Configures custom error pages for common HTTP errors.
- **Logging.** Defines the default logging behavior for websites.
- **SSL Settings.** Defines the default SSL settings for websites.

Default website configuration in IIS

Individual websites can override all, or parts of, the global server configuration. You can create a web.config file in the website's root folder. IIS Manager displays the contents of that file in a graphical form when you select a website in IIS Manager.

A website configuration page resembles the global server configuration page. It's organized into sections and it might contain additional sections that aren't included in the global server configuration. As you install role services or optional components, new configuration items might become available in each website's configuration page.

In addition to the configuration item icons, you will notice an **Actions** pane on the right side of the website configuration page. This pane includes several options that relate to the website's main configuration in addition to several management actions that you can perform in the website. These options and actions include the following:

- Exploring the website's file structure
- Editing the permissions of the website
- Modifying the website's basic settings
- Restarting or stopping the website

Websites can start or stop independently of one another and independently of the server. That is, IIS can continue to run some websites, while leaving other websites in a stopped status. A stopped website will not respond to incoming HTTP requests. Restarting a website stops it and then starts it again by using a single action. Stopping a website stops all processes that are responding to HTTP requests for that website, resetting the website and releasing any memory that the website was using.

Explore the default IIS folders

When you install IIS on a server, IIS creates a set of default folders. These are usually created on the system disk (**drive C** for most servers) and include the following:

- **\inetpub.** This is the root folder for the default IIS folder structure.

- **\Inetpub\Custerr.** This folder contains customizable error pages.
- **\Inetpub\Logs.** This folder contains one subfolder for each website that has logging enabled. Under those subfolders, you will find daily log files.
- **\Inetpub\Wwwroot.** This folder is the root folder for the default website content.

When you create new websites, you can decide to store their root folders under **\Inetpub** or you can store them in any other location that is available to the server. For example, your organization might adopt a standard server configuration that includes an additional disk (such as **drive E**) that you can use for website content. You can also store Website content in shared folders on the network, which enables you to place them on file servers.

Inside the default **\Inetpub\Wwwroot** folder, you will find files that create the default IIS webpage. This webpage is created only for the default website and it enables you to quickly verify that a new IIS installation works correctly. These files usually include the webpage **listart.htm** and a logo graphic **Welcome.png**.

By default, IIS is configured to use **listart.htm** as a default webpage. That is, if someone sends an HTTP request to the default website, but that user does not specify a specific webpage, IIS will respond by sending **listart.htm**. However, IIS defines other default webpages that have a higher priority than **listart.htm**. These other webpages include **Default.htm** and **Index.html**. Most web applications include one of these other default webpages. Therefore, when you add a web application's content to the default website's folder, the web application's default page will usually be used instead of **listart.htm**. Many administrators will leave **listart.htm** in the folder so that they can use it to verify the functionality of the website. It is a usual practice to remove **listart.htm** before putting the website into production.

Create a simple webpage

Although the default website includes a simple webpage that you can use to verify site functionality, other websites that you create will not. Therefore, it is useful to know how to create simple webpages that you can use to test website functionality.

You can use Windows File Explorer to create a new file. Open the website's root folder, right-click or access the context menu in the folder, and then select the command to create a new text file. Name the file **Default.htm** or similar.

After you create the new file, open it in Notepad or another text editor. You don't have to create a complex webpage. Add the following text to create a functional webpage:

```
<body>
  <h1>IIS 10.0 running on Windows Server 2019</h1>
  <p>This is a test page</p>
</body>
</html>
```

Save the file and verify that the complete filename is **Default.htm**. You can use any filename, but by using this name, you enable IIS to serve the webpage as a default webpage, so you don't have to specify the webpage by name.

Demonstration: Create and configure a new site in IIS

In this demonstration you will learn how to:

- Install the **Web Server** role using PowerShell.
- Verify the installation of the **Web Server** role.
- Create a new site in IIS (Microsoft Internet Information Service) and verify it.

Preparation steps

The required virtual machines (**WS-011T00A-SEA-DC1**, **WS-011T00A-SEA-ADM1**, **WS-011T00A-SEA-SVR1**, and **WS-011T00A-SEA-CL1**) should be running after the previous demonstration.

Demonstration steps

Install the Web Server role using PowerShell

1. On **SEA-ADM1**, open a Windows PowerShell command prompt, enter the following command, and then select Enter:

```
Install-WindowsFeature -ComputerName SEA-SVR3 -name Web-Server -IncludeManagementTools
```

Wait for the command to complete, which should take approximately 1 minute.

Verify the installation of the Web Server role

1. On **SEA-ADM1**, open Microsoft Edge, and in the address bar, enter `http://SEA-SVR3`.
2. Verify that the IIS displays the default webpage.
3. In the address bar, enter `http://172.16.10.14`.
4. Verify that IIS displays the default webpage.

Configure a website in IIS and verify it

1. Open the **Internet Information Services (IIS) Manager** console.
2. In **Internet Information Services (IIS) Manager**, in the console tree, navigate to **SEA-ADM1/Sites**, and then select **Default Web site**.
3. Configure site bindings by selecting **Contoso VPN as SSL Certificate**. When prompted, select **Yes**.
4. Close the **Internet Information Services (IIS) Manager** console.

After you complete this demonstration, revert all virtual machines that are running.

Test your knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

Which Windows kernel-mode driver is responsible for listening for incoming TCP (Transmission Control Protocol) requests in IIS?

- HTTP.sys
- www.dll
- webhost.sys
- svchost.exe

Question 2

Which two services are used by IIS (Microsoft Internet Information Service)? Choose two.

- Application Identity
- Software Protection
- World Wide Web Publishing Service
- Network Location Awareness
- Windows Process Activation Service

Question 3

Which server role must you choose to install IIS?

- Volume Activation Service
- Application Server role
- Remote Access
- Web Server (IIS)

Question 4

Which folder is the root folder for the default website content when you install IIS?

- \inetpub\Custerr
- \inetpub
- \inetpub\wwwroot
- \inetpub\logs

Module 10 lab and review

Lab: Deploying network workloads

Scenario

The employees in the IT department at Contoso need to be able to access server systems outside of business hours to correct issues that arise during weekends or holidays. Some of the employees are using computers that aren't members of the `contoso.com` domain. Other users are running non-Windows operating systems on their computers. To enable remote access for these users, you will provide remote access to Windows Admin Center, secure it with Web Application Proxy, and deploy a secure virtual private network (VPN) solution using the Secure Socket Tunneling Protocol VPN protocol.

You are a web server administrator for Contoso and your company is preparing to deploy a new intranet web application on an internal web server. You need to verify the server configuration and install Internet Information Service. The website must be accessible using a friendly Domain Name System name and all web connections to and from the server must be encrypted.

Objectives

After completing this lab, you'll be able to:

- Deploy and configure Web Application Proxy
- Implement a VPN solution
- Deploy and configure a web server

Estimated time: 60 minutes

Module review

Use the following questions to check what you've learned in this module.

Question 1

What are the requirements for the Windows 10 client using a device tunnel with Always On VPN (virtual private network)? Choose all that apply.

- Windows 10 Enterprise Edition
- Domain membership
- Group Policy
- Windows 10 Professional Edition
- A computer authentication certificate

Question 2

Why should you use only the IKEv2 (Internet Key Exchange version 2) or SSTP (Secure Socket Tunneling Protocol) VPN protocols with Always On VPN?

Question 3

What is the name of the script you can use to create the two configuration files for the Always On VPN client profile?

- VPN_Profile.ps1
- MakeProfile.ps1
- AlwaysOnVPNProfile.ps1
- VPN_Profile.xml files

Question 4

Does Always On VPN require IPv6 as was the case with DirectAccess?

Answers

Question 1

Which remote access feature in Windows Server automatically connects clients to the internal network when outside the office?

- Network Policy Server (NPS)
- Web Application Proxy
- Router
- DirectAccess

Explanation

DirectAccess is the correct answer. The DirectAccess feature in Windows Server automatically connects clients to the internal network when they are outside the office.

Question 2

Which information is required when publishing a web app with Web Application Proxy? (Choose three.)

- External URL of the applications
- A certificate from a private certification authority
- URL of the backend server
- Type of preauthentication

Explanation

The information required when publishing a web app with Web Application Proxy includes the external URL of the applications, the URL of the backend server, and the type of preauthentication.

Question 1

Which two types of VPN (virtual private network) connections does Windows Server 2019 support? Choose two.

- Remote access
- IKEv2 (Internet Key Exchange version 2)
- Point-to-site
- Site-to-site
- VPN reconnect

Explanation

Remote access and site-to-site are the correct answers. All other answers are incorrect. IKEv2 is an authentication protocol and VPN reconnect is a feature that leverages IKEv2. A point-to-site VPN is commonly used when you want to connect directly to Azure from your machine.

Question 2

Which types of site-to-site VPNs can you create in Windows Server? Choose three.

- PPTP (Point-to-Point Tunneling Protocol)
- IKEv2
- L2TP (Layer 2 Tunneling Protocol)
- SSTP (Secure Socket Tunneling Protocol)

Explanation

PPTP, IKEv2 and L2TP are the correct answers. All other answers are incorrect.

Question 3

Which VPN tunneling protocol should you use for your mobile users?

- PPTP
- L2TP
- SSTP
- IKEv2

Explanation

IKEv2 is the correct answer. All other answers are incorrect even though they could be used. 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. Furthermore it supports VPN reconnect. Consider a scenario where you are using a laptop that is running Windows 10. When you travel to work by train, you connect to the internet with a wireless mobile broadband card and then establish a VPN connection to your organization's network. When the train passes through a tunnel, you lose the internet connection. After the train emerges from the tunnel, the wireless mobile broadband card automatically reconnects to the internet.

Question 4

You configure your VPN server to use IP addresses from a DHCP server on your private network to assign those addresses to remote clients. How many IP addresses is it going to lease at a time?

- 2
- 10
- 25
- 5

Explanation

The correct answer is 10. All other answers are incorrect. The remote access VPN server will lease a block of 10 addresses at a time from the DHCP server and then assign those addresses to remote clients.

Question 1

Which of the following is a RADIUS (Remote Authentication Dial-In User Service) client?

- VPN (virtual private network) Server
- Wireless access point
- Windows 10 client
- Windows Server 2019 member server
- Remote Desktop (RD) Gateway server

Explanation

VPN Server, wireless access point and Remote Desktop (RD) Gateway server are the correct answers. A Windows 10 client or a Windows Server 2019 member server is not a RADIUS client. When using NPS (Network Policy Server) as a RADIUS server, you configure network access servers (NASs), such as wireless access points, VPN servers, and Remote Desktop (RD) Gateway servers, which are known as RADIUS clients in NPS.

Question 2

Which authentication protocols should you use in a production environment? Choose two.

- SPAP (Shiva Password Authentication Protocol)
- EAP
- PAP (Password Authentication Protocol)
- MS-CHAP
- CHAP
- MS-CHAP v2

Explanation

EAP and MS-CHAP v2 are the correct answers. All other answers are incorrect. You should not use PAP, SPAP, CHAP or MS-CHAP in a production environment as they are considered highly insecure.

Question 3

What kind of policies can you create on a Network Policy Server? Choose two.

- Connection Request Policies
- Group Policies
- Network Policies
- Configuration Policies

Explanation

Connection Request Policies and Network Policies are the correct answers. 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. All other answers are incorrect.

Question 1

Which of the following infrastructure components does Always On VPN require? Choose all that apply.

- Remote access server (virtual private network)
- Azure Virtual Network Gateway
- Group Policy
- Windows 10 clients
- Public Key Infrastructure
- Network Policy Server

Explanation

Remote access server (virtual private network or VPN), Windows 10 clients, public key infrastructure, and Network Policy Server are the correct answers. Azure Virtual Network Gateway is incorrect because Always on VPN is not supported on that platform. Even though you could use Group Policy for autoenrollment of certificates for Always On VPN users and clients, it is not required.

Question 2

Which methods does Always ON VPN support to create and manage a VPN profile? Choose three.

- Group Policy
- PowerShell
- Intune
- Microsoft Endpoint Configuration Manager

Explanation

PowerShell, Intune, and Microsoft Endpoint Configuration Manager are the correct answers.

Question 3

What is the name of the configuration item used to configure an Always On VPN profile?

