Microsoft

OFFICIAL MICROSOFT LEARNING PRODUCT

20342B

Advanced Solutions of Microsoft Exchange Server 2013

MCT USE ONLY. STUDENT USE PROHIBITED

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

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

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

© 2014 Microsoft Corporation. All rights reserved.

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

Product Number: 20342B

Part Number (if applicable): X18-52919

Released: 07/2014

MCT USE ONLY. STUDENT USE PROHIBITED

MICROSOFT LICENSE TERMS
MICROSOFT INSTRUCTOR-LED COURSEWARE

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

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

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

1. DEFINITIONS.

a. "Authorized Learning Center" means a Microsoft IT Academy Program Member, Microsoft Learning Competency Member, or such other entity as Microsoft may designate from time to time.

b. "Authorized Training Session" means the instructor-led training class using Microsoft Instructor-Led Courseware conducted by a Trainer at or through an Authorized Learning Center.

c. "Classroom Device" means one (1) dedicated, secure computer that an Authorized Learning Center owns or controls that is located at an Authorized Learning Center's training facilities that meets or exceeds the hardware level specified for the particular Microsoft Instructor-Led Courseware.

d. "End User" means an individual who is (i) duly enrolled in and attending an Authorized Training Session or Private Training Session, (ii) an employee of a MPN Member, or (iii) a Microsoft full-time employee.

e. "Licensed Content" means the content accompanying this agreement which may include the Microsoft Instructor-Led Courseware or Trainer Content.

f. "Microsoft Certified Trainer" or "MCT" means an individual who is (i) engaged to teach a training session to End Users on behalf of an Authorized Learning Center or MPN Member, and (ii) currently certified as a Microsoft Certified Trainer under the Microsoft Certification Program.

g. "Microsoft Instructor-Led Courseware" means the Microsoft-branded instructor-led training course that educates IT professionals and developers on Microsoft technologies. A Microsoft Instructor-Led Courseware title may be branded as MOC, Microsoft Dynamics or Microsoft Business Group courseware.

h. "Microsoft IT Academy Program Member" means an active member of the Microsoft IT Academy Program.

i. "Microsoft Learning Competency Member" means an active member of the Microsoft Partner Network program in good standing that currently holds the Learning Competency status.

j. "MOC" means the "Official Microsoft Learning Product" instructor-led courseware known as Microsoft Official Course that educates IT professionals and developers on Microsoft technologies.

k. "MPN Member" means an active Microsoft Partner Network program member in good standing.

MCT USE ONLY. STUDENT USE PROHIBITED

l. "Personal Device" means one (1) personal computer, device, workstation or other digital electronic device that you personally own or control that meets or exceeds the hardware level specified for the particular Microsoft Instructor-Led Courseware.

m. "Private Training Session" means the instructor-led training classes provided by MPN Members for corporate customers to teach a predefined learning objective using Microsoft Instructor-Led Courseware. These classes are not advertised or promoted to the general public and class attendance is restricted to individuals employed by or contracted by the corporate customer.

n. "Trainer" means (i) an academically accredited educator engaged by a Microsoft IT Academy Program Member to teach an Authorized Training Session, and/or (ii) a MCT.

o. "Trainer Content" means the trainer version of the Microsoft Instructor-Led Courseware and additional supplemental content designated solely for Trainers' use to teach a training session using the Microsoft Instructor-Led Courseware. Trainer Content may include Microsoft PowerPoint presentations, trainer preparation guide, train the trainer materials, Microsoft One Note packs, classroom setup guide and Pre-release course feedback form. To clarify, Trainer Content does not include any software, virtual hard disks or virtual machines.

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

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

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

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

MCT USE ONLY. STUDENT USE PROHIBITED

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

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

    provided you comply with the following:
iii. you will only provide access to the Licensed Content to those individuals who have acquired a valid license to the Licensed Content,
iv. you will ensure that each End User attending an Authorized Training Session has their own valid licensed copy of the Microsoft Instructor-Led Courseware that is the subject of the Authorized Training Session,
v. you will ensure that each End User provided with a hard-copy version of the Microsoft Instructor-Led Courseware will be presented with a copy of this agreement and each End User will agree that their use of the Microsoft Instructor-Led Courseware will be subject to the terms in this agreement prior to providing them with the Microsoft Instructor-Led Courseware. Each individual will be required to denote their acceptance of this agreement in a manner that is enforceable under local law prior to their accessing the Microsoft Instructor-Led Courseware,
vi. you will ensure that each Trainer teaching an Authorized Training Session has their own valid licensed copy of the Trainer Content that is the subject of the Authorized Training Session,
vii. you will only use qualified Trainers who hold the applicable Microsoft Certification credential that is the subject of the Microsoft Instructor-Led Courseware being taught for your Authorized Training Sessions,
viii. you will only use qualified MCTs who also hold the applicable Microsoft Certification credential that is the subject of the MOC title being taught for all your Authorized Training Sessions using MOC,
ix. you will only provide access to the Microsoft Instructor-Led Courseware to End Users, and
x. you will only provide access to the Trainer Content to Trainers.

MCT USE ONLY. STUDENT USE PROHIBITED

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

i. Each license acquired on behalf of yourself may only be used to review one (1) copy of the Microsoft Instructor-Led Courseware in the form provided to you. If the Microsoft Instructor-Led Courseware is in digital format, you may install one (1) copy on up to three (3) Personal Devices. You may not install the Microsoft Instructor-Led Courseware on a device you do not own or control.

ii. For each license you acquire on behalf of an End User or Trainer, you may either:

1. distribute one (1) hard copy version of the Microsoft Instructor-Led Courseware to one (1) End User attending the Private Training Session, and only immediately prior to the commencement of the Private Training Session that is the subject matter of the Microsoft Instructor-Led Courseware being provided, **or**
2. provide one (1) End User who is attending the Private Training Session with the unique redemption code and instructions on how they can access one (1) digital version of the Microsoft Instructor-Led Courseware, **or**
3. you will provide one (1) Trainer who is teaching the Private Training Session with the unique redemption code and instructions on how they can access one (1) Trainer Content,

provided you comply with the following:

iii. you will only provide access to the Licensed Content to those individuals who have acquired a valid license to the Licensed Content,

iv. you will ensure that each End User attending an Private Training Session has their own valid licensed copy of the Microsoft Instructor-Led Courseware that is the subject of the Private Training Session,

v. you will ensure that each End User provided with a hard copy version of the Microsoft Instructor-Led Courseware will be presented with a copy of this agreement and each End User will agree that their use of the Microsoft Instructor-Led Courseware will be subject to the terms in this agreement prior to providing them with the Microsoft Instructor-Led Courseware. Each individual will be required to denote their acceptance of this agreement in a manner that is enforceable under local law prior to their accessing the Microsoft Instructor-Led Courseware,

vi. you will ensure that each Trainer teaching an Private Training Session has their own valid licensed copy of the Trainer Content that is the subject of the Private Training Session,

vii. you will only use qualified Trainers who hold the applicable Microsoft Certification credential that is the subject of the Microsoft Instructor-Led Courseware being taught for all your Private Training Sessions,

viii. you will only use qualified MCTs who hold the applicable Microsoft Certification credential that is the subject of the MOC title being taught for all your Private Training Sessions using MOC,

ix. you will only provide access to the Microsoft Instructor-Led Courseware to End Users, and

x. you will only provide access to the Trainer Content to Trainers.

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

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

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

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

MCT USE ONLY. STUDENT USE PROHIBITED

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

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

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

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

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

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

a. **Pre-Release Licensed Content.** This Licensed Content subject matter is on the Pre-release version of the Microsoft technology. The technology may not work the way a final version of the technology will and we may change the technology for the final version. We also may not release a final version. Licensed Content based on the final version of the technology may not contain the same information as the Licensed Content based on the Pre-release version. Microsoft is under no obligation to provide you with any further content, including any Licensed Content based on the final version of the technology.

b. **Feedback.** If you agree to give feedback about the Licensed Content to Microsoft, either directly or through its third party designee, you give to Microsoft without charge, the right to use, share and commercialize your feedback in any way and for any purpose. You also give to third parties, without charge, any patent rights needed for their products, technologies and services to use or interface with any specific parts of a Microsoft technology, Microsoft product, or service that includes the feedback. You will not give feedback that is subject to a license that requires Microsoft to license its technology, technologies, or products to third parties because we include your feedback in them. These rights survive this agreement.

c. **Pre-release Term**. If you are an Microsoft IT Academy Program Member, Microsoft Learning Competency Member, MPN Member or Trainer, you will cease using all copies of the Licensed Content on the Pre-release technology upon (i) the date which Microsoft informs you is the end date for using the Licensed Content on the Pre-release technology, or (ii) sixty (60) days after the commercial release of the technology that is the subject of the Licensed Content, whichever is earliest ("**Pre-release term**"). Upon expiration or termination of the Pre-release term, you will irretrievably delete and destroy all copies of the Licensed Content in your possession or under your control.

MCT USE ONLY STUDENT USE PROHIBITED

4. **SCOPE OF LICENSE**. The Licensed Content is licensed, not sold. This agreement only gives you some rights to use the Licensed Content. Microsoft reserves all other rights. Unless applicable law gives you more rights despite this limitation, you may use the Licensed Content only as expressly permitted in this agreement. In doing so, you must comply with any technical limitations in the Licensed Content that only allows you to use it in certain ways. Except as expressly permitted in this agreement, you may not:
   - access or allow any individual to access the Licensed Content if they have not acquired a valid license for the Licensed Content,
   - alter, remove or obscure any copyright or other protective notices (including watermarks), branding or identifications contained in the Licensed Content,
   - modify or create a derivative work of any Licensed Content,
   - publicly display, or make the Licensed Content available for others to access or use,
   - copy, print, install, sell, publish, transmit, lend, adapt, reuse, link to or post, make available or distribute the Licensed Content to any third party,
   - work around any technical limitations in the Licensed Content, or
   - reverse engineer, decompile, remove or otherwise thwart any protections or disassemble the Licensed Content except and only to the extent that applicable law expressly permits, despite this limitation.

5. **RESERVATION OF RIGHTS AND OWNERSHIP**. Microsoft reserves all rights not expressly granted to you in this agreement. The Licensed Content is protected by copyright and other intellectual property laws and treaties. Microsoft or its suppliers own the title, copyright, and other intellectual property rights in the Licensed Content.

6. **EXPORT RESTRICTIONS**. The Licensed Content is subject to United States export laws and regulations. You must comply with all domestic and international export laws and regulations that apply to the Licensed Content. These laws include restrictions on destinations, end users and end use. For additional information, see www.microsoft.com/exporting.

7. **SUPPORT SERVICES**. Because the Licensed Content is "as is", we may not provide support services for it.

8. **TERMINATION.** Without prejudice to any other rights, Microsoft may terminate this agreement if you fail to comply with the terms and conditions of this agreement. Upon termination of this agreement for any reason, you will immediately stop all use of and delete and destroy all copies of the Licensed Content in your possession or under your control.

9. **LINKS TO THIRD PARTY SITES**. You may link to third party sites through the use of the Licensed Content. The third party sites are not under the control of Microsoft, and Microsoft is not responsible for the contents of any third party sites, any links contained in third party sites, or any changes or updates to third party sites. Microsoft is not responsible for webcasting or any other form of transmission received from any third party sites. Microsoft is providing these links to third party sites to you only as a convenience, and the inclusion of any link does not imply an endorsement by Microsoft of the third party site.

10. **ENTIRE AGREEMENT.** This agreement, and any additional terms for the Trainer Content, updates and supplements are the entire agreement for the Licensed Content, updates and supplements.

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

MCT USE ONLY. STUDENT USE PROHIBITED

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

12. LEGAL EFFECT. This agreement describes certain legal rights. You may have other rights under the laws of your country. You may also have rights with respect to the party from whom you acquired the Licensed Content. This agreement does not change your rights under the laws of your country if the laws of your country do not permit it to do so.