- AlwaysOn.conf
- ProfileXML
- OMA-DM
- PEAP

Explanation

ProfileXML is the correct answer. It's the name of an URI node within the VPNv2 CSP. The "ProfileXML" node is used to configure Windows 10 VPN client settings and the XML data contains all information needed to configure a Windows 10 Always On VPN profile.

All other answers are incorrect. Open Mobile Alliance Device Management (OMA-DM) is a protocol used for mobile device management. AlwaysOn.conf doesn't exist and PEAP is used for providing secure communication.

Question 1

Which Windows kernel-mode driver is responsible for listening for incoming TCP (Transmission Control Protocol) requests in IIS?

- HTTP.sys
- www.dll
- webhost.sys
- svchost.exe

Explanation

HTTP.sys is the correct answer. All other answers are incorrect.

Question 2

Which two services are used by IIS (Microsoft Internet Information Service)? Choose two.

- Application Identity
- Software Protection
- World Wide Web Publishing Service
- Network Location Awareness
- Windows Process Activation Service

Explanation

World Wide Web Publishing Service and Windows Process Activation Service are the correct answers. All other answers are incorrect. IIS consists of various components that each perform functions for the web server and application such as reading configuration files and listening for requests made IIS. These components are either services or protocol listeners. The services include World Wide Web Publishing Service (WWW service) and Windows Process Activation Service (WAS).

Question 3

Which server role must you choose to install IIS?

- Volume Activation Service
- Application Server role
- Remote Access
- Web Server (IIS)

Explanation

Web Server (IIS) is the correct answer. All other answers are incorrect. The Remote Access role is used when you want to install a VPN server and the Volume Activation Service role is used when you want to install the Key Management Service (KMS). The option Application Server role does not exist.

Question 4

Which folder is the root folder for the default website content when you install IIS?

- \Inetpub\Custerr
- \Inetpub
- \Inetpub\Wwwroot
- \Inetpub\Logs

Explanation

\Inetpub\Wwwroot is the correct answer. All other answers are incorrect. \Inetpub is the root folder for the default IIS folder structure. \Inetpub\Custerr is the folder that contains customizable error pages. And \Inetpub\Logs is in the folder that contains one subfolder for each website that has logging enabled. Under those subfolders, you will find daily log files.

Question 1

What are the requirements for the Windows 10 client using a device tunnel with Always On VPN (virtual private network)? Choose all that apply.

- Windows 10 Enterprise Edition
- Domain membership
- Group Policy
- Windows 10 Professional Edition
- A computer authentication certificate

Explanation

Windows 10 Enterprise Edition, domain membership, and a computer authentication certificate are the correct answers. Windows 10 Professional Edition doesn't support tunnel mode and Group Policy is required even though it could be used for autoenrollment of the computer authentication certificates for Always On VPN clients.

Question 2

Why should you use only the IKEv2 (Internet Key Exchange version 2) or SSTP (Secure Socket Tunneling Protocol) VPN protocols with Always On VPN?

Because they are modern VPN protocols and considered secure. IKEv2 is designed for mobility and considered the most secure but could be blocked in firewalls. SSTP is also considered quite secure and is usually not blocked in firewalls because it uses port 443.

Question 3

What is the name of the script you can use to create the two configuration files for the Always On VPN client profile?

- VPN_Profile.ps1
- MakeProfile.ps1
- AlwaysOnVPNProfile.ps1
- VPN_Profile.xml files

Explanation

MakeProfile.ps1 is the correct answer. Running the MakeProfile.ps1 script will create the two configurations, VPN_Profile.ps1 and VPN_Profile.xml. The AlwaysOnVPNProfile.ps1 is fictitious.

Question 4

Does Always On VPN require IPv6 as was the case with DirectAccess?

Always On VPN works equally well with both IPv4 and IPv6, but is not dependent on IPv6 like DirectAccess.



Module 11 Server and performance monitoring in Windows Server

Overview of Windows Server monitoring tools

Lesson overview

Windows Server provides a range of tools to monitor an operating system and the applications on a computer. Use these tools to configure your system for efficiency and troubleshoot problems. Small and midsize organizations can use the monitoring tools in Windows Server to monitor their server infrastructure. However, enterprise organizations that deploy a more complex IT infrastructure will need a more complex monitoring and management solution, such as Microsoft System Center.

Lesson objectives

After completing this lesson, you'll be able to:

- Describe how **Task Manager** works.
- Describe the features of **Performance Monitor**.
- Describe the role of **Resource Monitor**.
- Describe **Reliability Monitor**.
- Describe **Event Viewer**.
- Describe how to monitor servers with **Server Manager**.
- Describe how to monitor servers with the **Windows Admin Center**.
- Explain how to use Sysinternals tools to monitor Windows Server.

Overview of Task Manager

Task Manager in Windows Server provides information to help you identify and resolve performance-related issues.

Note: **Task Manager** only enables monitoring of a local server.

Task Manager includes the following tabs:

- **Processes.** The **Processes** tab displays a list of running programs, which are subdivided into applications and internal processes of the Windows operating system. For each running process, this tab displays a summary of processor and memory usage.
- **Performance.** The **Performance** tab displays a summary of central processing unit (CPU) usage, memory usage, and network statistics.
- **Users.** The **Users** tab displays resource consumption on a per-user basis. You can also expand the user view to observe more detailed information about the specific processes that a user is running.
- **Details.** The **Details** tab lists all the running processes on the server, providing statistics about the CPU, memory, and consumption of other resources. You can use this tab to manage the running processes. For example, you can stop a process, stop a process and all its related processes, and change process priority values. By changing a process's priority, you determine how much of the CPU's resources the process can consume. By increasing the priority of a process, you allow the process to request more of the CPU's resources.
- **Services.** The **Services** tab provides a list of running Windows services and related information. The tab indicates whether a service is running and displays the process identifier (PID) of the running service. You can start and stop services by using the list on the **Services** tab.

You might consider using **Task Manager** when a performance-related problem arises. For example, you might examine the running processes to determine if a specific program is using excessive CPU resources.

Note: Always remember that **Task Manager** displays a snapshot of current resource consumption—you might need to examine historical data to determine a true picture of a server's performance and response under load.

Overview of Performance Monitor

Performance Monitor is a Microsoft Management Console (MMC) snap-in that you can use to obtain system performance information. Use this tool to analyze the performance effects that applications and services have on a computer, to obtain an overview of system performance, or to collect detailed information for troubleshooting.

Performance Monitor includes the following features:

- **Monitoring Tools.** This section contains **Performance Monitor**, which provides a visual display of built-in Windows performance counters, either in real time or as historical data. **Performance Monitor** includes the following features:
 - Performance counters: **Performance Monitor** uses performance counters to measure a system's state or activity. The operating system includes some performance counters, and individual applications might include other performance counters. **Performance Monitor** requests the current value of performance counters at a specified time interval, which by default is one second. Add performance counters to **Performance Monitor** by copying the counters or by creating a custom data collector set.
 - Multiple graph views: **Performance Monitor** features multiple graph views that enable you to visually review performance log data.
 - Custom views: Create custom views in **Performance Monitor** and then export them as data collector sets for use with performance and logging features.

- **Data collector sets.** A *data collector set* is a custom set of performance counters, event traces, and system configuration data. After you create a combination of data collectors that describe useful system information, save them as a data collector set and then run and observe the results.
- A data collector set organizes multiple data collection points into a single, portable component. Use a data collector set on its own, group it with other data collector sets, incorporate it into logs, or observe it in **Performance Monitor**.
- Configure a data collector set to generate alerts when it reaches thresholds. You can also configure a data collector set to run at a scheduled time, for a specific length of time, or until it reaches a predefined size. For example, run a data collector set for 10 minutes every hour during working hours to create a performance baseline. You can also set a data collector to restart when the collection reaches a set limit so that **Performance Monitor** creates a separate file for each interval. Scheduled data collector sets collect data regardless of whether you start **Performance Monitor**.
- Use data collector sets and **Performance Monitor** to organize multiple data collection points into a single component that you can use to review or log performance. **Performance Monitor** also includes default data collector set templates to help system administrators begin the process of collecting performance data.
- In **Performance Monitor**, under the **Data Collector Sets** node, use the **User Defined** node to create your own data collector sets. Specify the objects and counters that you want to include in the set for monitoring. To help you select appropriate objects and counters, use the following templates provided for monitoring:
 - **System Diagnostics.** This template selects objects and counters that report the status of hardware resources, system response times, and processes on the local computer, along with system information and configuration data. The report provides guidance on ways to optimize the computer's responsiveness.
 - **System Performance.** This template generates reports that detail the status of local hardware resources, system response times, and processes.
 - **WDAC Diagnostics.** This template enables you to trace the debug information for Windows Data Access Components.
 - **Basic.** This template creates a simple collector that you can add to later. It includes a processor performance counter, a simple configuration trace, and a Windows kernel trace object.
- **Reports.** Use the **Reports** feature to observe and generate reports from a set of counters that you create by using data collector sets. **Performance Monitor** creates a new report automatically every time a data collector set runs.

Common performance counters

Add many different performance counters to **Performance Monitor**. Some performance counters aren't used often. The following table contains the commonly used performance counters.

Table 1: Common performance counters

Counter	Usage
PhysicalDisk% Disk Time	This counter measures the percentage of time the disk was busy during the sample interval. If this counter rises more than 85 percent, the disk system is saturated.
PhysicalDisk\Avg. Disk Queue Length	This counter indicates how many I/O operations are waiting for the hard drive to become available. If the value is larger than two times the number of spindles, the disk itself might be the bottleneck. If this counter indicates a possible bottleneck, consider measuring the Avg. Disk Read Queue Length and Avg. Disk Write Queue Length to determine if read or write operations are causing the bottleneck.
Memory\Pages per Second	This counter measures the rate at which pages are read from or written to the disk to resolve hard page faults. If the value is greater than 1,000 as a result of excessive paging, a memory leak might exist.
Processor% Processor Time	This counter measures the percentage of elapsed time that the processor spends running a non-idle thread. If the percentage is greater than 85 percent, the processor is overwhelmed, and the server might require a faster processor.
System\Processor Queue Length	This counter indicates the number of threads in the processor queue. The server doesn't have enough processor power if the value is more than two times the number of central processing units (CPUs) for an extended period.
Network Interface\Bytes Total/Sec	This counter measures the rate at which bytes are sent and received over each network adapter, including framing characters. The network is saturated if more than 70 percent of the interface is consumed.
Network Interface\Output Queue Length	This counter measures the length of the output packet queue, in packets. Network saturation exists if this value is more than two.

Note: If your server is configured with solid state disks (SSDs), these disk counters are less relevant. It's unlikely that disk bottlenecks will occur in these configurations.

Overview of Reliability Monitor

Windows Server installs **Reliability Monitor** by default. It monitors hardware and software issues that occur during the selected time interval. Based on the number and type of issues, it assigns a number called a *stability index* that indicates the server's reliability. The stability index ranges from 1 to 10, where 1 represents the least-stable server state and 10 represents the most stable state. By using the stability index, administrators can quickly evaluate a server's reliability. Any issue that effects the server has the potential to change the value of the stability index.

Find the **Reliability Monitor** by accessing **Control Panel**, browsing to **Security and Maintenance**, and then selecting **Maintenance**. **Reliability Monitor** is represented with a **View Reliability History** link. By selecting this link, a **Reliability Monitor** window displays:

- A reporting history of the stability index values from previous days or weeks. The stability index information about application failures, Windows operating system failures, miscellaneous failures, and warnings are available.
- A reliability details table that has the source of the issue, summary information, the date, and the action taken.
- A group of actions that you can perform, which are represented as links in the console and include:
 - Saving the reliability history to an XML file. Use this option if you want to keep track of older reliability history information.
 - Starting the **Problem Reports** console. You can use this to monitor issues related to specific applications. For each problem that **Reliability Monitor** detects, options in the console allow you to get more details about the problem, check online for a solution, or to delete the reported problem information.
 - Checking for a solution for all reported problems. You can use this option if you want **Reliability Monitor** to connect to the internet to find online information about resolving all the reported problems.

Overview of Event Viewer

Event Viewer provides access to Windows Server event logs. Event logs provide information about system events that occur within the Windows operating system. These events include information, warnings, and error messages about Windows components and installed applications.