13. DISCLAIMER OF WARRANTY. THE LICENSED CONTENT IS LICENSED "AS-IS" AND "AS AVAILABLE." YOU BEAR THE RISK OF USING IT. MICROSOFT AND ITS RESPECTIVE AFFILIATES GIVES NO EXPRESS WARRANTIES, GUARANTEES, OR CONDITIONS. YOU MAY HAVE ADDITIONAL CONSUMER RIGHTS UNDER YOUR LOCAL LAWS WHICH THIS AGREEMENT CANNOT CHANGE. TO THE EXTENT PERMITTED UNDER YOUR LOCAL LAWS, MICROSOFT AND ITS RESPECTIVE AFFILIATES EXCLUDES ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT.

14. LIMITATION ON AND EXCLUSION OF REMEDIES AND DAMAGES. YOU CAN RECOVER FROM MICROSOFT, ITS RESPECTIVE AFFILIATES AND ITS SUPPLIERS ONLY DIRECT DAMAGES UP TO US$5.00. YOU CANNOT RECOVER ANY OTHER DAMAGES, INCLUDING CONSEQUENTIAL, LOST PROFITS, SPECIAL, INDIRECT OR INCIDENTAL DAMAGES.

This limitation applies to

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

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

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

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

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

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

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

MCT USE ONLY STUDENT USE PROHIBITED

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

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

Revised July 2013

MCT USE ONLY. STUDENT USE PROHIBITED

Welcome!

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

- **Microsoft Certified Trainers and Instructors**—Your instructor is a technical and instructional expert who meets ongoing certification requirements. And, if instructors are delivering training at one of our Certified Partners for Learning Solutions, they are also evaluated throughout the year by students and by Microsoft.

- **Certification Exam Benefits**—After training, consider taking a Microsoft Certification exam. Microsoft Certifications validate your skills on Microsoft technologies and can help differentiate you when finding a job or boosting your career. In fact, independent research by IDC concluded that 75% of managers believe certifications are important to team performance[1]. Ask your instructor about Microsoft Certification exam promotions and discounts that may be available to you.

- **Customer Satisfaction Guarantee**—Our Certified Partners for Learning Solutions offer a satisfaction guarantee and we hold them accountable for it. At the end of class, please complete an evaluation of today's experience. We value your feedback!

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

Sincerely,

Microsoft Learning
www.microsoft.com/learning

Microsoft | Learning

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

MCT USE ONLY. STUDENT USE PROHIBITED

Acknowledgements

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

Damir Dizdarevic – Content Developer

Damir Dizdarevic is an MCT, Microsoft Certified Solutions Expert (MCSE), Microsoft Certified Technology Specialist (MCTS), and a Microsoft Certified Information Technology Professional (MCITP). He is a manager and trainer of the Learning Center at Logosoft d.o.o., in Sarajevo, Bosnia and Herzegovina. Damir has more than 17 years of experience on Microsoft platforms and he specializes in Microsoft Windows Server®, Exchange Server, security, and virtualization. He has worked as a subject matter expert and author on many Microsoft Official Courses (MOC) courses, mostly on Exchange and Windows Server topics, and has published more than 400 articles in various IT magazines, such as *Windows ITPro*. He's also a frequent and highly rated speaker on most of Microsoft conferences in South and Eastern Europe. Additionally, he is a Microsoft Most Valuable Professional and a president of MSCommunity user group in Bosnia. His blog about MS technologies can be found at: http://dizdarevic.ba/ddamirblog.

Robert Genes – Content Developer

Robert Genes is a messaging architect and a Microsoft Certified Master for Microsoft Exchange Server 2010. As the manager of Genes Messaging Solutions, he has worked in different Exchange Server projects in Southern Germany. Robert specializes in Exchange Server and has more than 10 years of Exchange experience.

Siegfried Jagott – Content Developer

Siegfried Jagott is a Principal Consultant and Team Lead for the Messaging and Collaboration team at Atos Germany. He is an award-winning author of *Microsoft Exchange Server 2010 Best Practices* (Microsoft Press), and has authored and technically reviewed several Microsoft Official Curriculum (MOC) courses on various topics such as MOC 10165: Updating Your Skills from Microsoft Exchange Server 2003 or Exchange Server 2007 to Exchange Server 2010 SP1. He has coauthored various books on Windows, Microsoft System Center Virtual Machine Manager, and Exchange, and is a frequent presenter on these topics at international conferences such as IT & Dev Connections Spring 2012 in Las Vegas. Siegfried has planned, designed, and implemented some of the world's largest Windows® and Exchange Server infrastructures for international customers. He received an MBA from Open University in England, and has been an MCSE since 1997.

Stan Reimer – Content Developer

Stan Reimer is president of S. R. Technical Services Inc., and he works as a consultant, trainer, and author. Stan has extensive experience consulting on Active Directory® and Exchange Server deployments for some of the largest companies in Canada. Stan is the lead author for two Active Directory books for Microsoft Press®. For the last ten years, Stan has been writing courseware for Microsoft Learning, specializing in Active Directory and Exchange Server courses. Stan has been a Microsoft® Certified Trainer (MCT) for 14 years.

Thomas Poett – Content Developer

Thomas Poett is a Managing Consultant responsible for planning, designing, and troubleshooting Enterprise Unified Communication solutions. He is a Microsoft Most Valuable Professional (MVP) for Microsoft Lync®. His experience in working with Microsoft Unified Communication technologies enabled him to design the world's largest Enterprise Voice project based on Microsoft Exchange 2010 and Lync 2010. He has experience working with all versions of Microsoft Exchange Server since version 5.0 and on

MCT USE ONLY. STUDENT USE PROHIBITED

implementing Unified Communication technologies with Exchange 2000 Conferencing Server. Besides Enterprise Voice, he also has a strong interest in Private Cloud designs based on Microsoft technologies.

Anthony Steven – Technical Reviewer

Anthony is a Principal Technologist with Content Master and is based in Cumbria in the United Kingdom. He has created courseware, deployed and maintained messaging servers, delivered training courses, and provided strategic IT consultancy on Exchange Server 4.0 to Exchange Server 2013.

Martin Coetzer – Subject Matter Expert

Martin Coetzer is a Portfolio Architect with the Microsoft Learning eXperiences team. He is responsible for managing the Office 365, Exchange, Lync, SharePoint, Office and Dynamics certification portfolios. Prior to this Martin was a consultant responsible for architecting and deploying Microsoft technologies at medium to large customers around the world.

MCT USE ONLY. STUDENT USE PROHIBITED

Contents

Module 1: Microsoft Exchange Server 2013 Site Resilience

Module Overview **01-1**

Lesson 1: Site Resilience and High Availability in Exchange Server 2013 **01-2**

Lesson 2: Planning a Site Resilient Implementation **01-12**

Lesson 3: Implementing Site Resilience **01-24**

Lab: Designing and Implementing Site Resiliency **01-27**

Module Review and Takeaways **01-33**

Module 2: Planning Virtualization for Microsoft Exchange Server 2013

Module Overview **2-1**

Lesson 1: Planning a Hyper-V Deployment to Exchange Server 2013 **2-2**

Lesson 2: Virtualizing Exchange Server 2013 Server Roles **2-12**

Lab: Planning the Virtualization of Exchange Server Roles **2-19**

Module Review and Takeaways **2-27**

Module 3: Overview of Exchange Server 2013 Unified Messaging

Module Overview **3-1**

Lesson 1: Overview of Telephony Technologies **3-2**

Lesson 2: Unified Messaging in Exchange Server 2013 **3-9**

Lesson 3: Unified Messaging Components **3-14**

Lab: Unified Messaging Overview **3-21**

Module Review and Takeaways **3-23**

Module 4: Designing and Implementing Exchange Server 2013 Unified Messaging

Module Overview **4-1**

Lesson 1: Designing a Unified Messaging Deployment **4-2**

Lesson 2: Deploying and Configuring Unified Messaging Components **4-13**

Lesson 3: Designing and Implementing Exchange Server 2013 UM Integration with Lync Server 2013 **4-22**

Lab: Designing and Implementing Exchange Server 2013 Unified Messaging **4-29**

Module Review and Takeaways **4-39**

MCT USE ONLY. STUDENT USE PROHIBITED

Module 5: Designing and Implementing Message Transport Security

Module Overview **5-1**

Lesson 1: Overview of Messaging Policy and Compliance Requirements **5-2**

Lesson 2: Designing and Implementing Transport Compliance **5-5**

Lesson 3: Designing and Implementing AD RMS Integration with Exchange Server 2013 **5-15**

Lab: Designing and Implementing Message Transport Security **5-25**

Module Review and Takeaways **5-31**

Module 6: Designing and Implementing Message Retention

Module Overview **06-1**

Lesson 1: Overview of Messaging Records Management and Archiving **06-2**

Lesson 2: Designing In-Place Archiving **06-6**

Lesson 3: Designing and Implementing Message Retention **06-11**

Lab: Designing and Implementing Message Retention **06-18**

Module Review and Takeaways **06-24**

Module 7: Designing and Implementing Messaging Compliance

Module Overview **07-1**

Lesson 1: Designing and Implementing Data Loss Prevention **07-2**

Lesson 2: Designing and Implementing In-Place Hold **07-10**

Lesson 3: Designing and Implementing In-Place eDiscovery **07-14**

Lab: Designing and Implementing Messaging Compliance **07-19**

Module Review and Takeaways **07-25**

Module 8: Designing and Implementing Administrative Security and Auditing

Module Overview **08-1**

Lesson 1: Designing and Implementing RBAC **08-2**

Lesson 2: Designing and Implementing Split Permissions **08-14**

Lesson 3: Planning and Implementing Audit Logging **08-19**

Lab: Designing and Implementing Administrative Security and Auditing **08-25**

Module Review and Takeaways **08-31**

MCT USE ONLY. STUDENT USE PROHIBITED

Module 9: Managing Exchange Server 2013 with Exchange Management Shell

Module Overview **9-1**

Lesson 1: Overview of Windows PowerShell 3.0 **9-2**

Lesson 2: Managing Exchange Server Recipients by Using the Exchange Management Shell **9-7**

Lesson 3: Using Windows PowerShell to Manage Exchange Server **9-17**

Lab: Managing Microsoft Exchange Server 2013 by Using Exchange Management Shell **9-23**

Module Review and Takeaways **9-27**

Module 10: Designing and Implementing Integration with Microsoft Exchange Online

Module Overview **10-1**

Lesson 1: Planning for Exchange Online **10-2**

Lesson 2: Planning and Implementing the Migration to Exchange Online **10-10**

Lesson 3: Planning to Coexist with Exchange Online **10-15**

Lab: Designing Integration with Exchange Online **10-25**

Module Review and Takeaways **10-27**

Module 11: Designing and Implementing Messaging Coexistence

Module Overview **11-1**

Lesson 1: Designing and Implementing Federation **11-2**

Lesson 2: Designing Coexistence Between Exchange Server Organizations **11-8**

Lesson 3: Designing and Implementing Cross-Forest Mailbox Moves **11-15**

Lab: Implementing Messaging Coexistence **11-23**

Module Review and Takeaways **11-30**

Module 12: Designing and Implementing Microsoft Exchange Server Upgrades

Module Overview **12-1**

Lesson 1: Planning the Upgrade from Previous Exchange Server Versions **12-2**

Lesson 2: Implementing the Upgrade from Previous Exchange Versions **12-16**

Lab: Upgrading from Exchange Server 2010 to Exchange Server 2013 **12-23**

Module Review and Takeaways **12-35**

MCT USE ONLY. STUDENT USE PROHIBITED

Lab Answer Keys

Module 1 Lab: Designing and Implementing Site Resiliency **L01-1**

Module 2 Lab: Planning the Virtualization of Exchange Server Roles **L02-9**

Module 3 Lab: Unified Messaging Overview **L03-17**

Module 4 Lab: Designing and Implementing Exchange Server 2013 Unified Messaging **L04-19**

Module 5 Lab: Designing and Implementing Message Transport Security **L05-31**

Module 6 Lab: Designing and Implementing Message Retention **L06-39**

Module 7 Lab: Designing and Implementing Messaging Compliance **L07-45**

Module 8 Lab: Designing and Implementing Administrative Security and Auditing **L08-53**

Module 9 Lab: Managing Microsoft Exchange Server 2013 by Using Exchange Management Shell **L09-61**

Module 10 Lab: Designing Integration with Exchange Online **L10-65**

Module 11 Lab: Implementing Messaging Coexistence **L11-67**

Module 12 Lab: Upgrading from Exchange Server 2010 to Exchange Server 2013 **L12-77**

MCT USE ONLY. STUDENT USE PROHIBITED

MCT USE ONLY. STUDENT USE PROHIBITED

About This Course

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

Course Description

This course will provide you with the knowledge to design and implement a Microsoft® Exchange Server 2013 messaging environment. This course will teach you how to design and configure advanced components in an Exchange Server 2013 deployment, and it will provide guidelines, best practices, and considerations that will help you optimize your Exchange Server deployment. This course is an instructor-led course and will provide students with the knowledge and skills to design, manage, and configure Unified Messaging, site resiliency, advanced security, compliance, archiving, discovery solutions, coexistence, hybrid scenarios, migration, and federation.