Event Viewer provides categorized lists of essential Windows log events, including application, security, setup, and system events. **Event Viewer** also provides log groupings for individual installed applications and specific Windows component categories. Individual events provide detailed information about the type of event that occurred. When an event occurs, **Event Viewer** provides details about the source of the event and detailed technical information to assist you in troubleshooting the event.

Additionally, **Event Viewer** allows you to consolidate logs from multiple computers onto a centralized server by using subscriptions. Finally, you can configure **Event Viewer** to run a specific action when a specified type of event occurs. This might include sending an email message, opening an application, or running a script.

Event Viewer in Windows Server contains the following important features:

- The ability to review multiple logs. You can filter for specific events across multiple logs. This makes it easier to investigate issues and troubleshoot the problems that might appear in several logs.
- The ability to customize views. You can use filtering to narrow searches to only the events in which you're interested, and you can save these filtered views.
- The ability to configure scheduled, automated tasks to run in response to events. **Event Viewer** is integrated with **Task Scheduler**.
- The ability to create and manage event subscriptions. You can collect events from remote computers and then store them locally.

Note: To collect events from computers, you must create an inbound rule in Windows Firewall to permit Windows Event Log Management.

Event Viewer tracks information in several different logs. These logs provide detailed information, including:

- A description of the event.
- An event ID number.
- The component or subsystem that generated the event.
- Information, warning, or error status.
- The time of the event.
- The user's name on whose behalf the event occurred.
- The computer on which the event occurred.
- A link to Microsoft Support or Microsoft Knowledge Base for more information about the type of event.

Windows Server logs

The following table lists several of the built-in **Event Viewer** logs.

Table 1: Windows Server logs

Built-in log	Description and use
Application log	This log contains errors, warnings, and informational events that pertain to the operation of applications such as Microsoft Exchange Server, the Simple Mail Transfer Protocol (SMTP) service, and other applications.
Security log	This log reports the results of auditing if you enable it. Audit events report success or failure depending on the event. For example, the log would report success or failure depending on whether a user was able to access a file.
Setup log	This log contains events that relate to application setup.
System log	Windows components and services log general events and classify them as errors, warnings, or information. The Windows operating system predetermines the events that system components log.
Forwarded events	This log stores events that Windows components collect from remote computers. To collect events from remote computers, you must create an event subscription.

Application and service logs

The **Applications and Services Logs** node stores events from a single application or component rather than events that might have system-wide effects. This category of logs includes four subtypes:

- Admin

- Operational
- Analytic
- Debug

Admin logs are of interest to administrators and support personnel who use **Event Viewer** to troubleshoot problems. These logs provide guidance about how to respond to issues. The events found in the Admin logs indicate a problem and a well-defined solution upon which an administrator can act.

Events in the Operational log are also useful for IT professionals, but they're likely to require more interpretation. You can use operational events to analyze and diagnose a problem or occurrence and to trigger tools or tasks based on the problem or occurrence.

Analytic and Debug logs aren't very user-friendly. Analytic logs store events that trace an issue, and they often log a high volume of events. Developers use Debug logs when they're debugging applications. By default, both Analytic and Debug logs are hidden and disabled.

Windows log files are 20,480 kilobytes (KB) in size, except the Setup log, which is 1,024 KB. The operating system overwrites events in the log files where necessary. If you want to clear a log manually, you must sign into the server as a local administrator.

If you want to configure event log settings centrally, you can do so by using Group Policy. Open the **Group Policy Management Editor** for your selected Group Policy Object, and then browse to **Computer Configuration\Policies\Administrative Templates\Windows Components\Event Log Service**.

For each log, you can define the following properties:

- The location of the log file
- The maximum size of the log file
- Automatic backup options
- Permissions on the logs
- Behavior that occurs when the log is full

Monitor a server with Server Manager

Organizations typically have multiple servers, both physical and virtual, that they must monitor. The number of servers in an organization depends on the organization's size and the complexity of its IT infrastructure. The most efficient way to monitor multiple servers is to deploy management and monitoring software that provides a centralized dashboard where administrators will be able to monitor all IT infrastructure components.

Depending on the size of the organization and the complexity of its IT infrastructure, monitoring software can be classified in two ways:

- Enterprise management and monitoring solutions, such as the Microsoft System Center suite of tools.
- Small and midsize organization monitoring solutions, such as **Server Manager**.

Windows Server installs **Server Manager** by default. You can also install it as a console on a Windows 10 client computer. It helps monitor both local and remote servers, collects monitoring data from specific servers, and presents the data in a centralized dashboard. Administrators can monitor up to 100 servers by using **Server Manager**. If you must monitor more than 100 servers, consider an enterprise monitoring solution such as System Center.

Server Manager can monitor both Desktop Experience and Server Core editions of Windows Server. Configuration for remote management and monitoring is enabled by default, but you can change it by

using **Server Manager** and Windows PowerShell on the monitored server. **Server Manager** doesn't support monitoring of the Windows client operating system.

When using **Server Manager**, you can perform the following monitoring tasks on remote servers:

- Adding remote servers to a pool of servers that **Server Manager** will monitor. Administrators can choose which servers to monitor.
- Creating custom groups of monitored servers. Administrators can group monitored servers in **Server Manager** by using different criteria, such as department, city, or country/region. Grouping servers helps organizations assign different administrators to monitor different groups of servers.
- Starting different tools on remote servers. Administrators can start different tools remotely, such as **Microsoft Management Console (MMC)** for monitoring different types of data or using PowerShell Remoting on remote servers. This ensures that administrators don't have to sign in locally to a server to perform different management tasks, such as starting a service.
- Determining server status and identifying critical events. **Server Manager** displays servers with critical issues on the centralized dashboard in the color red. This alerts administrators to start troubleshooting the issue immediately.
- Analyzing or troubleshooting different types of issues. You can configure centralized console monitoring information to display by type, such as Active Directory Domain Services, Domain Name System (DNS), Microsoft Internet Information Services (IIS), and remote access. This enables administrators to find an issue and begin troubleshooting it. The centralized console also provides general monitoring information that displays on the console as **All Servers**.
- Monitoring the status of the Best Practices Analyzer (BPA) tool. The BPA compares current server role configuration with recommended settings from Microsoft based on best practices. The centralized dashboard in **Server Manager** displays the results of the BPA from all the monitored servers.

Monitor a server with Windows Admin Center

Windows Admin Center is a web-based console that you can use to manage computers that are running Windows Server and Windows 10. Typically, you use **Windows Admin Center** to manage servers instead of using Remote Server Administration Tools.

Windows Admin Center works with any browser that's compliant with modern standards, and you can install it on computers that run Windows 10 and Windows Server.

Note: You can't install **Windows Admin Center** on a server computer that's configured as an Active Directory Domain Services (AD DS) domain controller.

With a few exceptions, **Windows Admin Center** supports almost all current Windows Server and Windows 10 administrative functionality.

Note: **Windows Admin Center** functionality updates regularly.

To use **Windows Admin Center**, you must first download and install it. You can download it from the Microsoft Download Center. After downloading and installing **Windows Admin Center**, you must enable TCP port 6516 on the local firewall.

Note: You're prompted to select a TCP port during setup.

The first time you run **Windows Admin Center**, you'll receive a prompt to select a certificate. Be sure to select the **Windows Admin Center** client certificate. Initially, only the local host computer displays as an available connection. You add more connections from the **Windows Admin Center** homepage. To do this, select **Add**, and then select the type of device that you want to add. The types are:

- Servers

- Windows PCs
- Failover clusters
- Hyperconverged clusters

After adding devices to the console, you can select a device to manage by selecting it from the **All connections** list. After selecting a device, you can manage it by using one of the preinstalled solutions. **Windows Admin Center** comes with a number of preinstalled solutions, which include:

- **Server Manager**
- **Computer Management**
- **Cluster Creation**
- **Cluster Manager**

Using the **Server Manager** feature in **Windows Admin Center**, you can manage various Windows Server components and features, including the following:

- Overview of current server health
- Observe and configure Active Directory objects and settings
- Enable and configure Azure Backup
- Enable and configure Azure File Sync
- Manage certificates
- Manage containers
- Manage devices
- Configure and manage the Dynamic Host Configuration Protocol (DHCP) service and the Domain Name System (DNS) service
- Configure and manage firewall and network settings
- Manage local users and groups
- Manage roles and features
- Manage scheduled tasks
- Manage services and storage
- Monitor performance statistics with **Performance Monitor**
- Manage Windows updates
- Observe the **System Insights** feature of Windows Server

Note: **System Insights** gives you increased insight into the functioning of your servers.

System Insights

This new feature is only available on Windows Server 2019. It uses a machine learning model to analyze Windows Server system data, including performance counters and events. These insights can help you predict future resource requirements in terms of computing, networking, and storage.

Note: You can use **Windows Admin Center** or Windows PowerShell to manage System Insights.

Use Windows Sysinternals tools to monitor servers

The Windows Sysinternals suite provides a collection of advanced investigative and troubleshooting tools. For server troubleshooting, the most useful tools are typically those that help identify process problems or performance issues.

Process tools

The Process Explorer and Process Monitor tools are part of the Windows Sysinternals suite:

- Process Explorer. This tool enables you to determine the currently active processes on a Windows computer. Depending on the mode, Process Explorer enables you to observe the:
 - Handles that the selected process has opened.
 - Dynamic-link libraries (DLLs) and memory mapped files that the process has loaded.
- Process Monitor. This is an advanced tool for monitoring a Windows computer. It displays real-time file system, registry, and process/thread activity. Process Monitor includes monitoring and filtering capabilities.

Performance tools

You can use one or more of the following tools to monitor performance:

- Contig. This tool enables you to quickly defragment your frequently used files.
- DiskMon. This tool enables the computer to capture all hard disk activity, and it acts like a software disk activity light in the system tray.
- PageDefrag. This tool enables you to defragment your paging files and registry hives.
- Process Explorer. This tool enables you to determine the files, registry keys, and other objects that processes have opened and the DLLs they have loaded. This tool also reveals who owns each process.
- Process Monitor. This tool enables you to monitor the file system, registry, process, thread, and DLL activity in real time.

Test Your Knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

If you wanted to observe the performance of the processor in your computer over a period, which tool would you use?

Question 2

*Which port does **Windows Admin Center** typically use?*

Using Performance Monitor

Lesson overview

You can use **Performance Monitor** to collect, analyze, and interpret performance-related data about your organization's servers. This enables you to make informed capacity-planning decisions. However, to make informed decisions, it's important to know how to establish a performance baseline, how to use data collector sets, and how to use reports to help you compare performance data to your baseline.

Lesson objectives

After completing this lesson, you'll be able to:

- Explain what a baseline is.
- Describe data collector sets.
- Describe how to capture counter data with a data collector set.
- Describe how to configure an alert.
- Describe how to observe **Performance Monitor** reports.
- Identify the key parameters that you should track when monitoring network infrastructure services.
- Identify considerations for monitoring virtual machines.
- Explain how to use **Windows Admin Center** to monitor server performance.

Overview of baselines, trends, and capacity planning

By calculating performance baselines for your server environment, you can interpret real-time monitoring information more accurately. A baseline for a server's performance indicates what performance-monitoring statistics indicate during normal use. You can establish a baseline by monitoring performance statistics over a specific period. When an issue or symptom occurs in real time, you can compare baseline statistics to real-time statistics and then identify anomalies.

Trends analysis

You should consider the value of performance data carefully to ensure that it reflects your server environment. Additionally, you should gather performance data that you can use to plan for business or technological growth, and then create upgrade plans. You might be able to reduce the number of servers that are in operation after measuring performance and assessing the required environment.

By analyzing performance trends, you can predict when existing capacity is likely to be exhausted. Review historical analysis along with your business requirements, and then use this data to determine when more capacity is required. Some peaks are associated with one-time activities, such as extremely large orders. Other peaks occur on a regular basis, such as monthly payroll processing. These peaks could make a capacity increase necessary to meet the demands of an increased number of employees.

Capacity planning

Planning for future server capacity is a best practice for all organizations. Planning for business changes often requires more server capacity to meet targets. By aligning your IT strategy with your business strategy, you can support business objectives. Furthermore, you should consider virtualizing your environment to reduce the number of physical servers that you require. You can consolidate servers by implementing the Hyper-V role in a Windows Server environment.

Capacity planning focuses on assessing server workload, the number of users that a server can support, and the ways to scale systems to support more workload and users in the future. New server applications and services affect the performance of an IT infrastructure. These services could receive dedicated hardware, although they often use the same local area network (LAN) and wide area network (WAN) infrastructure. Planning for future capacity should include all hardware components and how new servers, services, and applications affect the existing infrastructure. Factors such as power, cooling, and rack space are often overlooked during initial planning for capacity expansion. You should consider how your servers can scale up and out to support an increased workload.

Tasks such as upgrading to a newer version of Windows Server might affect the performance of your servers and network. An update can sometimes cause problems with applications that might be incompatible with Windows Server. Careful performance monitoring before and after applying updates can help identify and rectify these problems.

An expanding business can require an infrastructure to support a growing number of users. You should consider your organization's current and anticipated business requirements when purchasing hardware. This will help you to meet future business requirements by increasing the number of servers or by adding capacity to existing hardware when needed.

Additional capacity requirements can include:

- Adding more servers.
- Adding hardware.
- Reducing application loads.
- Reducing the number of users that connect to a server. You can do this by distributing users to multiple servers.

Understanding bottlenecks

A performance bottleneck occurs when a computer is unable to service requests for a specific resource. The resource might be a key component such as a disk, memory, processor, or network. Alternatively, the shortage of a component within an application package might cause the bottleneck. By regularly using performance-monitoring tools and comparing the results to your baseline and historical data, you can often identify performance bottlenecks before they affect users.

After identifying a bottleneck, you must decide how to remove it. Your options for removing a bottleneck include:

- Running fewer applications.
- Adding resources to the computer.

A computer that suffers from a severe resource shortage might stop processing user requests. This requires immediate attention. However, if your computer experiences a bottleneck but still works within acceptable limits, you might decide to defer any changes until you resolve the situation or have an opportunity to take corrective action.

Analyzing key hardware components

The four key hardware components are processor, disk, memory, and network. By understanding how your operating system uses these components and how they interact with one another, you'll better understand how to optimize server performance.

Processor

Processor speed is an important factor in determining your server's overall computing capacity. Processor speed can be defined as the number of operations that can be performed in a measured period. For example, a billion processor cycles per second is one gigahertz (GHz). Servers with multiple processors and processors with multiple cores generally perform processor-intensive tasks with greater efficiency and speed than a single processor or single-core processor computers. Processor architecture is also important.

Disk

Server hard disks store programs and data. Consequently, the throughput of hard disks affects the speed of a workstation or server, especially when the workstation or server is performing disk-intensive tasks. Most hard disks have moving parts, and it takes time to position the read/write heads over the appropriate disk sector to retrieve the requested information. Furthermore, disk controller performance and configuration also affect overall disk performance. By selecting faster disks and using disk arrays such as Redundant Array of Independent Disks (RAID) to optimize access times, you can alleviate the potential for a disk subsystem to create a performance bottleneck.

You should also remember that the data on a disk moves into the memory before it's used. If there is a surplus of memory, the Windows Server operating system creates a file cache for items that were recently written to or read from the disks. Installing more memory in a server can often improve disk subsystem performance because accessing the cache is faster than moving the information into memory.

Note: You can also improve disk performance by implementing solid-state drives (SSDs) or tiered storage.

Memory

Programs and data load from a disk into memory before a program manipulates the data. In servers that run multiple programs or where datasets are extremely large, increasing the amount of installed memory can help improve server performance.

Windows Server uses a memory model in which it doesn't reject memory requests by applications that exceed the computer's total available memory. Rather, it performs *paging* for these requests. During paging, Windows Server moves data and programs in memory that are currently not in use by the processors to the paging file, which is an area on the hard disk. This frees up physical memory to satisfy the excess requests. However, if a hard disk is comparatively slow, it has a negative effect on workstation performance. You can reduce the need for paging by adding more memory.

Network

The network is a critical part for performance monitoring because many network applications depend on the performance of network communications. Poor network performance can cause slow or unresponsive applications and server functionality. Therefore, network capacity planning is very important. While planning network capacity, you must consider bandwidth capacity and the capacity of any network

devices, such as router and switch capacity. In many cases, optimizing the configuration of network devices such as switches or routers improves the performance of a network and network applications.

What are data collector sets?

A *data collector set* is a custom set of performance counters, event traces, and system configuration data. A data collector set organizes multiple data-collection points into a single portable component. You can use a data collector set on its own or group it with other data collector sets. You can also incorporate a data collector set into logs or observe it in **Performance Monitor**. You can configure a data collector set to generate alerts when it reaches thresholds configured in performance counters.

Although it's useful to analyze current performance activity on a server computer, you might find it more useful to collect performance data over a set period and then analyze and compare it with data that you gathered previously. You can use this comparison to determine resource usage to plan for growth and to identify potential performance problems.

You can also configure a data collector set to run at a scheduled time for a specific length of time or until it reaches a predefined size. For example, you can run a data collector set for 10 minutes every hour during your working hours to create a performance baseline. You can also set a data collector to restart when it reaches set limits so that it creates a separate file for each interval.

You can configure a schedule for performance monitoring when configuring a data collector set. Scheduling options on the **Schedule** tab of the data collector set's properties window. The schedule monitoring options that you can select include beginning date, expiration date, and start time. You can also choose the day of the week you want performance monitoring to run.

After you've created a combination of data collectors that describe useful system information, you can save them as a data collector set, run the set, and observe the results.

Data collector sets can use the following types of data collectors:

- Performance counters. This data collector provides server performance data.
- Event trace data. This data collector provides information about system activities and events, and it's often useful for troubleshooting.
- System configuration information. This data collector allows you to record the current state of registry keys and to record changes to those keys.
You can create a data collector set from a template, from an existing set of data collectors in a **Performance Monitor** view, or by selecting individual data collectors, and then setting each individual option in the data collector set properties.

Demonstration: Capture counter data with a data collector set

In this demonstration, you'll learn how to:

- Create a data collector set.
- Create a load on the server.
- Analyze the resulting data in a report.

Demonstration steps

Create a data collector set

1. Switch to **SEA-ADM1**, and then sign in as **Contoso\Administrator** with the password **Pa55w.rd**.
2. Open **Performance Monitor**.
3. Create a new **User Defined** data collector set with the following key counters:
 - **Processor% Processor Time**
 - **Memory\Pages/sec**
 - **PhysicalDisk% Disk Time**
 - **PhysicalDisk\Avg. Disk Queue Length**
 - **System\Processor Queue Length**
 - **Network Interface\Bytes Total/sec**
4. Start the data collector set.

Create a disk load on the server

1. Open a Windows PowerShell command prompt, and then use the **fsutil** command to create a large file with a size of 104,857,600 bytes.
2. Copy the file to the **SEA-DC1** server to generate network load.
3. Create a new copy of the large file on the local hard disk by copying it from **SEA-DC1**.
4. Delete all the newly created files.

Analyze the resulting data in a report

1. Switch to **Performance Monitor**, and then stop the data collector set.
2. Select **Performance Monitor**, and then select **View Log Data**.
3. Add the data that you collected in the data collector set to the chart.
4. Change the view to **Report**.

Demonstration: Configure an alert

What are alerts?

Alert is Windows Server operating system functionality that notifies you when certain events occur or when certain performance thresholds are reached. You can configure alerts in Windows Server as network messages or as events that log in the application event log. You can also configure alerts to start applications and performance logs. You can configure alerts when you create data collectors by selecting the **Performance Counter Alert** type of data collector.

When you create an alert, configure the following settings:

- **Alert when.** This is the alert-threshold setting for a specific performance counter.

- **Alert Action.** This setting specifies whether to log an entry in the application event log or to start another data collector set.
- **Alert Task.** This setting specifies which command task to trigger when the alert threshold is reached. Additionally, you might specify command parameters if applicable.

In this demonstration, you'll learn how to:

- Create a data collector set with an alert counter.
- Generate a server load that exceeds the configured threshold.
- Examine the event log for the resulting event.

Demonstration steps

Create a data collector set with an alert counter

1. Create a new **User Defined** data collector set.
2. Select the **Performance Counter Alert** option, and then add only the **Processor% Processor Time** counter.
3. Set the threshold to be above **10** percent and generate an entry in the event log when this condition is met.
4. Start the data collector set.

Generate a server load that exceeds the configured threshold

1. Open File Explorer, and then browse to **C:\Labfiles\Mod11**.
2. Double-click or select the **CPUSTRES64.EXE** program, and then select Enter.
3. Configure the highlighted thread for Busy (**75%**).
4. When the tool has run for a minute, stop it, and then close CPUSTRES64.

Examine the event log for the resulting event

- Open **Event Viewer**, and then examine the Diagnosis-PLA log for performance alerts.

Demonstration: View reports in Performance Monitor

In this demonstration, you'll learn how to access a performance report.

Demonstration steps

View a performance report

1. In the navigation pane, expand **Reports/User Defined/SEA-ADM1 Performance**.
2. Expand the folder under **SEA-ADM1 Performance**.

Note: The data collector set's previous collection process generated this report. You can change from the chart view to any other supported view.

3. If the report isn't displaying, select the **Refresh** button on the toolbar, and then repeat step 2.
4. Close all open windows.

Monitor network infrastructure services

Network infrastructure services are an essential foundation of many other server-based services. Therefore, you should correctly configure them and ensure that they run optimally.

Your organization can benefit in several ways by gathering performance-related data on your network infrastructure services.

- Optimizing network-infrastructure server performance. Providing a performance baseline and trend data enables you to help your organization optimize your network infrastructure servers' performance.
- Troubleshooting servers. When server performance degrades, either over time or during periods of peak activity, you can help identify possible causes and take corrective action. This can help you quickly bring a service back within the limits of your service level agreement.

Monitoring DNS

Domain Name System (DNS) provides name-resolution services on a network. You can monitor the Windows Server DNS Server role to determine the following aspects of your DNS infrastructure, including the:

- General DNS server statistics, including the number of overall queries and responses that a DNS server is processing.
- User Datagram Protocol (UDP) or Transmission Control Protocol (TCP) counters, which measure DNS queries and responses that the DNS server processes by using either of these transport protocols.
- Dynamic update and secure dynamic-update counters for measuring registration and update activity that dynamic clients generate.
- Memory-usage counter for measuring a system's memory usage and memory-allocation patterns that are created by operating the server computer as a DNS server.
- Recursive lookup counters for measuring queries and responses when the DNS Server service uses recursion to look up and fully resolve DNS names on behalf of requesting clients.
- Zone transfer counters, including specific counters for measuring all zone transfer (AXFR), incremental zone transfer (IXFR), and DNS zone-update notification activity.

Note: You can also perform basic DNS monitoring by using the **DNS** console.

Monitoring DHCP

The Dynamic Host Configuration Protocol (DHCP) service provides dynamic IP configuration services for your network, and it provides data on a DHCP server, including the:

- The Average Queue Length counter indicates the current length of a DHCP server's internal message queue. This number represents the number of unprocessed messages that the server receives. A large number might indicate heavy server traffic.
- The Milliseconds per packet counter is the average time that a DHCP server uses to process each packet that it receives. This number varies depending on the server hardware and its I/O subsystem. A

spike indicates a problem with the I/O subsystem becoming slower or because of intrinsic processing overhead on the server.

Note: You can also perform basic DHCP monitoring by using the **DHCP Manager** console.

Considerations for monitoring virtual machines

Server virtualization has been part of the Windows Server operating system since the release of Windows Server 2008 and the introduction of the Hyper-V role. Many organizations have migrated some or all of their server workloads to virtual machines (VMs) that are running virtualization servers. From a monitoring perspective, it's important to remember that servers running as guest VMs consume resources in the same way as physical host-server computers.

With Hyper-V server virtualization, you can create separate VMs and run them concurrently by using the resources of the operating system that's running on a single physical server. The operating systems running within each VM are guests, while the computer that's running Hyper-V is the host.

Guest VMs function as physical computers. Guest VMs that are hosted on the same hypervisor remain independent of one another. If the host server has enough resources, you can simultaneously run multiple VMs that are using different operating systems on a host server.

When you create a VM, you configure characteristics that define the available resources for that guest. These resources include memory, processors, disk configuration, and network adapter configuration. These VMs operate within the boundaries of the resources that you allocate to them, and they can suffer from the same performance bottlenecks as host servers. Therefore, it's important to monitor VMs in the same way you monitor host servers.