Audience

This course is intended for people aspiring to be enterprise-level messaging administrators. Others who may take this course include IT generalists and help desk professionals who want to learn about Exchange Server 2013. People taking the course are expected to have at least 3 years of experience working in the IT field—typically in the areas of network administration, help desk, or system administration. Students are expected to have experience with Exchange Server 2013 or previous versions of Exchange Server.

This course is also intended as preparation material for IT professionals who are looking either to take the exam 70-342A: Microsoft Exchange Server 2013, Advanced Solutions itself, or as part of the requirement for the MCSE: Microsoft Exchange Server 2013 certification.

Student Prerequisites

This course requires that you meet the following prerequisites:

- Passed 70-341: Core Solutions of Microsoft Exchange Server 2013, or equivalent
- Minimum of two years of experience working with Exchange Server
- Minimum of six months of experience working with Microsoft Exchange Server 2010 or Exchange Server 2013
- Minimum of two years of experience administering the Windows Server® operating system, including Windows Server 2008 R2 or Windows Server 2012
- Minimum of two years of experience working with Active Directory® Domain Services
- Minimum of two years of experience working with name resolution, including Domain Name System (DNS)
- Experience working with certificates, including public key infrastructure (PKI) certificates
- Experience working with Windows® PowerShell

Course Objectives

After completing this course, students will be able to:

- Design and implement site resiliency for Exchange Server 2013.
- Plan a virtualization strategy for Exchange Server 2013 roles.
- Describe the basic concepts of Unified Messaging in Exchange Server 2013.
- Design and implement Exchange Server 2013 Unified Messaging.

MCT USE ONLY. STUDENT USE PROHIBITED

- Design and implement message transport security.
- Design and implement message retention in Exchange Server 2013.
- Design and implement messaging compliance.
- Design and implement administrative security in an Exchange Server 2013 environment.
- Use Windows PowerShell 3.0 to manage Exchange Server 2013.
- Design and implement integration with Exchange Online.
- Design and implement messaging coexistence.
- Design and implement Exchange Server upgrades.

Course Outline

The course outline is as follows:

- **Module 1**, "Designing and Implementing Site Resilience" describes how to design and implement site resiliency for Exchange Server 2013.
- **Module 2**, "Planning Virtualization for Microsoft Exchange Server 2013" explains how to plan a virtualization strategy for Exchange Server 2013 roles.
- **Module 3**, "Overview of Exchange Server 2013 Unified Messaging" explains the basic concepts of Unified Messaging in Exchange Server 2013.
- **Module 4**, "Designing and Implementing Exchange Server 2013 Unified Messaging" describes how to how to design and implement Exchange Server 2013 Unified Messaging.
- **Module 5**, "Designing and Implementing Message Transport Security" explains how to design and implement message transport security.
- **Module 6**, "Designing and Implementing Message Retention" explains how to design and implement message retention in Exchange Server 2013.
- **Module 7**, "Designing and Implementing Messaging Compliance" explains how to design and implement messaging compliance.
- **Module 8**, "Designing and Implementing Administrative Security and Auditing" explains how to design and implement administrative security in an Exchange Server 2013 environment.
- **Module 9**, "Managing Exchange Server 2013 with Exchange Management Shell" explains how to use Windows PowerShell 3.0 to manage Exchange Server 2013.
- **Module 10**, "Designing and Implementing Integration with Microsoft Exchange Online" explains how to design and implement integration with Exchange Online.
- **Module 11**, "Designing and Implementing Messaging Coexistence" explains how to design and implement messaging coexistence.
- **Module 12**, "Designing and Implementing Exchange Server Upgrades" explains the options and procedures for upgrading a current Exchange Server environment to Exchange Server 2013

MCT USE ONLY. STUDENT USE PROHIBITED

Course Materials

The following materials are included with your kit:

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



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

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



Additional Reading: Student Course files on the http://www.microsoft.com/learning/en/us/companion-moc.aspx Site: includes the Allfiles.exe, a self-extracting executable file that contains all required files for the labs and demonstrations.

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

MCT USE ONLY. STUDENT USE PROHIBITED

Virtual Machine Environment

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

Virtual Machine Configuration

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

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

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

Virtual machine	Role
20342B-LON-DC1	Domain controller in the Adatum.com domain
20342B-LON-MBX1	Exchange Server 2013 Mailbox server
20342B-LON-CAS1	Exchange Server 2013 Client Access server
20342B-LON-MBX2	Exchange Server 2013 Mailbox server
20342B-LON-CAS2	Exchange Server 2013 Client Access server
20342B-LON-LY1	Microsoft ® Server 2013 Standard Edition server
20342B-LON-CL1	Windows 8 computer in the Adatum.com domain
20342B-LON-CL2	Windows 8 computer in the Adatum.com domain
20342B-Trey-DC1	Domain controller in the TreyResearch.net domain
20342B-Trey-EX1	Exchange Server 2010 server in the TreyResearch.net domain
20342B-TREY-EX13	Member server in the TreyResearch.net domain

Software Configuration

The following software is installed on each student LUC-CL1 VM:

- Windows Server 2012
- Windows® 8
- Microsoft Office 2013
- Exchange Server 2013, Cumulative Update 1
- Lync® Server 2013

MCT USE ONLY. STUDENT USE PROHIBITED

Course Files

The files associated with the labs in this course are located in the <install_folder>\Labfiles\LabXX folder on the student computers.

Classroom Setup

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

Course Hardware Level

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

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

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

MCT USE ONLY. STUDENT USE PROHIBITED

MCT USE ONLY. STUDENT USE PROHIBITED

Module 1

Microsoft Exchange Server 2013 Site Resilience

Contents:

Module Overview **01-1**

Lesson 1: Site Resilience and High Availability in Exchange Server 2013 **01-2**

Lesson 2: Planning a Site Resilient Implementation **01-12**

Lesson 3: Implementing Site Resilience **01-24**

Lab: Designing and Implementing Site Resiliency **01-27**

Module Review and Takeaways **01-33**

Module Overview

To design and deploy site resilience for a Microsoft® Exchange Server 2013 installation, you must plan every aspect of the implementation. The Mailbox and Client Access server roles have distinct ways to achieve high availability and site resilience. In addition, you must make other services resilient, such as Domain Name System (DNS), network connectivity, and Active Directory® Domain Services (AD DS). This module examines what you must consider to design and deploy a resilient site.

Objectives

After completing this module, you will be able to:

- Describe site resilience.
- Plan a site resilient Exchange Server 2013 deployment.
- Implement a site resilient Exchange Server 2013 deployment.

MCT USE ONLY. STUDENT USE PROHIBITED

Lesson 1
Site Resilience and High Availability in Exchange Server 2013

Messaging is a mission critical tool to many businesses. An organization's email system can contain customer information, legal correspondence, and other valuable information. If this information is lost or unavailable, business processes are interrupted. It is no wonder why organizations put so much effort in to keeping email data safe and available despite any number of failures. To reduce this effort, Exchange Server 2013 has several new features that make a site resilient solution easier to design, deploy, and manage. Many of these simplifications result from changes in namespace planning and in the site failover process.

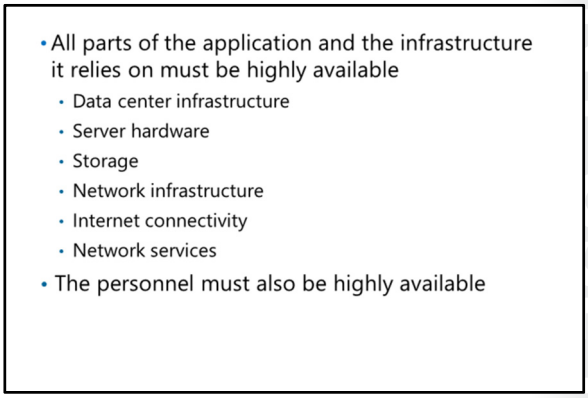
Lesson Objectives

After this lesson, you will be able to:

- Describe high availability improvements in Exchange Server 2013.
- Describe database availability groups.
- Describe site resilience.
- Describe improvements in site resilience in Exchange Server 2013.

Components of High Availability

A high availability configuration withstands failures by reducing, mitigating, or eliminating failure points. To build a highly available solution, you must examine all the components of the solution and find these potential failure points. This approach also applies to building a highly available Exchange Server 2013 deployment. Before you decide which of the high availability features of Exchange Server 2013 to use, consider the other components that affect the overall availability of your Exchange Server 2013 solution.



The following are some of the additional components that you must consider:

- Data center infrastructure. The servers must have sufficient power and cooling capacity, which must also be highly available. You can make power highly available by ensuring that an alternate power source—such as an uninterruptable power supply (UPS) and a generator—is available if the electricity goes out. You can make cooling capacity highly available by using multiple cooling units that have sufficient capacity to keep the data center cool if one unit fails. In the case of a catastrophic failure, you can use an alternate data center location that is geographically distributed.
- Server hardware. To make server hardware highly available, there must be redundant components in the server. Redundant components include power supplies, network adapters, processors, and memory. Error-correction code memory helps resolve minor errors in memory.
- Storage. To make storage highly available on a single server, you can use a version of Redundant Array of Independent Disks (RAID). RAID uses redundancy or parity information to help ensure that a server can survive the loss of at least one hard drive without losing any data. If multiple servers are

MCT USE ONLY. STUDENT USE PROHIBITED

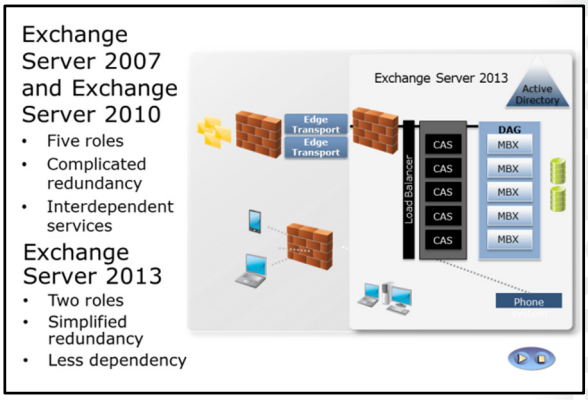
available, you can replicate data between servers. Data that is replicated between servers can survive the loss of an entire server, rather than just the loss of a hard drive. You can also use a combination of RAID and server replication to provide a highly available storage solution.