Note: In addition to monitoring guest VMs, always remember that you must monitor the host that runs them.

Microsoft provides a tool, Hyper-V Resource Metering, that enables you to monitor resource consumption on VMs. Hyper-V Resource Metering allows you to track the resource utilization of VMs hosted on Windows Server computers that have the Hyper-V role installed.

With Hyper-V Resource Metering, you can measure the following parameters on individual Hyper-V VMs:

- Average graphics processing unit (GPU) use
- Average physical memory use, including:
 - Minimum memory use
 - Maximum memory use
- Maximum disk-space allocation
- Incoming network traffic for a network adapter
- Outgoing network traffic for a network adapter

By measuring how much of these resources each VM uses, an organization can bill departments or customers based on their hosted VM use rather than charging a flat fee per VM. An organization with only internal customers can also use these measurements to observe patterns of use and to plan future expansions.

You perform resource-metering tasks by using Windows PowerShell cmdlets with the Hyper-V module for Windows PowerShell. There's no graphical user interface (GUI) tool to perform this task. You can use the following cmdlets to perform resource metering tasks:

- **Enable-VMResourceMetering.** This cmdlet starts collecting data on a per-VM basis.

- **Disable-VMResourceMetering.** This cmdlet disables resource metering on a per-VM basis.
- **Reset-VMResourceMetering.** This cmdlet resets VM resource-metering counters.
- **Measure-VM.** This cmdlet displays resource-metering statistics for a specific VM.

Performance monitoring with Windows Admin Center

You can use **Windows Admin Center** to perform many of the same tasks that you might perform by using **Server Manager**. These include the following features:

- The **Overview** tab helps you observe current performance details similar to **Task Manager**.
- The **Performance Monitor** tab allows you to compare performance counters for Windows operating systems, apps, or devices in real time.

Using Windows Admin Center

You can use the following procedure to monitor the four core resources on a remote computer by using **Windows Admin Center**:

1. Open **Windows Admin Center**, and then sign in as necessary.
2. Select the appropriate remote server.
3. Select the **Overview** tab, and then in the details pane, scroll down and review the details for:
 - CPU
 - Memory
 - Ethernet

If you also want to monitor disk performance, you must enable disk metrics.

Note: Enabling disk metrics can affect performance.

To monitor disk performance:

1. In the details pane, select **Enable Disk Metrics** on the menu bar.
2. When prompted, select **Yes**.
3. In the details pane next to **Ethernet**, you can observe the details for installed disks.

Note: After you finish, disable disk metrics.

Note: Remember that you are monitoring real-time statistics. For more meaningful data, collect data by using **Performance Monitor**.

Using Performance Monitor

Use the following procedure to access **Performance Monitor**:

1. Open **Windows Admin Center**, and then sign in as necessary.
2. Select the appropriate remote server.
3. In the navigation pane, select **Performance Monitor**.
4. In the details pane, select **Blank workspace**.

5. Select **Add counter**, and then add the desired performance **Objects, Instances**, and **Counters**.
6. Choose the Graph type: **Line, Report, Min-Max**, and **Heatmap**.
7. To define the workspace name, time range, and other details, select **Settings**.
8. Save the workspace for future use.

Note: You can use **Windows Admin Center** to add the same objects, counters, and instances that you can add with **Performance Monitor**, which this module discussed earlier.

Test Your Knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

What's the purpose of creating a baseline?

Question 2

*To use **Windows Admin Center** to measure server performance, you need to measure disk performance. What must you do?*

Monitoring event logs for troubleshooting

Lesson overview

Event Viewer provides a convenient and accessible location for you to observe events that occur. Windows Server records events in one of several log files based on the type of event that occurs. To support your users, you should know how to access event information quickly and conveniently, and you should know how to interpret the data in the event log.

Lesson objectives

After completing this lesson, you'll be able to:

- Describe how to use **Server Manager** to review event logs.
- Explain what a custom view is.
- Describe how to create a custom view.
- Explain what event subscriptions are.
- Describe how to configure an event subscription.

Use Server Manager to review event logs

Server Manager provides a centralized location in which you can store and access event logs for multiple remote servers that you're monitoring. **Server Manager** provides a monitoring and troubleshooting solution in which administrators can review, in one console, information regarding specific events from different servers and applications. This is more efficient compared to viewing event logs by connecting to a specific server from a remote location.

You can review **Server Manager** event logs for all servers, for a specific server, or per server role, such as Active Directory Domain Services (AD DS), Domain Name System (DNS), or remote access. You can choose different event log views from the **Server Manager** navigation pane:

- **Local Server.** This view displays event logs that are on the local server where **Server Manager** is running. By default, **Application**, **Security**, and **System** event logs are displayed.
- **All Servers.** This view displays event logs from all servers that **Server Manager** is monitoring.
- **AD DS, DNS, and Remote Access.** This view displays event logs from all servers that **Server Manager** is monitoring and that have specific server roles installed, such as AD DS, DNS, or the Remote Access role. These logs display specific information that the AD DS, DNS, or the Remote Access server roles generate.
- **Roles and Server Groups tiles in Server Manager Dashboard.** To display the events for a specific server role, you can also choose an events link in a specific server group tile, such as the **AD DS** tile, **DNS** tile, or **Remote Access** tile, in **Server Manager Dashboard**.

You can further customize event log views by:

- Creating queries for specific types of events that must display. You can save these queries and use them later when you're searching for events that are defined in the query criteria.
- Configuring event data that needs to display. You can choose what type of events to display, such as **Critical**, **Error**, **Warning**, and **Informational**. Additionally, you can choose the event log files from where the events will display, such as Application, Directory Service, DNS Server, Security, System, and Setup.

What is a custom view?

Event logs contain vast amounts of data and narrowing the set of events to just those events that interest you can be a challenge. Custom views allow you to query and sort just the events that you want to analyze. You can also save, export, import, and share these custom views.

Event Viewer allows you to filter specific events across multiple logs and display all events that might relate to an issue that you're investigating. To specify a filter that spans multiple logs, you must create a custom view. You create custom views in the Action pane in **Event Viewer**.

You can filter custom views based on multiple criteria, including the:

- Time that the event logged.
- Event level, including errors or warnings.
- Logs from which to include events.
- Specific event IDs to include or exclude.
- User context of the event.
- Computer on which the event occurred.

Demonstration: Create a custom view

This demonstration, you'll learn how to:

- Examine server roles custom views.
- Create a custom view.

Demonstration steps

Examine server roles custom views

- In **Event Viewer**, examine predefined **Server Roles** custom views.

Create a custom view

1. Create a new custom view to select the following event types:
 - **Critical**
 - **Warning**
 - **Error**
2. Select the following logs:
 - **System**
 - **Application**
3. Name the custom view **Contoso Custom View**.
4. Observe the resulting filtered events in the details pane.

What are event log subscriptions?

Event log subscriptions enables a single server to collect copies of events from multiple systems. Using the Windows Remote Management (WinRM) and Windows Event Collector services (Wecsvc), you can collect events in the event logs of a centralized server, where you can analyze them together with the event logs of other computers that are being collected on the same central server.

Subscriptions can either be collector-initiated or source computer-initiated:

- **Collector-initiated.** A collector-initiated subscription, or a *pull subscription*, identifies all the computers from which the collector will receive events and typically pulls events from these computers. In a collector-initiated subscription, the subscription definition is stored and maintained on the collector computer. You use pull subscriptions when you must configure many of the computers to forward the same types of events to a central location. In this manner, you must define and specify only one subscription definition to apply to all computers in the group.
- **Source computer-initiated.** In a source computer-initiated subscription, or a *push subscription*, source computers push events to the collector. In a source computer-initiated subscription, you create and manage the subscription definition on the source computer, which is the computer that's sending events to a central source. You can define these subscriptions manually or by using Group Policy. You create push subscriptions when each server is forwarding a set of events different from what the other servers are forwarding or when you must maintain control over the event-forwarding process at the source computer. This might be the case when you must make frequent changes to the subscription.

To use event log subscriptions, you must configure the forwarding and the collecting computers. The event-collecting functionality depends on the WinRM service and Wecsvc. Both of these services must be running on computers that are participating in the forwarding and collecting process.

Enabling subscriptions

To enable subscriptions, perform the following tasks:

1. On each source computer, run the following command at an elevated command prompt to enable WinRM:

```
winrm quickconfig
```
2. On the collector computer, enter the following command at an elevated command prompt to enable Wecsvc:

```
wecutil qc
```
3. Add the computer account of the collector computer to the local **Event Log Readers** group on each of the source computers.

Note: You can also use the **Event Monitor** console and Group Policy to enable and configure event subscriptions.

Demonstration: Configure an event subscription

In this demonstration, you'll learn how to:

- Configure a source computer.
- Configure a collector computer.

- Create and review the subscribed log.

Demonstration Steps

Configure a source computer

1. Switch to **SEA-ADM1**, and if necessary, sign in as **Contoso\Administrator** with the password **Pa55w.rd**.
2. Run the **winrm quickconfig** command.
Note: The service is already running.
3. Open **Computer Management**, and then add the **SEA-CL1** computer as a member of the domain local **Event Log Readers** group.

Configure a collector computer

1. Switch to **SEA-CL1**, and then open a **Command Prompt** window.
2. Run the **wecutil qc** command.

Create and review the subscribed log

1. Switch to **Event Viewer**.
2. Create a new subscription to collect the following types of events from **SEA-ADM1**:
 - **Collector initiated**
 - **Source computer SEA-ADM1**
 - **All events types**
 - **Last 30 days**

Test Your Knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

What group memberships must you change to establish an event subscription?

Question 2

*On which computer must you run the **wecutil qc** command when establishing an event subscription?*

Module 11 lab and review

Lab 11: Monitoring and troubleshooting Windows Server

Scenario

Contoso, Ltd is a global engineering and manufacturing company with its head office in Seattle, Washington, in the United States. An IT office and datacenter are in Seattle to support the Seattle location and other locations. Contoso recently deployed a Windows Server 2019 server and client infrastructure.

Because the organization deployed new servers, it's important to establish a performance baseline with a typical load for these new servers. You've been asked to work on this project. Additionally, to make the process of monitoring and troubleshooting easier, you decided to perform centralized monitoring of event logs.

Objectives

After completing this lab, you'll be able to:

- Establish a performance baseline.
- Identify the source of a performance problem.
- Review and configure centralized event logs.

Estimated time: 40 minutes

Module review

Use the following questions to check what you've learned in this module.

Question 1

*What significant counters should you monitor in **Performance Monitor**?*

Question 2

Why is it important to monitor server performance periodically?

Question 3

Why should you use performance alerts?

Answers

Question 1

If you wanted to observe the performance of the processor in your computer over a period, which tool would you use?

*Although **Task Manager** and **Resource Monitor** provide performance detail, you can't observe data for a long period. **Performance Monitor data collector sets** would be better.*

Question 2

Which port does **Windows Admin Center** typically use?

TCP port 6516.

Question 1

What's the purpose of creating a baseline?

You can use a baseline to compare current performance with historic performance data.

Question 2

To use **Windows Admin Center** to measure server performance, you need to measure disk performance. What must you do?

You must enable disk metrics to collect data about disk performance.

Question 1

What group memberships must you change to establish an event subscription?

*You must add the computer account of the collector computer to the local **Event Log Readers** group on each of the source computers.*

Question 2

On which computer must you run the **wecutil qc** command when establishing an event subscription?

*You run that command on the collector computer to enable the **Webserv** service.*

Question 1

What significant counters should you monitor in **Performance Monitor**?

*You should monitor **Processor\% Processor Time**, **System\Processor Queue Length**, **Memory\Pages/sec**, **Physical Disk\% Disk Time**, and **Physical Disk\Avg. Disk Queue Length**.*

Question 2

Why is it important to monitor server performance periodically?

By monitoring server performance, you can perform capacity planning, identify and remove performance bottlenecks, and assist with server troubleshooting.

Question 3

Why should you use performance alerts?

By using alerts, you can react more quickly to emerging performance-related problems, perhaps before they impinge on users' productivity.



Module 12 Upgrade and migration in Windows Server

AD DS migration

Lesson overview

You can use Windows Server 2019 for domain controllers in Active Directory Domain Services (AD DS). You can upgrade an existing AD DS forest to use domain controllers running Windows Server 2019 or migrate to a new AD DS forest. Most organizations upgrade the existing AD DS forest. However, if you decide to migrate to a new AD DS forest, you can use the Active Directory Migration Tool (ADMT).

Lesson objectives

After completing this lesson, you'll be able to:

- Compare upgrading an AD DS forest and migrating to a new AD DS forest.
- Describe how to upgrade an existing AD DS forest.
- Describe how to migrate to a new AD DS forest.
- Describe ADMT.

Upgrade vs. migration

Unlike most new versions of Windows Server, Windows Server 2019 doesn't introduce new domain and forest functional levels. The highest available domain and forest functional levels are Windows Server 2016. If your Active Directory Domain Services (AD DS) forest is at the Windows Server 2016 functional level, you still might want to implement domain controllers running Windows Server 2019 to use the latest operating system.

Most organizations add domain controllers running Windows Server 2019 to their existing AD DS forest. Using Domain Controllers prior to Windows Server 2016 allows you to raise the domain level and the forest level for AD DS. When you upgrade an AD DS forest, there's no downtime for users or resources, and the same AD DS structure is maintained.

Migrating to a new AD DS forest is more complex than upgrading an existing AD DS forest. When you perform an AD DS migration, you create a new AD DS forest with new domain controllers running Windows Server 2019. The migration process involves moving users, groups, computers, servers, and applications to the new AD DS forest. During the migration, you need to plan for coexistence and ensure that users maintain access to resources in the source and destination environment.

Migration to a new AD DS forest is typically driven by a need to restructure AD DS rather than a need to use domain controllers running Windows Server 2019 or a higher AD DS functional level. Migration is often used to merge multiple domains into a single domain for simplified management, but this can be done within a single AD DS forest. Common reasons to migrate to a new AD DS forest include:

- An acquisition or merger.
- Divestiture of a company or business unit.
- Need to rename the AD forest or domain.

Note: You can rename a domain by using `Random.exe` without performing an AD DS migration, but this is a risky operation, and many organizations prefer to perform a migration instead. There are also some apps that fail after a domain rename. For example, if you have an Exchange Server in the AD DS forest, you can't rename the domain.

Upgrade a previous version of AD DS to Windows Server 2019

To raise your Active Directory Domain Services (AD DS) forest and domains to the Windows Server 2016 functional level, you must ensure that domain controllers are running Windows Server 2016 or Windows Server 2019. If you have domain controllers running older versions of Windows Server, you must remove or upgrade them. However, before you remove the older domain controllers, you must add new domain controllers running Windows Server 2016 or Windows Server 2019. If your AD DS forest and domains are already running at the Windows Server 2016 functional level, you can still upgrade to using Windows Server 2019 for your domain controllers.

To add a domain controller running Windows Server 2019, the existing AD DS domain level and forest level must be at a minimum of Windows Server 2008. Even if all domain controllers are running newer versions of Windows Server, the domain level and forest level requirement must be met.

Domain controllers running Windows Server 2019 use Distributed File System (DFS) Replication for SYSVOL replication. If you haven't updated SYSVOL replication in the domain to use DFS Replication, you must do that before adding the first domain controller running Windows Server 2019. Domain controllers running Windows Server 2008 and newer can use DFS Replication for SYSVOL replication.

You can verify that the SYSVOL replication has been updated to DFS Replication by running **`dfsrmig / getmigrationstate`**. A status of Eliminated indicates that file replication service (FRS) isn't in use.

A final prerequisite for adding Windows Server 2019 domain controllers is preparing the forest and the domain. Forest preparation includes extending the schema. To do this, use the `Adprep.exe` tool included on Windows Server 2019 installation media.

To prepare the forest, run:

```
Adprep /forestprep
```

To prepare each domain in the AD DS forest, run:

```
Adprep /domainprep
```

To prepare the forest, you must be a member of Enterprise Admins and Schema Admins. To prepare a domain, you must be a member of Domain Admins.

Note: If you don't manually prepare the forest and domains, the domain controller promotion process performs this step automatically.

After you add new domain controllers running Windows Server 2019, you can remove domain controllers running the previous version of Windows Server. Within each domain, after all domain controllers are running Windows Server 2016 or newer, you can raise the domain level to Windows Server 2016. After all domains in the forest have been raised to the Windows Server 2016 level, the forest can be raised to the Windows Server 2016 level.

Before you remove domain controllers running previous versions of Windows Server, you need to understand how they're being used. You must plan for updating any clients using those servers for services other than Windows authentication. For example, if an application is configured to use a specific domain controller for Lightweight Directory Access Protocol (LDAP) authentication, you must reconfigure the application to use a different domain controller. DNS is another common service provided by domain controllers that you should consider.

Migrate to AD DS in Windows Server 2019 from a previous version

Migrating to a new Active Directory Domain Services (AD DS) forest requires a high level of effort and extensive planning. A key concern during the migration process is ensuring that users retain access to resources in the source AD DS forest and the destination AD DS forest. In a large organization, the planning and migration process can take many months.

Careful planning is required before you can migrate to a new AD DS forest. Planning includes:

- Select new names. Most migration tools require the source and destination AD DS forests to have unique domain names, NetBIOS names, and user principal name (UPN) suffixes.
- Plan organizational unit (OU) structure. You can replicate the OU structure in the source AD DS forest, but this is an opportunity to restructure to better meet organizational needs if you choose to.
- Plan Group Policies. Evaluate which Group Policy settings should be applied in the new AD DS forest and how they'll be applied in the new OU structure.
- Identify objects to migrate. Depending on the scope of the project, only a subset of users, groups, computers, and servers might be migrated. For these objects, you also need to determine which attributes should be migrated or excluded.
- Identify apps that will be migrated. Apps can have many interdependent components, so you must also identify their components. Only after accurate identification can you plan how each app will be migrated.

After planning is complete, you can create the new AD DS forest by using new domain controllers running Windows Server 2019. For some apps, you'll need to perform schema extensions to prepare the new forest to host those apps. You might also need to install new instances of apps before migrating app data.

To support a phased migration that allows users to access resources in the source and destination, you must configure a forest trust. This allows users in the source forest to access resources migrated to the target forest. It also allows users migrated to the target forest to access resources in the source forest. After a user is migrated to the target forest, the **sidHistory** attribute is used to maintain access to resources in the source forest. The **sidHistory** attribute on the migrated user is updated to include the security identifier (SID) of the source user. This allows the migrated user to access resources in the source

forest to which the source user had been assigned access. Populating **sidHistory** allows the new migrated user to impersonate the source user for security purposes.

To use **sidHistory**, you must disable SID filtering on the forest trust. SID filtering removes SIDs from the local domain that weren't issued by a domain controller in the local AD DS forest.

Password synchronization is another important part of migrating. To simplify the migration process for users, you should migrate passwords along with the user accounts.

Overview of Active Directory Migration Tool

To move objects between Active Directory Domain Services (AD DS) forests and domains, Microsoft provides the Active Directory Migration Tool (ADMT). You can use this tool when consolidating domains within a forest or when migrating to a new AD DS forest. This tool hasn't been updated recently, but it still works with Windows Server 2019.

ADMT can perform the following functions:

- Migrate user accounts
- Migrate service accounts
- Migrate groups
- Migrate computer accounts
- Translate security for local profiles

If you're defining a new naming structure for users or groups, you can implement that as part of the migration process. Use an include file, which identifies the desired source objects and the destination names. Computer accounts can't be renamed during migration.

When you migrate user objects, all attributes are migrated by default. You can filter some attributes out in the migration process if the attributes value is no longer valid. For example, attributes that retired applications configured could be filtered. If there are application-specific attributes, you'll need to research the best approach for migrating them.

ADMT installation

ADMT installs on a member server with Desktop Experience in the target AD DS forest. Before installing ADMT, you must install Microsoft SQL Server to hold migration information. You can use SQL Server Express for ADMT, but you should monitor the size of the database to ensure that it doesn't reach its maximum limit and stop functioning.

If you're using password synchronization, you must install Password Export Server (PES) on a domain controller in the source. PES is responsible for exporting user password hashes from the source to the target. If you don't use PES, migrated user accounts are configured with a new password stored in a text file. Without password synchronization, you need a process to distribute new passwords to users.

Migration accounts

Migration accounts are user accounts in the source; they target forests with enough permissions to perform migration tasks. Accounts that are members of Domain Admins in the source and target forests will work, but you can create accounts with only the necessary permissions delegated for specific tasks such as migrating users or computers. The migration account in the target forest must be a local administrator on the member server where ADMT is installed.

Security translation

Translating security in local profiles on client computers allows a migrated user to retain access to the same profile on the local computer as the source user. For users, this means that their profiles stays intact with all their app configurations.

Note: ADMT was developed to work with Windows 7 and earlier. It hasn't been updated to work with Window 8 or Windows 10. This means that profile translation might not work properly, depending on the configuration of computers and apps in your organization. For detailed information about using ADMT, obtain Active Directory Migration Tool (ADMT) Guide: Migrating and Restructuring Active Directory Domains from the Microsoft Download Center.

Test your knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

Which of the following are true about upgrading Active Directory Domain Services (AD DS)? Choose two.

- The highest available forest functional level is Windows Server 2012 R2.
- All domain controllers must be running Windows Server 2019 or Windows Server 2016 before you can raise the forest functional level to Windows Server 2016.
- Upgrading AD DS to use domain controllers running Windows Server 2019 can be done without any downtime.
- Upgrading AD DS is very complex and requires downtime.

Question 2

What is the minimum forest functional level required to support a domain controller running Windows Server 2019?

- Windows Server 2003 R2
- Windows Server 2008
- Windows Server 2008 R2
- Windows Server 2012
- Windows Server 2016

Storage Migration Service

Lesson overview

You can use Storage Migration Service to migrate files and file shares from existing file servers to new servers running Windows Server 2019. If you choose to move the identity from the source server to the destination server, the client connectivity to the shares is preserved during the migration. Before you create a job that identifies how the migration will be performed, configure an orchestrator server to manage the migration process. In the job, specify which volumes and shares you need to migrate.

Lesson objectives

After completing this lesson, you'll be able to:

- Describe Storage Migration Service and its usage scenarios.
- Identify the requirements for using Storage Migration Service.
- Describe how to migrate a server with storage migration.
- List the considerations for using Storage Migration Service.

Storage Migration Service overview and usage scenarios

Storage Migration Service migrates data from multiple sources to either an on-premises Windows Server or a virtual machine (VM) in Microsoft Azure. The graphical management interface is integrated as part of Windows Admin Center. The primary use case for Storage Migration Service is to migrate an existing file server to a new file server. If you're migrating storage to Windows Azure, Storage Migration Service can automate the creation of a VM as the destination for the data.

Note: Storage Migration Service can't be used to combine multiple file shares into a single destination file share.

A key benefit of Storage Migration Service is that it assigns the identity of the source server to the target server, including the server name and the server IP addresses. This means that clients configured to access a share on the source server can automatically begin using the migrated data on the target servers. You don't need to update drive mappings or file share names in scripts.

Storage Migration Service can also migrate local user accounts. This can be useful if you have local user accounts created for administrative access or applications.

The general process for using Storage Migration Service is:

1. Inventory source servers
2. Transfer data
3. Cut over identities

After cutting over, the source servers are still functional but aren't accessible to users and apps at the original names and IP addresses. The files are still available to the administrators if required, and you can decommission the source servers when you're ready.

Storage migration requirements

The console for Storage Migration Service is a desktop computer or server running Windows Admin Center. Windows Admin Center provides the user interface for configuring Storage Migration Service but doesn't manage the work. Alternatively, you can configure Storage Migration Service by using Windows PowerShell cmdlets.

Orchestrator server

To manage the migration work, you need an orchestrator server running Windows Server 2019. This is the server you install the Storage Migration Service feature on. If you're migrating only one server, you can use the destination server as the orchestrator server. If you're migrating multiple servers, you should use a dedicated orchestrator server.

The orchestrator server should be in the same Active Directory Domain Services (AD DS) domain as the source and destination computers. The cutover process works across domains, but the source fully qualified domain name (FQDN) can't be migrated to a different domain.