- Network infrastructure. To make a local area network (LAN) highly available, you must introduce redundant components. Within a LAN, this resilience typically requires redundant switches. Even moderately priced switches include redundant configurations. To make the network connectivity for any individual computer fault-tolerant, you must configure redundant network adapters on the computer. This is a standard feature in most mid-level and higher servers. High availability for a wide area network (WAN) is typically the responsibility of the WAN service provider. However, if you are using private links for your WAN, you can create redundant paths through it.
- Internet connectivity. For highly available Internet access, you must have redundant Internet connectivity. Ideally, use two different Internet service providers (ISPs) and two different physical connectivity methods. For example, one ISP can be land-based, and the other wireless. If you use these methods, a problem that affects one ISP is unlikely to affect the other. Many firewalls and routers can use one connection for Internet connectivity and can fail over to another connection if the primary service fails. For incoming email, you must use multiple mail exchanger (MX) resource records, with one record pointing at the IP address allocated by each ISP.
- Network services. AD DS and DNS are two services that must be highly available to support highly available Exchange Server 2013 organizations. To make AD DS highly available, you should have multiple domain controllers and global catalog servers. Depending on the size of a location, there may be multiple domain controllers and global catalog servers in a single location. To make internal DNS highly available, you must have multiple DNS servers which have DNS information synchronized between them. By default, the DNS zones for AD DS are Active Directory–integrated, and replicated between all domain controllers in the forest.
- Personnel. Failures do not happen when it is convenient for everyone to be available. You must also consider that an event that causes a primary site failure may also impact any staff located at or near that location. The people that manage, maintain, and repair the applications must also be highly available. To accomplish this, put in place comprehensive documentation and an automation plan. You can also mitigate potential issues by cross-training staff members so that multiple people share knowledge, and by hiring employees who live near the data centers.

  Question: Which infrastructure is highly available in your organization?

MCT USE ONLY. STUDENT USE PROHIBITED

High Availability in Exchange Server 2013



Microsoft Exchange Server 2007 and Microsoft Exchange Server 2010 are a part of many successful, highly available solutions. However, with five interdependent server roles, configuration and upgrades are more challenging. For example, interdependencies between roles require a load balancer to maintain session affinity for inter-role communications. This requirement significantly complicates the configuration, limits scalability, and increases the cost of a highly available Exchange deployment. These complications and added steps lead many administrators to choose to deploy Exchange Server without using load balancers, which reduces the overall effectiveness of the high availability features.

In Exchange Server 2013, the number of server roles is reduced from five roles to just two: the Client Access server role and the Mailbox server role. This change means that the separate functions now run on every server. The following table shows how the new server roles in Exchange Server 2013 compare to the server roles in Exchange Server 2010.

Exchange Server 2010 role	Exchange Server 2013 role	Functionality
Mailbox	Mailbox	All
Client Access	Client Access and Mailbox	Client Access: Authentication, proxying, and redirection Mailbox: Protocols, APIs, Microsoft Outlook® Web App rendering, business logic, Mailbox replication service, remote procedure call (RPC) over HTTP, MAPI over HTTP
Hub Transport	Mailbox	All
Unified Messaging	Client Access and Mailbox	Client Access: Unified Messaging Call Router Mailbox: Unified Messaging
Edge Transport	Edge Transport is available in Exchange Server 2013 SP1	None

This consolidation of roles shifts all core functionality to the Mailbox server role. The Mailbox server no longer must rely on other roles to send, receive, or transcode email; to render Microsoft Outlook® Web App content; or to perform other functions. In Exchange Server 2013, the local instance of these services always handles all of these functions. Load balancers do not need to track session affinity for communication from the client computer to the Client Access servers or from the Client Access servers to the Mailbox servers, because only the Mailbox server that has the active database copy performs these functions.

Client Access Availability

High availability for Client Access servers starts with deploying multiple Client Access servers. Next, you need to configure either hardware-based or software-based load balancing, such as the Network Load Balancing (NLB) feature in the Windows Server® 2012 operating system. You can also create multiple host

MCT USE ONLY. STUDENT USE PROHIBITED

(A) resource records in DNS for your Client Access servers and configure round-robin DNS. Round-robin DNS enables you to distribute network connections across the different Client Access servers, but it does not provide load balancing or automatic failover. Load balancing spreads client requests between the Client Access servers. If one Client Access server becomes unavailable, the remaining Client Access servers handle those client requests.

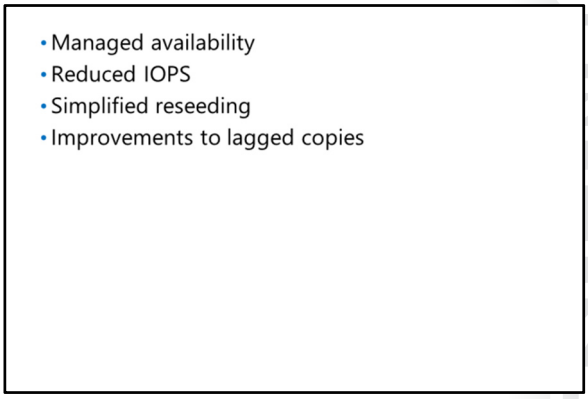
Mailbox Server Availability

You can make Mailbox servers highly available by using database availability groups (DAGs). A DAG is a collection of servers that provides the infrastructure for replicating and activating database copies, along with the supporting services.

High availability in Exchange Server 2013 is simplified in the following ways:

- Simplified Exchange role architecture. Exchange Server 2013 reduces the number of server roles to two. Each Mailbox server hosts its own services for message transport, database store, and business logic. A separate high availability configuration for each of these services is no longer available, and a Mailbox server no longer depends on other servers to carry out these activities.
- Client Access simplification. In Exchange Server 2010, the Client Access server role requires a complicated load balancing configuration, due to the number and types of communication protocols. Exchange Server 2013 eliminates this complexity.

Improvements to High Availability in Exchange Server 2013



Beyond the simplified role architecture, Exchange Server 2013 also introduces several other improvements to high availability. Managed availability monitors for and recovers from problems that involve access to messaging services. Even if a service for the Windows® operating system is running or the Event log has no errors, the users might still be experiencing problems. Rather than focusing on whether services are running, managed availability monitors the health of the user experience. This approach improves availability because, rather than just checking whether a database is mounted, managed availability tests the functionality of Outlook Web App and communication protocols and remediates problems when possible. For example, if a user cannot send an email from Outlook Web App, managed availability tries to recover the transport service. If the problem continues, the active database may be switched to another server in the DAG. Intelligent decisions maintain a healthy user experience, thus improving availability in a way that is noticeable to users.

DAG enhancements also improve the availability in Exchange Server 2013. For example, passive node copies consume roughly half the number of disk input/output operations per second (IOPS) as the primary copy does. This change enables faster reseed operations, so you can deploy multiple databases to a single disk. Another improvement is the addition of the Autoreseed feature, which reduces the complexity of the reseed process by automatically reseeding failed database copies to a spare disk. If a failure occurs, it is important that a healthy copy of the database be activated on a server that is functional. This is why the best copy selection process includes checking the health status from the managed availability service.

Each of the last several releases of Exchange Server has reduced the amount of IOPS needed to support user mailboxes. These improvements have opened the possibility of using cheaper, less reliable storage, if

MCT USE ONLY. STUDENT USE PROHIBITED

you take proper precautions to ensure quick recovery from an error. In this manner, Exchange Server Products can recover from many storage failures automatically. Automatic recovery steps can include restarting the server if storage systems become unresponsive. This action can enable the server to recover automatically from problems or to take steps to fix the problem before an administrator can continue troubleshooting and recovering.

The activation times of lagged copies decrease as well, because they are integrated with Safety Net. Regardless of which version of Exchange you have deployed, you should never activate your only lagged copy. If you do, you eliminate the protection that the lagged copy provides. If you ever plan to activate a lagged copy, be sure to have two lagged copies available. If you activate a lagged copy in Exchange Server 2010, the lagged copy must replay all lagged transaction logs before it can be mounted and become available to end users. If the lagged copy has a seven day replay lag, it can take hours to replay the outstanding transaction logs and mount the database. To use Safety Net for improved lagged copy activation, you must configure Safety Net retention to be equal to or longer than the replay lag. To activate the lagged copy using Safety Net data, discard the lagged transaction logs and mount the database. The database immediately mounts, and Safety Net redelivers any missing email.

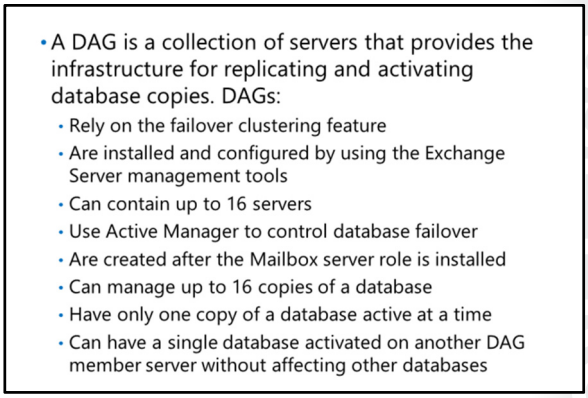
Exchange Server 2013 also includes several other enhancements to lagged copies that improve high availability by enabling automatic replay of transaction logs to the lagged copy in several critical situations, as follows:

- If Exchange Server detects a corruption issue in the lagged copy, it automatically begins to replay the log and to update the corrupt page by using the active copy.
- If a low disk space threshold is reached, logs are replayed to the lagged copy to free up space.

If there are three or fewer healthy copies of the database for 24 hours, the lagged copy is automatically replayed to make it ready for use in case another copy is lost.

What are DAGs?

A DAG is a collection of servers that provides the infrastructure for replicating and activating database copies. Changes in the active database are continuously replicated to each of the passive database copies within the DAG. DAGs have the following characteristics:



- They require the Windows Server failover clustering feature, although all installation and configuration tasks occur with the Exchange Administration Center or Exchange Management Shell. Even though a DAG requires the failover clustering feature, Exchange Server does not use it to manage database failover. Instead, it uses an internal component, Active Manager, to manage failover. Windows Server failover clustering is used to detect some failures, such as a server failure.
- DAGs use an improved version of the continuous replication technology that Exchange Server 2007 introduced. These improvements support the new high availability features, such as database copies and database mobility. Continuous replication is explained later in this lesson.
- You can use DAGs to add and remove Mailbox servers at any time. You do not need to decide on the DAG membership during installation.

MCT USE ONLY. STUDENT USE PROHIBITED

- Because DAGs use the failover clustering feature, you must install Exchange Server 2013 on Windows Server 2012 Datacenter Edition, Windows Server 2012 Standard Edition, or Windows Server 2012 R2, or on Windows Server 2008 R2 Enterprise Edition or Windows Server 2008 Datacenter Edition.
- You can move a single database between servers in the DAG without affecting other databases.
- You can add up to 16 servers to a DAG, which means that you can create up to 16 copies of a database. The database copies must be stored in the same path on all servers. For example, if you store Mailbox Database 1 in D:\Mailbox\DB\Mailbox Database 1\ on LON-MBX01, you must also store it in D:\Mailbox\DB\Mailbox Database 1\ on all other servers that host copies of Mailbox Database 1.
- DAGs define the boundary for replication, because only servers within the DAG can host database copies. You cannot replicate database information to Mailbox servers outside the DAG.

The active database copy uses continuous replication to keep the passive copies synchronized based on their replay lag-time setting. A DAG leverages the failover clustering feature in Windows Server. However, the DAG relies on the Active Manager server to maintain the status of all of the databases hosted in DAGs. The following are database characteristics:

- A single database can fail over or switch over between DAG servers. However, it is active on only one server at a time.
- At any given time, a copy of the database is either the replication source or the replication target, but not both.
- A server may not host more than one copy of a given database.
- All database copies must be stored in the same path on each server.
- You can configure database copies as lagged copies. This configuration delays the application of updates received from the active database for a configured time period of up to 14 days.
- Not all databases must have the same number of copies. In a 16-node DAG, one database can have 16 copies, while another database that is not redundant can contain one active copy.

A database failover occurs if failures cause the active database to go offline. Either a single server failure or something specific to a database may cause the failure. A switchover occurs if an administrator intentionally coordinates moving the active database from one server to another.

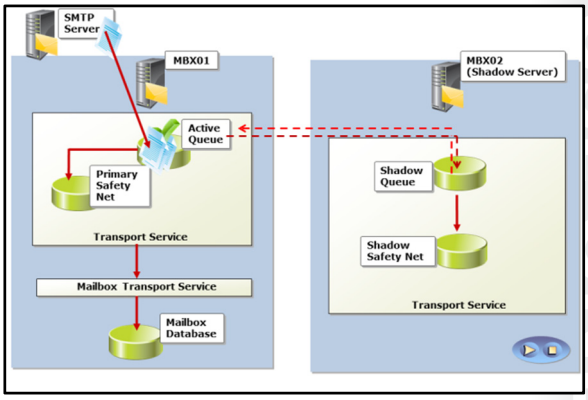
The primary Active Manager in a DAG determines which copies are active and which are passive. It is also responsible for processing topology change notifications and for reacting to server failures. A standby Active Manager provides information to other components of Exchange Server about which server hosts the active copy of a mailbox database. For example, the Client Access server communicates with the Active Manager to determine which DAG server has the active database for a specific mailbox for a user. A standby Active Manager also detects local database and local information store failures. If the database is replicated, the standby Active Manager reacts to failures by sending a request to the primary Active Manager to initiate a failover.

You must manually create the database copies on each DAG member. When creating database copies using the **Add-MailboxDatabaseCopy** cmdlet, you can also specify the following properties:

- **ActivationPreference**. A lower number means that Active Manager gives a higher preference to that copy when determining which database copy to activate.
- **ReplayLagTime**. This setting specifies how long the transaction log files are held before being replayed on the database copy. The default replay lag is zero, or disabled, and the maximum is 14 days.
- **TruncationLagTime**. This value specifies the amount of time to wait before replayed logs are truncated. The default truncation lag is zero, or disabled, and the maximum lag is 14 days.

MCT USE ONLY. STUDENT USE PROHIBITED

Understanding Transport High Availability



Shadow redundancy is a feature introduced with Exchange Server 2010 that makes a copy of a message available if a Mailbox server crashes before messages are committed to the databases. Exchange Server 2013 improves this feature by automatically creating a redundant copy of any message it receives before it acknowledges successful receipt to the sending Simple Mail Transfer Protocol (SMTP) server. In Exchange Server 2013, it no longer matters if a sending server supports shadow redundancy, because now a shadow copy is automatically created every time. By default, a shadow copy of a message is removed after two days.

The main goal of shadow redundancy is to maintain two copies of a message within a transport high availability boundary while the message is in transit. A transport high availability boundary is one of the following:

- A DAG, for Mailbox servers that are members of a DAG. This boundary includes a DAG that spans multiple Active Directory sites.
- An Active Directory site, for Mailbox servers that do not belong to a DAG.

Where and when the redundant copy of the message is created depends on where the message came from and where it is going. There are three situations where shadow redundancy protects messages:

- Messages received from outside a transport high availability boundary.
- Messages sent outside a transport high availability boundary.
- Messages received from the Mailbox Transport Submission service from a Mailbox server within the transport high availability boundary.

Note: Shadow redundancy never tracks shadow messages across a transport high availability boundary.

How Shadow Redundancy Works

The following is an example of how shadow redundancy in a DAG works:

1. An SMTP server connects to the Transport service on a Mailbox server where the active database of the target recipient is mounted, and it transmits a message. After the message is received, the session stays active.
2. The Transport service opens a new SMTP session to a Transport service on another Mailbox server in the same DAG to create a redundant copy of the message. If the DAG spans multiple Active Directory sites, a Mailbox server in another Active Directory site is preferred by default. The copy of the message is the shadow message, and the Mailbox server that holds it is the shadow server for the primary server. The message exists in a shadow queue on the shadow server.
3. After the message is successfully transmitted to the shadow server, the server acknowledges receipt of the message to the SMTP server and closes the connection.

Note: If the Mailbox server is not member of a DAG, any mailbox server in the same Active Directory site is used as a shadow server.

MCT USE ONLY. STUDENT USE PROHIBITED

When Shadow Messages are Removed

After the server successfully transmits the message to the database, the server updates the discard status of the message. The discard status is essentially a message that contains of list of messages that are being monitored. A successfully delivered message does not need to be retained in a shadow queue. After the shadow server knows that the primary server has successfully transmitted the message to the next hop, the shadow server moves the shadow message from the shadow queue into the Safety Net.

How Message Recovery Works

If a Mailbox server has an outage due to a hardware failure, each Mailbox server that has shadow messages queued for it assumes ownership of those messages. When the server comes back online, it tries to resubmit the messages. All messages are then redelivered to their destinations. This results in duplicate delivery of the messages, but Exchange Server automatically detects duplicate messages and does not add them to the database again. It only adds the messages that are not already in the database.

Safety Net

Safety Net is a special message queue that is available in the Transport service on every Mailbox server. By default, this queue stores up to two days of messages that are successfully delivered to a mailbox database. Safety Net helps protect against Mailbox server failures in which transaction logs are lost. If a failure occurs and some transaction logs are not replicated to the passive copy, you can use Safety Net to redeliver messages.

Safety Net in Exchange Server 2013 improves the transport dumpster in Exchange Server 2010 in the following ways:

- Safety Net is redundant and uses shadow redundancy to provide a shadow Safety Net queue on another server. Shadow redundancy does not keep another copy of the message, as does the transport dumpster in Exchange Server 2010. If the primary Safety Net is unavailable for more than 12 hours, the resubmit requests become shadow resubmit requests, and messages are redelivered from the shadow Safety Net.
- Safety Net does not require a DAG. It uses the same server that shadow redundancy uses to store a shadowed Safety Net copy.

How Safety Net Works

Safety Net works as follows when shadow redundancy is finished:

1. The Transport service on the primary server processes the primary message. The Mailbox Transport service delivers the message to the local mailbox database. Then the message is moved from the queue to the primary Safety Net queue.
2. The shadow server frequently polls the primary server for the discard status of the primary message. After the discard status is received, the shadow server moves the message from the shadow queue to the shadow Safety Net queue.

MCT USE ONLY STUDENT USE PROHIBITED

What is Site Resilience?

A site resilient Exchange solution is one which is deployed in two or more data centers and is designed to remain functional even if one or more sites are offline. An entire site may go offline due to events such as a power outage, a network failure, or a natural disaster. In some cases, the secondary data center is a site dedicated to disaster recovery. In other cases, the secondary data center might be another company location that has sufficient capacity to handle services in the event of a primary site failure.

- Site resilience is the ability of a messaging system to survive the failure of a site
  - An alternate data center can be located at another company location that has sufficient capacity
  - A DAG can span Active Directory sites
  - Other roles and services must already exist in the site

A successful site resilient design requires more than deploying an extra Exchange server in the failover site. In addition, it requires that a least a subset of the services from the primary site also be available in the secondary site. Services include AD DS, DNS, Exchange Server 2013 Mailbox and Client Access server roles, and network connectivity. However, these services have little value if the mailbox data is not available in the secondary site.

Although a site resilient design employs multiple high availability techniques, configuring high availability features does not make a deployment site resilient. You must design a site resilient solution to be highly available.

Question: Does your organization plan for site resilience as part of its disaster recovery planning?

Discussion: Deploying a Site Resilient Configuration

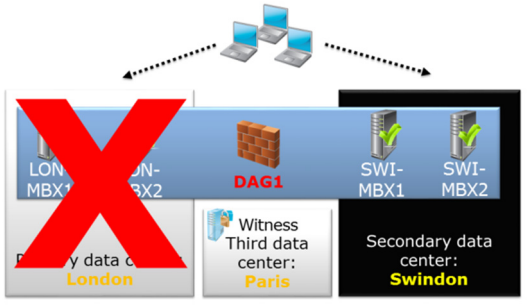
Discuss the following questions:

- Why might you implement site resilience?
- What might cause a site failure?
- What should you consider when designing a site resilient solution?

- Discuss the following questions:
  - Why might you implement site resilience?
  - What might cause a site failure?
  - What should you consider when designing a site resilient solution?

Site Resilience in Exchange Server 2013

At the simplest level, site resilience enables a user to access their email even if the primary data center is unavailable. Exchange Server site resilience is achieved by deploying and configuring Mailbox and Client Access servers in each site. Users can access their mailbox contents and send messages, because database copies are maintained in a secondary site by using a DAG and because Client Access servers are configured in the secondary site. A third site can be used to provide enough votes to maintain quorum if a



MCT USE ONLY. STUDENT USE PROHIBITED

single site fails.

If the primary data center fails, the databases are activated in the secondary data center. Client computers reconnect to the Client Access servers in the secondary site because they resolve the IP address of the secondary site from global server load balancing or simply by requesting a new valid IP address from DNS.

The new role architecture in Exchange Server 2013 enables separate recovery of the Client Access and Mailbox servers. If an Exchange Server 2010 site resilient configuration loses the Client Access Array, or if the majority of the DAG members are lost, a manual data center switchover is required. In Exchange Server 2013, these switchovers are not required in most scenarios.

For example if Client Access servers in the primary site are unavailable, the clients are redirected to the Client Access servers in the secondary datacenter. These Client Access servers proxy the communication back to the user's Mailbox server in the primary site. Instead of recovering the entire service, the clients automatically reconnect, and you can focus on fixing the Client Access servers in the primary site. This also works similarly if the Mailbox servers in the primary site become unavailable; but enough votes are still available in the DAG to maintain a quorum. The clients remain connected to the Client Access servers in the primary site, and the Client Access servers' proxy connections to the active mailbox copies in the secondary site.

MCT USE ONLY. STUDENT USE PROHIBITED

Lesson 2
Planning a Site Resilient Implementation

After gaining a basic understanding of the components of a site resilient Exchange Server deployment, you can begin to plan a deployment. The planning process requires you to gather the appropriate information and to make design decisions based on that information.

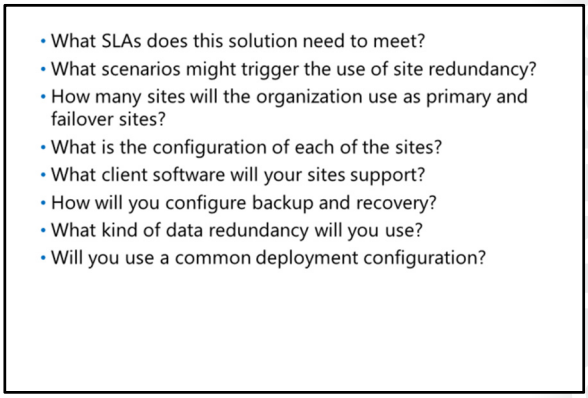
Lesson Objectives

After this lesson, you will be able to:

- Define the basic site resilient design options.
- Describe namespace considerations for a site resilient deployment.
- Plan a site resilient DAG architecture.
- Describe transport considerations for a site resilient deployment.