Note: To support migrating data to or from a Windows failover cluster, you must install the Failover Clustering tools on the orchestrator server.

The minimum recommended hardware requirements for an orchestrator server are:

- 2 CPU cores
- 2 GB of memory

Source servers

Source servers can be running Windows Server 2003 or newer versions of Windows Server. This includes Windows Small Business Server and Windows Server Essentials. However, because Windows Small Business Server and Windows Server Essentials are domain controllers, data migration is supported but server name migration isn't. Windows Server 2012 or newer failover clusters are also supported sources.

Linux servers configured with Samba are also supported sources. Tested Samba versions include 3.6, 4.2, 4.3, 4.7, and 4.8 with multiple Linux distributions.

Destination servers

Destination servers can be running Windows Server 2012 R2 or newer. However, migration performance is approximately doubled when using Windows Server 2019 or Windows Server Semi-Annual Channel with the Storage Migration Service proxy installed. If you're migrating to an Azure virtual machine (VM), Storage Migration service can create the VM automatically based on specifications that you provide.

The minimum recommended hardware requirements for a destination server are:

- 2 CPU cores
- 2 GB of memory

Security

To migrate data, the necessary firewall rules must be enabled. These firewall rules might be enabled already, but you should verify this. On the orchestrator server, you must enable the File and Printer

Sharing (SMB-In) firewall rule. On source and destination servers, the following firewall rules must be enabled:

- File and Printer Sharing (SMB-In)
- Netlogon Service (NP-In)
- Windows Management Instrumentation (DCOM-In)
- Windows Management Instrumentation (WMI-In)

You can perform migrations by using a single account that is an administrator on the source, destination, and orchestrator servers. Alternatively, you can have a source migration account and a destination migration account. The source migration account is an administrator on the source and orchestrator servers. The destination migration account is an administrator on the destination and orchestrator servers.

Migrate a server with storage migration

After you have prepared all the servers for the data migration, you can begin to perform the migration. You perform all tasks by using Storage Migration Service in Windows Admin Center. As part of your preparation, you can use this tool to install the Storage Migration Service on the orchestrator server.

Note: You can include multiple source and destination servers in a migration job, but there's a one-to-one relationship between source and destination servers.

Inventory source servers

All phases of the migration are controlled by a job that you create for the purpose. The first stage in the job is inventorying source servers. You can identify whether Windows servers or Linux servers are being migrated, and you'll be prompted for credentials to gather information from the source servers.

After you've provided the basic information, add the servers to be inventoried. If you know the server names, enter them in directly. Otherwise, you can search for servers in Active Directory Domain Services (AD DS). After all the servers are added, start a scan.

The scan of the source servers identifies:

- Shares
- Server configuration
- Network adapter configuration
- Volumes

Migrate data

To transfer data, you must enter credentials that have administrative permissions on the destination server. After you enter the credentials, add the destination server and scan it to identify the volumes present on the destination server.

If you want to create an Azure virtual machine (VM) instead of specifying an existing server, the following are required:

- A valid Azure subscription.
- An existing Azure Compute resource group where you have Create rights.
- An existing Azure Virtual Network (VNet) and subnet.

- An Azure Express Route or virtual private network (VPN) solution tied to the VNet and subnet that allows connectivity from this Azure infrastructure as a service (IaaS) VM to your on-premises network and all the computers being used for migration.

To identify where the source data is to be migrated, you must map source volumes to the volumes on the destination servers. You also must identify which shares you want to migrate. In most cases, you won't want to migrate administrative shares.

You can choose to migrate local users and groups from source servers to the destination server. If there are naming conflicts with existing local users and groups, you can specify whether to reuse existing accounts or rename existing accounts on the destination. If you're migrating data from a domain controller, you must specify that local users and groups won't be transferred.

Note: Migrated local users are assigned a randomly generated 127-character password.

After you have specified the options for data transfer, you can perform a validation. The validation ensures that everything is properly configured for the migration.

Storage Migration Service is designed to allow the transfer process to be performed multiple times. The first time transfers a full copy of the data, and subsequent transfers copy only changed files. You can use this functionality to perform an initial large copy and then perform a final transfer during a maintenance window.

Note: If files are found on the destination server during the initial file copy, they must be moved to a backup folder.

After the file transfer is complete, the files, shares, and their security permissions are migrated to the destination server. If you don't want to migrate the source server identities to the destination server, you can mark the migration complete at this point.

Cut over to the destination server

Cutting over to the destination server moves the identity information from the source server to the destination server. This includes the server name and IP address.

When you perform the cutover, you must specify:

- Which adapter on the destination server will be configured with the source IP addresses.
- The IP address to assign to the source server.
- The name to assign to the source server.

When you perform the cutover, user access to the migrated data will be interrupted as the servers are renamed and restarted. This process is typically performed during a maintenance window. After the cutover is complete, users and apps can access the data by using the same server names, IP addresses, and share names as they used before the migration.

After a cutover, there's no automated way to roll back to the source server. You can manually reconfigure the computer names and IP addresses on the source and destination servers to roll back.

Storage migration considerations

Using Storage Migration Service to migrate data works well for most, but not all, scenarios. When using Storage Migration Service, you should consider the following:

- Locked files aren't migrated. If apps lock the files, the files won't be migrated when in use.

- You can't migrate the identity of domain controllers. If you're migrating file shares on a domain controller and want to migrate the domain controller identity, you should demote that domain controller to be a member server.
- Windows system files won't move to the PreExistingData folder on the destination server. For example, if you migrate a J:\Windows folder on the source server to C:\Windows on the destination server, Storage Migration Service won't migrate the folder. Other system files and folders, such as Program Files, Program Data, and Users are also protected.
- Server consolidation isn't supported. Storage Migration Service doesn't include logic to understand the dependencies involved in migrating shares from multiple servers to a single server. For example, there's no mechanism to support multiple source identities being applied to the target server.
- Data on New Technology File Systems (NTFS) must be migrated to NTFS on the target server. You can't migrate data from NTFS to Resilient File Systems (ReFS).
- Previous file versions aren't migrated. Many file servers have volume shadow copy enabled to allow easy restores of deleted or accidentally modified files. The previous versions retained in the volume shadow copies aren't migrated by Storage Migration Service.

In most cases, Storage Migration Service provides good performance for data transfer. However, to optimize performance you can do the following:

- Use Windows Server 2019 with the Storage Migration Service Proxy service installed as the destination. This allows file transfers to be performed directly from source to destination instead of being copied through the orchestrator server.
- If you have enough network bandwidth, processor performance, and memory, increasing the number of threads used by Storage Migration Service Proxy might increase performance. By default, eight threads are allocated. You can increase the allocated threads by creating a FileTransferThreadCount value in the HKEY_Local_Machine\Software\Microsoft\SMSProxy key and setting a value up to 128.
- Add processor cores and memory. Monitor source, destination, and orchestrator computers to identify whether processor and memory capacity are bottlenecks.
- Create multiple jobs. Within a single job, source servers are processed one after the other. If you create multiple jobs, they can be performed in parallel. This is most effective when the Storage Migration Service Proxy is used on Windows Server 2019 destination servers.
- Use high-performance networking. High-performance networking features such as Server Message Block version 3.0 (SMB3) with Remote Direct Memory Access (RDMA) and SMB3 multichannel ensure that network performance doesn't become a bottleneck.
- Use high-performance storage. The type of disk being used affects storage performance. Ensure that the disk subsystem has enough performance to avoid being a bottleneck. In some cases, the antivirus software can cause poor disk performance.

Test your knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

You're using Storage Migration Service to migrate file shares from a source server to a destination server. If you want to use a single account to connect to the source and destination servers, which permissions need to be configured for the account? Choose three.

- Member of users on source server.
- Member of administrators on source server.
- Member of users on orchestrator server.
- Member of administrators on orchestrator server.
- Member of administrators on destination server.

Question 2

What information do you need to specify during the cutover phase of a job in Storage Migration Service? Choose three.

- The name to assign to the source server.
- The name to assign to the destination server.
- The IP address to assign to the source server.
- The IP address to assign on the destination server.
- The network adapter to configure on the destination server.

Windows Server Migration Tools

Lesson overview

You can use the Windows Server Migration Tools feature to sync the configuration information and data from a source server to a destination server. You use the cmdlets included in the Windows Server Migration Tools to export the configuration information and data from the source server and import the configuration on the destination server. For supported roles and features, this is much faster than manually reconfiguring the features on the destination server. However, you do need to manually install the roles or features on the destination server before importing the configuration.

Lesson objectives

After completing this lesson, you'll be able to:

- Describe the Windows Server Migration Tools.
- Describe how to install the Windows Server Migration Tools.
- Describe how to use the Windows Server Migration Tools to migrate configuration data.

What are Windows Server Migration Tools?

Windows Server Migration Tools are a set of Windows PowerShell cmdlets that migrate configuration information and data from a source server to a destination server. This is done primarily to migrate server roles and features from a server being retired to a new server running a newer operating system.

Roles and features that you can migrate include:

- IP configuration
- Local users and groups
- DNS
- DHCP
- Routing and Remote Access

Windows Server Migration Tools can also be used to migrate file shares. However, you should use Storage Migration Service instead.

Source servers must be running Windows Server 2008 or newer. However, you can't migrate from a server core installation of Windows Server that doesn't have the Microsoft .NET Framework, such as Windows Server 2008 Server Core. Also, the language on the source and destination servers must match.

Install Windows Server Migration Tools

To perform a migration by using Windows Server Migration Tools, you must install the cmdlets on the source and destination servers. For a computer running Windows Server 2019, you must install the Windows Server Migration Tools feature. You can install this feature by using graphical tools such as Windows Admin Center or Server Manager. Alternatively, you can also install Windows Server Migration Tools by using Windows PowerShell.

```
Install-WindowsFeature Migration
```

If the source computer is running an earlier version of Windows Server, you must create a deployment folder with installation files for the source server. To create a deployment folder, run **SmigDeploy.exe**

from the %Windir%\System32\ServerMigrationTools\ folder on the destination server. When you run SmigDeploy.exe, you must specify the:

- Architecture of the source server.
- Operating system of the source server.
- Path to store the deployment folder.

The following example creates a deployment share for a 64-bit version of Windows Server 2008 R2 in C:\Deploy that's named SMT_WS08R2_amd64:

```
SmigDeploy.exe /package /architecture amd64 /os WS08R2 /path C:\Deploy
```

Note: For detailed information about SmigDeploy.exe switches, use the /? option.

On the source server, you need to register the Windows Server Migration Tools by running SmigDeploy.exe with no options from the deployment folder.

The deployment folder can't be run from a network location. You must copy the deployment folder to a local drive on the source server or use removable storage.

Use Windows Server Migration Tools

The Windows Server Migration Tools copy role and feature configuration, but don't install the necessary roles and features on the destination server. Before you migrate a role or feature, you must install that role or feature on the destination server in preparation for the migration. If the destination server is prepared with the necessary roles and features, you can export the configuration from the source server and then import the configuration on the destination server.

Most Windows PowerShell cmdlets install as modules and are automatically available at a Windows PowerShell prompt. To use the Windows Server Migration Tools cmdlets, you must load a snap-in first, as demonstrated in the following example:

```
Add-PSSnapin Microsoft.Windows.Windows.ServerManager.Migration
```

The following table lists the Windows Server Migration Tools cmdlets.

Table 1: Windows Server Migration Tools cmdlets

Cmdlet	Description
Get-SmigServerFeature	Lists the Windows features that can be migrated from either the local computer or a migration store.
Export-SmigServerSetting	Exports the settings for the specified Windows features and operating system (OS) from the local computer to a migration store.
Import-SmigServerSetting	Imports the settings for the specified Windows features and OS from a migration store to the local computer.
Send-SmigServerData	Sends shares and data from the source server to a destination server.
Receive-SmigServerData	Receives shares and data from the source server.

Note: You should use Storage Migration Service instead of Send-SmigServerData and Receive-SmigServerData.

Export settings

Before you export settings from the source server, you can run the `Get-SmigServerFeature` cmdlet to verify which feature settings can be exported. This cmdlet provides the feature names and IDs that you must specify during the export.

When you run the `Export-SmigServerSetting` cmdlet on the source server, you specify which Windows features to export. You also have the option to specify that local users, local groups, and IP configuration should be exported. The migration store you create with `Export-SmigServerSetting` is encrypted and protected with a password that you specify during the export.