Planning Exchange Server 2013 Site Resiliency

The first step of planning a site resilient deployment is to define the requirements. To help define the requirements, you should ask the following questions:



- What service level agreements (SLAs) does this solution need to meet? The required level of service drives all the other design decisions. The higher the availability requirement in the SLA, the more levels of redundancy the design requires at both the primary site and the secondary, failover site.
- What scenarios might trigger the use of site redundancy? For example, will site redundancy support extended site failures? If so, the secondary site most likely needs to have capacity and services similar to the primary site.
- How many sites will the organization use as primary and failover sites? The number and location of the sites define where you locate the servers and how you assign their roles.
- What is the configuration of each of the sites? Define the number and type of mailboxes that will be active at each of these sites during both normal and failover situations. Also, define the length of time and the services that will be offered. For example, you may decide not to offer unified messaging in a failover site because of the added expense it requires. You may also decide to have only two copies of each mailbox database available in the secondary site, even though the primary site hosts three copies of the database.
- What client software will your sites support? Define which client software the sites need to support, such as the Microsoft Office Outlook messaging client, Outlook Web App, Internet Message Access Protocol 4 (IMAP4), Post Office Protocol Version 3 (POP3), and the Microsoft Exchange ActiveSync® technology.
- How will you configure backup and recovery? If you use additional storage for backup and recovery, you may need to replicate or somehow provide offsite storage to enable access to the backup data if

MCT USE ONLY. STUDENT USE PROHIBITED

the primary site is offline. If you use lagged database copies, do you need to provide lagged copies at each location to provide recovery if one of the sites is offline?

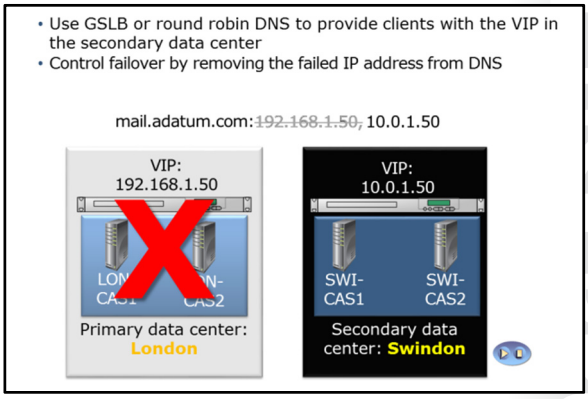
- What kind of data redundancy will you use? Will you use RAID-compatible or Just a Bunch of Disks (JBOD) storage for the databases? Also, how many copies of each database will be located at each site? Will the number of copies be the same for each mailbox, or will some mailboxes require more redundancy than others?

After you have collected the information for your project, you must determine the configuration you will use. The following site resilient designs are the most common:

- One active and one passive site. The most basic site resilient solution includes two data centers. During normal operation, the primary site hosts the active mailbox copies, and all clients connect to the primary site. Mailbox data is replicated to the secondary site along with configuration information that is stored in Active Directory. In this scenario, the secondary site usually has fewer servers, because it is intended to operate as the active site only temporarily during the recovery of the primary site. This same model applies to more than two data centers, but one or more passive data centers would still remain idle until a failover occurs.
- One active, one passive, and one witness site. This is similar to the first solution, but it is typically deployed if either site is designed to host all of the active mailboxes. A file share witness server is located in the third site to maintain quorum in the event one of the other sites is offline. This solution requires more hardware, but it can significantly simplify recovery.
- Two or more active sites. A more complicated solution has active mailboxes in two or more sites during normal operation. Mailbox data is replicated from the active site to another site. Active Directory data is replicated between all sites.

Planning Client Access for Site Resiliency

Multiple Client Access servers are deployed in a load balanced configuration. Because the Client Access server roles in Exchange Server 2013 are not remote procedure call (RPC) endpoints, you do not need to create a Client Access array object. Instead, you group Client Access servers by using network load balancing or DNS. A site resilient Client Access deployment requires Client Access servers to be located in each site. In each site, the deployed Client Access servers are load balanced by using network load balancing or round robin DNS.



Most of the Exchange Server 2013 services use HTTP as the communication protocol, for example Outlook Anywhere, a feature of Microsoft Exchange, MAPI over HTTPS, Exchange ActiveSync, Exchange Web Services, Outlook Web App, and Exchange Administration Center. Using HTTP as the communication protocol enables the clients to easily work with redundancy. An HTTP-based client can accept multiple IP addresses for each namespace. The client attempts to connect by using the first IP address, but if the client cannot connect after about 20 seconds it tries to connect to the next IP address in the list. If you lose the virtual IP address (VIP) for a set of load balanced Client Access servers, the clients reconnect automatically. To provide multiple IP addresses for a namespace, you configure DNS to provide multiple IP addresses to a client during name resolution. For example, if the client asks for webmail.adatum.com, the DNS server may return two or more IP addresses.

MCT USE ONLY. STUDENT USE PROHIBITED

Namespace Planning

Planning a site resilient Client Access deployment requires that you choose namespaces for the following Exchange services:

- Autodiscover.
- Outlook Web App.
- Exchange Control Panel.
- Exchange Web Services.
- Exchange ActiveSync.
- Offline address book (OAB).

In most deployments, you need to define just two namespaces: one for Autodiscover, and one for all of the other services. The fewer namespaces that are used, the easier it is to manage certificates. The table below shows an example namespace for a single site.

Exchange service	Namespace
Autodiscover	autodiscover.adatum.com
Outlook Web Access	webmail.adatum.com
Exchange Control Panel	webmail.adatum.com
Exchange Web Services	webmail.adatum.com
Exchange ActiveSync	webmail.adatum.com
Offline address book	webmail.adatum.com

In a single site configuration, both of the namespaces point to either the VIP for the Client Access servers or to each Client Access server that uses round robin DNS.

Client connections to mailboxes that are hosted in an Exchange Server 2013 DAG do not require a Client Access array namespace as they do in Exchange Server 2010. In Exchange Server 2013, clients connect to any available Exchange Client Access server, and then they are proxied to the DAG member that is hosting the user's active mailbox copy. In a site resilient configuration, you may use the same namespaces across both the primary and secondary sites. That way, if failover occurs, clients keep using the same namespace without requiring reconfiguration.

Single Namespace Site Failover

A single namespace across the sites does not need to be a single point of failure. There are two recommended ways to manage a single namespace across sites, either by using a global server load balancer (GSLB) or by using round robin DNS. Using a hardware device for GSLB costs more than round robin DNS, but there are a few features that may be available depending on the model and manufacturer. A GSLB functions more like DNS than a network load balancer. Rather than proxying network connections like a network load balancer, a GSLB responds to DNS queries based on the configuration criteria. For example, some GSLB devices have features to do the following:

- Respond based on service health. Similar to traditional load balancers, a GSLB can test each service to be sure it is healthy. If the service does not respond, the GSLB removes the service from possible responses until its health is restored.

MCT USE ONLY. STUDENT USE PROHIBITED

- Respond based on geography. The IP address for the site that is geographically closest to the request can be returned. If you deploy a site resilient configuration where locations are on different continents, this option may improve performance because users connect to the closest Client Access servers.

Round robin DNS may lack some features of GSLB, but it still fulfills basic failover requirements and is very simple to configure. To configure round robin DNS for site resilience, you create a DNS host record for the load balanced IP address in each site. If the client is connected to an IP address for a server that goes offline or that otherwise refuses TCP connections, the client reconnects by using the next IP address that DNS returns for that name. This client failover process takes at least 20 seconds for each attempted TCP connection, so if a site will be offline for an extended period of time, you should remove the unavailable site from DNS.

Multiple Namespace Site Failover

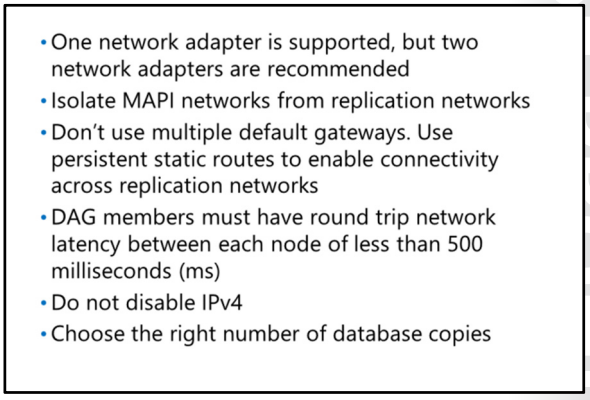
You can use separate namespaces to control traffic if both sites handle active mailbox databases and you want users to connect primarily to the Client Access servers in the primary site. The following table shows this different namespace for each site.

Exchange service	London site namespace	Swindon site namespace
Autodiscover	autodiscover.adatum.com	autodiscover.adatum.com
Outlook Web Access	lon-webmail.adatum.com	swi-webmail.adatum.com
Exchange Control Panel	lon-webmail.adatum.com	swi-webmail.adatum.com
Exchange Web Services	lon-webmail.adatum.com	swi-webmail.adatum.com
Exchange ActiveSync	lon-webmail.adatum.com	swi-webmail.adatum.com
Offline address book	lon-webmail.adatum.com	swi-webmail.adatum.com

Planning Mailbox Database Site Resilience

The general requirements for implementing a DAG are the following:

- You must implement DNS with a host record for each Exchange server. Dynamic updates for DNS are preferred.
- Each Mailbox server must be a member of the same domain. Mailbox servers in different Active Directory domains cannot be members of the same DAG.
- The Mailbox servers that are members of a DAG cannot also be domain controllers. This configuration is not supported.



- The computer name for the Mailbox server must be unique, and it must be 15 characters or less.

Consider the following for DAGs that are deployed for site resilience over a single site DAG:

MCT USE ONLY. STUDENT USE PROHIBITED

One network adapter is supported, but we recommend two network adapters. This way, you can configure a Messaging Application Programming Interface (MAPI) network and a separate replication network.

- You should isolate MAPI networks from replication networks, in order to keep network heartbeats from happening across network interfaces.
- Don't use multiple default gateways. Rather than configuring a default gateway on the replication network, configure persistent static routes to enable connectivity across replication networks.
- Regardless of their location, DAG members must have round trip network latency between each node of less than 500 milliseconds (ms). Lower latency improves replication performance.
- We recommend that you test and validate the network bandwidth and latency for the DAG networks to be sure that they can satisfy the deployment availability goals. Validation tests should take into account all traffic that traverses the networks between cross-site networks, such as database replication, Active Directory replication, client connectivity, message transport, and any other applications running on your network.
- You can use IPv6 only if IPv4 is also configured. You cannot disable IPv4.
- Automatic Private IP Addressing (APIPA) is not supported for DAG members.

Operating System Version

All members of a DAG must be running the same operating system version. All DAG members must be running either Windows Server 2008 R2 or Windows Server 2012. You cannot combine the two operating system versions within the same DAG. The join to the DAG operation fails if you try to join two different versions of the operating system.

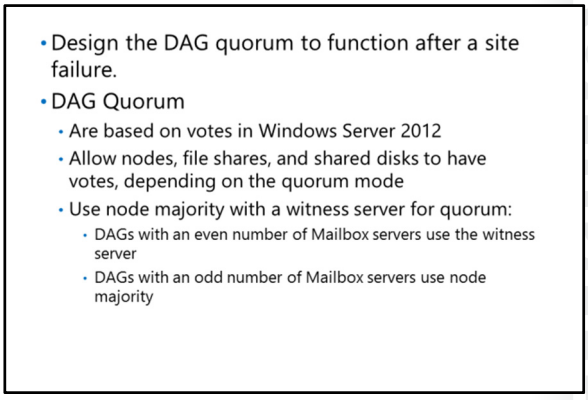
A DAG is based on the use of failover clustering in Windows Server 2012. Only Windows Server 2008 R2 Enterprise or Windows Server 2008 Datacenter editions and Windows Server 2012 Standard and Windows Server 2012 Datacenter editions include failover clustering. Therefore, you can use only these versions for DAG members.

Non-Microsoft Replication API

You may also use a non-Microsoft replication application programming interface (API) instead of the native replication feature. The non-Microsoft solution must replace all native continuous replication functionality. A non-Microsoft solution also must use a DAG as the structure to manage and activate mailbox database copies. Microsoft does not provide primary support for this type of solution.

What is a DAG Quorum?

A Windows Server failover cluster is a collection of servers that act as a single unit to coordinate running a service or application even if a failure occurs. Each server must coordinate with the cluster before trying to bring a service online. This coordination helps protect the application if servers fail. For example, consider a three node cluster where the network cable is unplugged from one of the network switches that connect the cluster nodes. In this case, the nodes cannot communicate with each other. How can the servers determine which node is healthy enough to host the clustered application?



MCT USE ONLY. STUDENT USE PROHIBITED

To help prevent problems caused by a split in the cluster, failover clusters use a voting algorithm to determine whether the cluster has enough votes to maintain a quorum. Because a given cluster has a specific set of nodes and a specific quorum configuration, the cluster determines how many votes are required. If the number of votes drops below the majority, the cluster cannot start. Nodes still listen for the presence of other nodes, in case another node appears on the network, but the nodes do not function as a cluster until a consensus is reached.

The cluster maintains the configuration and keeps track of which node is active and which nodes are passive. Additionally, the cluster decides which passive node to activate if the active node fails. The failover cluster quorum configuration used by an Exchange Server 2013 DAG determines how many failed nodes or failed storage and network components the cluster can sustain while continuing to function.

For example, if there are five votes in the cluster, the cluster continues to function as long as it has at least three available votes. The source of the votes in Exchange Server 2013 can be a node or a file share witness. If a majority of the votes is not available, or if only half of the votes are available, the cluster does not start. Additionally, if a majority of the nodes is not available, Exchange Server 2013 dismounts the databases.

In clusters with shared storage, connectivity to a shared disk can be used to define which nodes should potentially be active in the cluster. In a DAG, there is no central disk. Rather, a witness server is used to establish a quorum in DAGs that have an even number of nodes. In these cases, the witness server functions as an additional member of the DAG for determining the quorum. The witness server is a file share located on a server that is not a member of the DAG.

You can configure a Client Access server as a witness server, as long as it is not also a member of the DAG. Being a witness server adds only a minimal load on the server, and, because it is already under the control of the Exchange Server management group, you do not need to modify permissions. However, if the DAG witness server is not an Exchange Server computer, you need to add the Exchange Trusted Subsystem group as a member of the local Administrators group on the witness server. The witness server does not need to run the same version of Windows Server as the members of the DAG.

Site Resilient Quorum

In a site resilient deployment, you must design the DAG so that it can be recovered even if one of the sites is unavailable. This consideration means that you must determine how to restore the quorum even if all of the servers in either site are unavailable. For example, consider a six member DAG with three nodes in the primary site and three nodes in the secondary site. The witness server is in the primary site. If the secondary site becomes unavailable, there are still four votes available in the primary site. The three Mailbox servers in the primary site maintain a quorum by using the witness server. In this scenario, the three servers in the secondary site cannot access the witness server, so they lose quorum and the databases are dismounted. In this same configuration where you have a minority of nodes in the secondary site, you may still need to recover services in the secondary site if a failure takes the primary site offline. You can establish a majority in this case by configuring an alternate witness server in the secondary data center. This alternate witness server is only used if the secondary site initiates a recovery. After the recovery is run on the secondary site, the databases mount.

After recovery is run to create the quorum in the secondary site, the mailboxes are mounted and made available. This recovery in the secondary site can cause a problem if the servers in the primary site become available. A majority of votes are still in the primary site, so the servers can come online, establish a quorum apart from the servers in the secondary site, and mount the databases. This action results in a split-brain cluster, where servers in both sites have a quorum and have mounted the databases. The split-brain cluster is a problem because there is no way to reconcile the content in the two mounted databases. The Datacenter Activation Coordination (DAC) mode in Exchange Server 2013 employs the Datacenter Activation Coordination Protocol (DACP) to prevent split-brain conditions from occurring. You configure each DAG to use DAC mode, which is recommended for any DAG that has two or more members.

MCT USE ONLY. STUDENT USE PROHIBITED

Each time a DAG member starts, the DACP bit is set to 0, which indicates that mounting is not allowed. The DAG member communicates with other DAG members to find out their status. If the DAG member finds another DAG member that has the DACP bit set to 1, it sets its own DACP bit to 1, and now it can mount databases.

Note: To enable DAC mode for a DAG, use the command **Set-DatabaseAvailabilityGroup <DagName> –DatacenterActivationMode DagOnly** in the Exchange Management Shell.

To support DAC mode with two-node DAGs, the evaluation of whether a node can mount databases also includes the boot time of the alternate witness server in the alternate data center. You should never reboot the remaining single node in a two-node DAG and the alternate witness server at the same time, because if you do, DAC mode may prevent the single remaining node from starting databases. If this happens, you need to reset the DACP bit in the DAG by using the **Restore-DatabaseAvailabilityGroup** cmdlet.

Another way to mitigate split-brain clustering is to host the witness server on a third site, if the two primary sites are hosting an equal number of nodes. Because the witness server is located apart from the two sites with the DAG members, it can more reliably provide a tie breaking vote. If the nodes cannot access the witness server in the third site or the DAG members in the other site, they lose quorum and dismount the active mailbox databases.

Planning DAG configuration

Planning the DAG configuration must take into account maintaining quorum and providing the right amount of resources for the mailbox databases.

- Create fewer DAGs to reduce hardware requirements
- Create a DAG for each active site to reduce dependency on other sites.
- Use a third site to host a witness server to maintain quorum during a site outage.

In a typical two-site deployment that uses the secondary site only for failover scenarios, you can deploy a single DAG. You can locate an equal number of DAG members at each site, or you may choose to have fewer DAG members in the secondary site.

In a two-site deployment that has active mailbox databases in each site, you may choose to deploy two DAGs. If you do, you can configure the DAGs independently, and you help prevent the active mailbox databases from going offline if one of the sites becomes unavailable.

In a three site deployment, only two sites host Mailbox servers. Each site includes the same number of DAG members. The third site hosts the witness server, to help ensure that the loss either of the other two sites does not cause the loss of quorum.

In deployments where active mailbox databases are hosted in two sites, you must consider how the loss of one site affects the active mailbox databases in the other site. To help ensure that a data center outage does not affect the active mailbox databases, you can use a three site deployment or you can deploy multiple DAGs. If you deploy a DAG for each primary location, you provide additional control over the how quorum is established and maintained. The tradeoff is that this greater flexibility requires more servers.

MCT USE ONLY. STUDENT USE PROHIBITED

Planning Database Copies

You may have up to 16 copies of a database in a 16 member DAG. The number of database copies that you need depends on the amount and type of redundancy that you need. For example, if you are using Exchange Server native protection, instead of creating database backups, you need multiple database copies to help protect from failures.

- Do you want to activate the copy in the secondary site when you maintain DAG members in the primary site?
- Do you need redundancy in the secondary site, either for maintenance or to provide protection during a primary site failure?
- Will you be using Exchange native protection?
- Will you be using JBOD storage?
- Do you plan on activating a lagged copy?

Lagged Database Copies

A *lagged mailbox database copy* is a database that uses a delayed replay lag time to commit the log files to the database. This way, you can go back to a point in time up to 14 days in the past. By delaying the replay of logs to a database, you can recover the database to a particular point in the past. Lagged database copies can protect you from the following types of logical corruption:

- Database Logical Corruption. This corruption occurs when the database page's checksum matches, but the data on the pages is logically wrong. It can occur if the Extensible Storage Engine (ESE) attempts to write a database page and the operating system storage stack returns success, even though the data either never makes it to the disk or is written to the wrong place.
- Store Logical Corruption. This symptom indicates that data is added, deleted, or modified in a way that the user does not accept, so the user views it as a corruption. Single item recovery and retention hold provide some protection against this scenario, because all changed items are kept and therefore can be restored. However, particularly if large amounts of data change, it might be easier to recover the database to a point in time before the corruption occurred.
- Rogue Admin Protection. Malicious or rogue administrators might add, change, or remove data from the system in a way that the users see as undesirable. To help protect against this eventuality, you can place the lagged database copies on a server that is under separate administrative control.

You must determine the number and location of the lagged copies that the site resilient design requires. If you plan to use lagged copies as a backup, you never want to get to the point of activating your only lagged copy, because then you would no longer have a backup copy. In this case, you need to have enough database copies available so that you never need to activate the lagged copy, or you need to have multiple lagged copies.

Determining the Right Number of Database Copies

To have a site resilient database, you need to configure at least two database copies, one at each location. There are a number of reasons to configure more than two copies. The following table outlines several options for determining the number of database copies that you need.

Multisite redundancy types	Site 1: copies	Site 1: lagged copies	Site 2: copies	Site 2: lagged copies
Minimum	1	0	1	0
Site 1 redundancy	2	0	1	0

MCT USE ONLY. STUDENT USE PROHIBITED

Multisite redundancy types	Site 1: copies	Site 1: lagged copies	Site 2: copies	Site 2: lagged copies
Multisite redundancy	2	0	2	0
Multisite redundancy with native data protection	2	1	2	0
Multisite redundancy with multisite native data protection	2	1	2	1
JBOD deployment	3	0	3	0
JBOD deployment with multisite native data protection	3	1	3	1

To determine the number of copies that you need, answer the following questions:

- Do you want to activate the copy in the secondary site when you maintain DAG members in the primary site? If not, you should add at least one additional copy in the primary site.
- Do you need redundancy in the secondary site, either for maintenance or to provide protection during a primary site failure?
- Will you be using Exchange native protection? If so, do you need to provide the protection in both sites?
- Will you be using JBOD storage? If so, you should have at least three copies of the data in each site. This way, you can keep at least three good copies of the data in each site, even if a storage failure occurs.
- Do you plan to activate a lagged copy? If so, you should have at least two lagged copies in each site so that if you need to activate one lagged copy you still have another lagged copy available for protection.

Each copy adds to the storage, memory, processor, and network requirements. You should strive to deploy enough copies to meet your requirements, without deploying too many and causing resource problems.

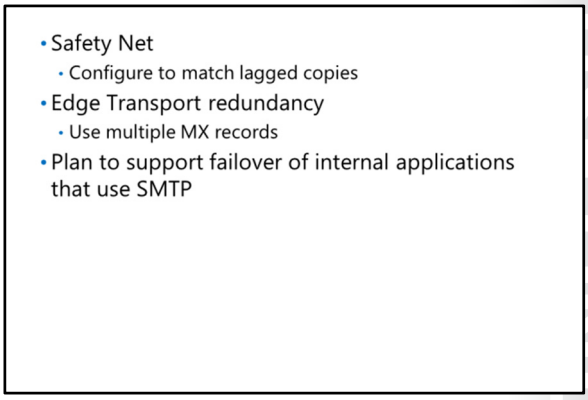
Blocked Database Copy Activation

In some instances, you do not want the database copies in the secondary site to activate automatically. For example, you may want to keep the database copies from automatically activating in a secondary site because you want a person or outside process to decide whether you should fail over to the secondary site. You can configure a database copy on a Mailbox automatically, or you can configure an entire Mailbox server to not activate any database copies that it hosts.

MCT USE ONLY. STUDENT USE PROHIBITED

Planning for Message Transport Site Resilience

Message transport now is the responsibility of the Mailbox server role. The mechanism for transport availability is called *shadow redundancy*. This redundancy causes messages to be persisted to two servers before the message is accepted from the sender. By default, if you deploy a DAG that crosses Active Directory sites, the shadow copy is created on a server in the second site. If the Mailbox server is not a DAG member, the shadow copy is created on another Mailbox server in the local Active Directory site, if possible.



Safety Net

Safety Net helps protect against Mailbox server failures if transaction logs are lost. If a failure occurs and some transaction logs are not replicated to the passive copy, you can use Safety Net to redeliver messages. Safety Net is a special message queue available in the Transport service on every Mailbox server. By default, this queue stores up to two days of messages that were successfully delivered to a mailbox database. If you are using lagged copies, you should configure Safety Net to store data for the same amount of time as the replay delay. This way, Safety Net can redeliver messages to the lagged copy, without needing to replay the lagged transaction logs.