Import settings

Before you import settings from a migration store, run the `Get-SmigServerFeature` cmdlet to verify which feature settings can be imported. Use this information to verify that the necessary Windows features install on the destination server.

When you run the `Import-SmigServerSetting` cmdlet on the destination server, you must provide the path to the migration store and the password to decrypt it. If you don't provide a password in the command, you're prompted for one. You also must identify which Windows features should be imported. In some cases, settings for Windows features must be migrated in a specific order. To guarantee that settings apply in the correct order, import the settings separately by running `Import-SmigServerSetting` multiple times.

You can choose to specify that local users, local groups, or IP configuration import from the migration store. When you import IP configuration, you must specify the hardware address of the IP configuration and the hardware address of the network card in the destination server.

Additional reading: For more information on obtaining detailed syntax information for the Windows Server Migration Tools cmdlets, refer to **ServerMigration**¹.

Test your knowledge

Use the following questions to check what you've learned in this lesson.

Question 1

Which of the following operating systems isn't supported in the Server Core configuration when using Windows Server Migration Tools?

- Windows Server 2019
- Windows Server 2008
- Windows Server 2012 R2
- Windows Server 2016
- Windows Server 2003

¹ <https://aka.ms/module-servermigration>

Question 2

How can you install the Windows Server Migration Tools on Windows Server 2019? Choose two.

- Run Install-WindowsFeature Migration.
- Download the Windows Server Migration Tools from the Microsoft Download website.
- Run SmigDeploy.exe.
- Use Windows Admin Center to install the Windows Server Migration Tools feature.

Module 12 lab and review

Lab: Migrating server workloads

Scenario

Contoso, Ltd. is an engineering, manufacturing, and distribution company. The organization is based in London, England and has major offices in Toronto, Canada and Sydney, Australia.

Because Contoso has been in business for many years, the existing servers include many versions of Windows Server. you're planning to migrate those services to servers running Windows Server 2019.

Objectives

After completing this lab, you'll be able to:

- Select a process to migrate server workloads.
- Plan how to migrate files with Storage Migration Service.

Estimated time: 20 minutes

Module review

Use the following to check what you've learned in this module.

Question 1

Which of the following are true about migrating to a new Active Directory Domain Services (AD DS) forest? Choose two.

- Migrating to a new AD DS forest is simpler than upgrading an existing AD DS forest.
- You must use the same domains and organizational unit (OU) structure in the new forest.
- You must use new domain names.
- All objects must be migrated from the source to the destination.
- During a phased migration, you can maintain access to resources in the source and destination AD DS forests.

Question 2

Which software is used to sync user passwords from a source domain to a target domain?

- ADMT
- Password Export Server
- DirSync
- BizTalk
- Users in the target domain are always assigned a new password.

Question 3

During an AD DS upgrade, which commands must you run before introducing the first domain controller running Windows Server 2019? Choose two.

- Adprep /domainprep
- Set-ADDomainMode -DomainMode Windows2016Domain
- Adprep /updateschema
- Adprep /forestprep
- Dfrsmig /getmigrationstate

Question 4

Which of the following are true of Storage Migration Service? Choose three.

- It can migrate the name of a source server to a destination server.
- It can migrate the IP address of a source server to a destination server.
- It can combine multiple source servers onto a single destination server.
- It can migrate domain users to a new target domain.
- It can migrate security permissions for files and folders to a destination server.

Question 5

Which of the following server operating systems can't have their identity migrated by Storage Migration Service? Choose two.

- Windows Server 2003
- Windows Small Business Server
- Linux servers configured with Samba
- Windows Server 2012 failover cluster
- Windows Server Essentials

Question 6

How does Storage Migration Service configure passwords for migrated users?

- Migrated users are assigned a randomly generated 127-character password.
- Migrated users are disabled and assigned a blank password.
- Migrated users are assigned a 12-character password that's stored in a text file on the destination server.
- Migrated users are assigned a password that you specify.
- Migrated users are assigned the same password that they had on the source server.

Question 7

Which two cmdlets from the Windows Server Migration Tools would you use to migrate Dynamic Host Configuration Protocol (DHCP) settings from a source server to a destination server? Choose two.

- Receive-SmigServerFeature
- Export-SmigServerFeature
- Get-SmigServerFeature
- Import-SmigServerFeature
- Send-SmigServerFeature

Answers

Question 1

Which of the following are true about upgrading Active Directory Domain Services (AD DS)? Choose two.

- The highest available forest functional level is Windows Server 2012 R2.
- All domain controllers must be running Windows Server 2019 or Windows Server 2016 before you can raise the forest functional level to Windows Server 2016.
- Upgrading AD DS to use domain controllers running Windows Server 2019 can be done without any downtime.
- Upgrading AD DS is very complex and requires downtime.

Explanation

The highest available forest and domain functional levels are Windows Server 2016. To use the Windows Server 2016 domain and forest functional levels, all domain controllers must be running Windows Server 2019 or Windows Server 2016. When you upgrade to using domain controllers running Windows Server 2019, no downtime is required because clients automatically begin using the new domain controllers when they're available.

Question 2

What is the minimum forest functional level required to support a domain controller running Windows Server 2019?

- Windows Server 2003 R2
- Windows Server 2008
- Windows Server 2008 R2
- Windows Server 2012
- Windows Server 2016

Explanation

The minimum forest and domain functional levels required to support a domain controller are Windows Server 2008.

Question 1

You're using Storage Migration Service to migrate file shares from a source server to a destination server. If you want to use a single account to connect to the source and destination servers, which permissions need to be configured for the account? Choose three.

- Member of users on source server.
- Member of administrators on source server.
- Member of users on orchestrator server.
- Member of administrators on orchestrator server.
- Member of administrators on destination server.

Explanation

A source migration account must be an administrator on the source server and the orchestrator server. A destination migration account must be an administrator on the destination server and the orchestrator server. To use a single account for migration, the account must be an administrator on the source server, destination server, and the orchestrator server.

Question 2

What information do you need to specify during the cutover phase of a job in Storage Migration Service? Choose three.

- The name to assign to the source server.
- The name to assign to the destination server.
- The IP address to assign to the source server.
- The IP address to assign on the destination server.
- The network adapter to configure on the destination server.

Explanation

During cutover, you must provide the name and IP address to assign to the source server. This is required because the original name and IP address of the source server are being assigned to the destination server. On the destination server, you must specify which network adapter will be configured with the IP address from the source server.

Question 1

Which of the following operating systems isn't supported in the Server Core configuration when using Windows Server Migration Tools?

- Windows Server 2019
- Windows Server 2008
- Windows Server 2012 R2
- Windows Server 2016
- Windows Server 2003

Explanation

To run the Windows PowerShell cmdlets provided by the Windows Server Migration Tools, the operating system must include the .NET Framework. Windows Server 2008 Core doesn't support running the .NET Framework or Windows PowerShell. Newer versions of Windows Server running as Server Core do support the .NET Framework and Windows PowerShell. Windows Server 2003 wasn't available in Server Core configuration.

Question 2

How can you install the Windows Server Migration Tools on Windows Server 2019? Choose two.

- Run Install-WindowsFeature Migration.
- Download the Windows Server Migration Tools from the Microsoft Download website.
- Run SmigDeploy.exe.
- Use Windows Admin Center to install the Windows Server Migration Tools feature.

Explanation

You can use the Install-WindowsFeature cmdlet to install the Windows Server Migration Tools. The name of the Windows Server Migration Tools feature is Migration when using the Install-WindowsFeature cmdlet. When you install the Windows Server Migration Tools feature by using Server Manager or Windows Admin Center, it uses the full name Windows Server Migration Tools in the user interface. The SmigDeploy.exe tool is used to create a deployment folder and can't be used to install the Windows Server Migration Tools feature. The Add-WindowsFeature cmdlet doesn't exist.

Question 1

Which of the following are true about migrating to a new Active Directory Domain Services (AD DS) forest? Choose two.

- Migrating to a new AD DS forest is simpler than upgrading an existing AD DS forest.
- You must use the same domains and organizational unit (OU) structure in the new forest.
- You must use new domain names.
- All objects must be migrated from the source to the destination.
- During a phased migration, you can maintain access to resources in the source and destination AD DS forests.

Explanation

The process of migrating to a new AD DS forest is complex and much harder than upgrading an existing AD DS forest. When you migrate to a new AD DS forest, you must use new domain names to avoid naming conflicts. The new AD DS forest doesn't need to mimic the OU structure of the existing domains, and you can combine multiple domains into a single domain. You can choose the specific objects that you want to migrate. Most migrations will be phased; during a phased migration you can maintain access to resources in both AD DS forests by using a forest trust and populating the `sidHistory` attribute.

Question 2

Which software is used to sync user passwords from a source domain to a target domain?

- ADMT
- Password Export Server
- DirSync
- BizTalk
- Users in the target domain are always assigned a new password.

Explanation

ADMT doesn't include native functionality to sync passwords from a source domain to a target domain. You need to install Password Export Server on a domain controller in the source domain.

Question 3

During an AD DS upgrade, which commands must you run before introducing the first domain controller running Windows Server 2019? Choose two.

- `Adprep /domainprep`
- `Set-ADDomainMode -DomainMode Windows2016Domain`
- `Adprep /updateschema`
- `Adprep /forestprep`
- `Dfrsmig /getmigrationstate`

Explanation

Before you introduce the first domain controller running Windows Server 2019 to an AD DS forest or domain, you need to prepare the forest and domain by using `Adprep`. This is required even if the forest and domain are already at the Windows Server 2016 functional level. The command `Adprep /forestprep` prepares the forest and includes schema extensions. The command `adprep /domainprep` prepares the domain.

Question 4

Which of the following are true of Storage Migration Service? Choose three.

- It can migrate the name of a source server to a destination server.
- It can migrate the IP address of a source server to a destination server.
- It can combine multiple source servers onto a single destination server.
- It can migrate domain users to a new target domain.
- It can migrate security permissions for files and folders to a destination server.

Explanation

Storage Migration Service can migrate file shares and their contents from a source server to a destination server. The data migration includes security permissions.

Storage Migration Service can migrate the identity from a source server to a destination server. The identity includes the server name and the IP addresses of the source server. However, it can't migrate the identities of multiple source servers to a single destination server, which means you can't combine multiple source servers onto a single destination server.

Storage Migration Service can migrate local users and groups from a source server to a destination server. The security permissions for domain users are preserved, but domain users aren't migrated.

Question 5

Which of the following server operating systems can't have their identity migrated by Storage Migration Service? Choose two.

- Windows Server 2003
- Windows Small Business Server
- Linux servers configured with Samba
- Windows Server 2012 failover cluster
- Windows Server Essentials

Explanation

Both Windows Small Business Server and Windows Server Essentials are configured as domain controllers.

Migrating the identity of domain controllers isn't possible. You can migrate the identity of Windows member servers and Linux servers configured with Samba.

Question 6

How does Storage Migration Service configure passwords for migrated users?

- Migrated users are assigned a randomly generated 127-character password.
- Migrated users are disabled and assigned a blank password.
- Migrated users are assigned a 12-character password that's stored in a text file on the destination server.
- Migrated users are assigned a password that you specify.
- Migrated users are assigned the same password that they had on the source server.

Explanation

When local users are migrated by Storage Migration Service, the new local user account is assigned a randomly generated 127-character password that isn't recorded anywhere. To begin using the local user account, you must assign it a known password.

Question 7

Which two cmdlets from the Windows Server Migration Tools would you use to migrate Dynamic Host Configuration Protocol (DHCP) settings from a source server to a destination server? Choose two.

- Receive-SmigServerFeature
- Export-SmigServerFeature
- Get-SmigServerFeature
- Import-SmigServerFeature
- Send-SmigServerFeature

Explanation

You should use `Export-SmigServerFeature` to save configuration information on the source server. Then you should use `Import-SmigServerFeature` to import configuration information on the destination server. These are the two cmdlets that are required.

The `Get-SmigServerFeature` cmdlet lets you notice which features can export from the source server. This can be useful, but it isn't required.

The `Send-SmigServerFeature` and `Receive-SmigServerFeature` cmdlets are used to migrate file shares. You shouldn't use these cmdlets with Windows Server 2019. Instead, you should use Storage Migration Service.