Edge Transport

To provide site resilience for Edge Transport servers, you must also deploy and configure at least one Edge Transport server at the second site. To enable message delivery to the servers at the secondary site, you can configure additional MX records in DNS. An MX record is a weighted pointer to the host name of the Edge Transport server. To redirect messages automatically to the alternate data center when the primary location is unavailable, you can configure multiple MX records. The priority setting for MX records determines the order in which they are used. The MX record with the lowest priority number is contacted first. The MX record for the alternate data center has a higher priority number than the MX record for the primary data center. With this configuration, mail servers attempt delivery to the primary data center first, and, if the primary data center is unavailable, the messages are delivered to the alternate data center.

Messages transported through the alternate data center automatically use the Edge Transport server in the alternate data center for message delivery, because it is the closest Edge Transport server.

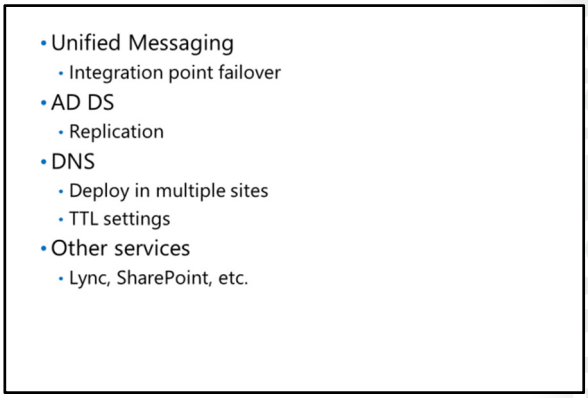
Considerations for SMTP-Based Applications

Some applications, which can include monitoring software, may send email messages directly to an Exchange server instead of determining the authoritative SMTP server by using DNS. You need to configure these applications so that they work after a site failure. If an application requires you to specify an SMTP server to send messages to, you may choose to use round robin DNS records. DNS can include IP addresses for servers in each location. If a site failure occurs, you can remove the records for the servers that are no longer responding. Alternatively, you can leverage global server load balancing to return the closest healthy SMTP server.

MCT USE ONLY. STUDENT USE PROHIBITED

Planning Site Resiliency for Other Services

You should also consider other services as you plan a site resilient configuration, such as AD DS and the network configuration.



Planning for Unified Messaging

The Mailbox server is the endpoint for Session Initiation Protocol (SIP) communications, so by planning a DAG with multiple members in two sites, the endpoint is redundant. The other part of a Unified Messaging (UM) deployment is the UM gateway, which can be more complicated to make highly available. If the UM gateway that Exchange Server 2013 works with is located in the primary datacenter, and the primary site is unavailable, UM fails.

Planning for AD DS

AD DS has a very simple site resilient model. To add redundancy to AD DS, you need to deploy additional domain controllers. You must plan to have an adequate number of domain controllers to support a failover. For example, if you are planning a secondary site, you need to have enough domain controllers available there to support authentication and lookup activity if a failover occurs, not just enough to handle the site while Exchange is active in the primary site. If you are using domain controllers running Windows Server 2008 R2 or newer, a best practice is to have at least one processor core for every eight Mailbox server processor cores. This ratio may be different in your deployment, so you should verify your sizing assumptions in a test environment.

Planning for Network Requirements

You need to plan for network redundancy and properly size for replication and cross site connectivity. If bandwidth between the sites is not adequate, the database copies won't be kept up to date, which could result in a loss of data. You can use the Mailbox Server Role Requirements Calculator to estimate the amount of bandwidth you need to support database replication.

Planning for DNS

DNS services for AD DS clients and internal Exchange users are critical to the overall operation of Exchange. For the servers in your domain, make sure that they can resolve DNS queries and reach AD DS services if one of the data centers is unavailable. For example, you do not want the Exchange servers or the domain controllers in the secondary data center to be using domain controllers in the primary site for DNS.

DNS is also critical for users who access email from outside the network and for those who send email to recipients in your organization. If the external DNS servers responsible for your domain are located in only one of your data centers and that data center fails, none of your external users will be able to resolve the IP address for essential services like Autodiscover. Also, those who try to send email to the users in your organization cannot look up the MX records to determine which servers to send email to. You may choose to deploy external DNS servers in multiple data centers or use an external provider to host your external DNS zones.

To reduce any disruption due to DNS changes, DNS records for client services such as Outlook Web App, Autodiscover, Outlook Anywhere, Exchange ActiveSync, IMAP4, POP3, SMTP, and Exchange Web Services should have a low Time to Live (TTL). The default TTL for a DNS zone hosted in Windows Server is one hour, which enables a client to cache an IP address for up to one hour before requesting the record again. One side effect of lowering the TTL is that the number of DNS lookups that are performed against the DNS servers increases, because the address must be retrieved again each time the TTL expires. We

MCT USE ONLY. STUDENT USE PROHIBITED

recommend that you set the TTL for client-facing DNS records to 5 minutes. Also, you should optimize replication between DNS servers so changes are replicated to all DNS servers quickly.

Planning for Other Services

Other services also connect to or rely on Exchange Server, and you must consider them when you create a site resilient design. Each of these applications, such as the Microsoft Lync® Server 2013 communications software, Microsoft SharePoint® Server 2013, email archiving, and email discovery, must be tested and configured to work in a site resilient design. Each application must be configured differently either to be site resilient or to continue to work if a site failure occurs.

MCT USE ONLY. STUDENT USE PROHIBITED

Lesson 3
Implementing Site Resilience

After planning a site resilient design, you must configure each of the components.

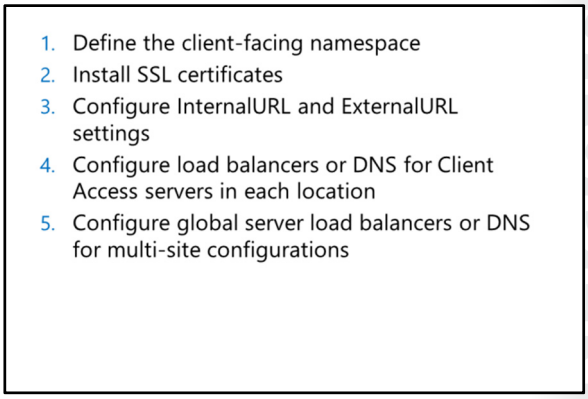
Lesson Objectives

After this lesson, you will be able to:

- Configure a Client Access server for site resilience.
- Configure a DAG for site resilience.
- Describe the site failover process.

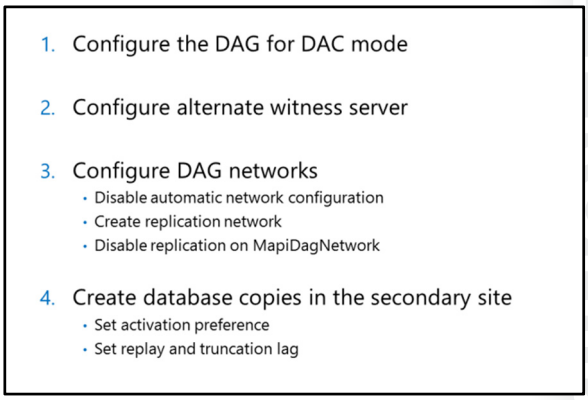
Configuring Site Resilient Client Connectivity

Configuring Client Access for a site resilient deployment is straightforward. To start, you define the client-facing namespace. This includes the fully qualified domain name (FQDN) of each service that will be hosted in each data center. Next, you install Secure Sockets Layer (SSL) certificates and configure InternalURL and ExternalURL settings on each of the Exchange services. In most site resilient deployments, at least two Client Access servers are deployed to each location. To provide high availability, you should configure the Client Access servers with either network load balancing or round robin DNS. Lastly, to achieve site resiliency, you must configure global server load balancers or DNS for multi-site configurations.



Configuring Database Site Resilience

To create a site resilient DAG, start as you would to create a single site DAG. In a two site deployment, create a DAG, define a witness server in the primary site, and define an alternate witness server in the secondary site. If you are deploying a DAG that uses three sites, define a single witness server in the third site. Next, configure the MAPI and replications networks, and add all DAG members. If you are deploying a DAG that has Mailbox servers running Windows Server 2012, you must pre-stage the cluster name object in Active Directory. Next, create mailbox databases, and then create additional database copies. You can distribute database copies across Mailbox servers in a flexible and granular way. You can replicate one, some, or all mailbox databases on a server.



You must specify the following information when creating a mailbox database copy:

- The name of the database you are copying.

MCT USE ONLY. STUDENT USE PROHIBITED

- The name of the Mailbox server that you want to host the database copy.
- An activation preference number. This is called a *preferred list sequence number*, and it represents the activation preference order of a database copy after a failure or outage of the active copy. You typically configure the databases in the secondary site with a higher activation preference, so they are activated only if the copies in the primary site are unavailable.
- The amount of time for log replay delay, in minutes. This value is the replay lag time, which sets how long to wait before the logs are committed to the database copy. Setting the value for replay lag time to 0 turns off log replay delay.
- The amount of time for log truncation delay, in minutes. This is the truncation lag time, which sets how long to wait before truncating committed transaction logs. Setting the value for truncation lag time to 0 turns off log truncation delay.

To enable Datacenter Activation Coordination (DAC), on DAG1 use the following command:

```
Set-DatabaseAvailabilityGroup DAG1 –DatacenterActivationMode DagOnly
```

Before you can configure the DAG networks to separate the MAPI and the replication networks, you must first disable automatic DAG network configuration. You can make this change in the Exchange Administration Center or in the Exchange Management Shell. After you enable manual configuration on the DAG, you can create the replication network and disable replication on the pre-created MapiDagNetwork.

Site Failure Switchover Process

Once the site resilient configuration is complete, you can test the switchover and failover processes. In a two-site deployment, if Site 1 goes offline and a majority of the DAG members is no longer available, you must reconfigure the DAG to reestablish quorum. To do this, first stop the DAG in the primary Active Directory site. This process temporarily stops the DAG members in the site from participating in the DAG. For example, to stop DAG1 in the London Active Directory site, run the following cmdlet:

1. Site 1 goes offline and quorum is lost
2. Stop the DAG in the primary site by running Stop-DatabaseAvailabilityGroup <DAGName> –ActiveDirectorySite: <primaryAdSite>
3. Stop the Cluster service on any servers in the primary site, if they are accessible
4. Restore the DAG in the secondary site by running Restore-DatabaseAvailabilityGroup <DAGName> –ActiveDirectorySite: <SecondaryADSite>

```
Stop-DatabaseAvailabilityGroup DAG1 –
ActiveDirectorySite London
```

Next, stop the Cluster service on any servers in the primary site, if the servers are accessible. This keeps these servers from attempting to rejoin the cluster while operating out of the secondary site. Finally, restore the DAG in the secondary Active Directory site. This causes the DAG members in the secondary site to reestablish quorum by using the available DAG members, including the alternate witness server. For example, to start DAG1 in the Swindon Active Directory site, run the following cmdlet:

```
Restore-DatabaseAvailabilityGroup DAG1 –ActiveDirectorySite Swindon
```

If you configured the servers in the secondary site to block automatic activation, you must remove the activation block before the database copies can activate.

After the primary site is recovered, you must perform a switchover process to restore services. First, you must reincorporate the DAG members in the recovered site by using the **Start-**

MCT USE ONLY. STUDENT USE PROHIBITED

DatabaseAvailabilityGroup cmdlet. For example, to start the switchover process for DAG1 back to the London Active Directory site, run the following cmdlet:

```
Start-DatabaseAvailabilityGroup DAG1 –ActiveDirectorySite London
```

After the passive copies are replicated back to the primary site, you can use the **Move-ActiveMailboxDatabase** cmdlet on the DAG members in the primary site. If you removed an activation block to keep the databases from activating in the secondary site, you need to reconfigure the activation block.

MCT USE ONLY. STUDENT USE PROHIBITED

Lab: Designing and Implementing Site Resiliency

Scenario

A. Datum has deployed a second data center at Swindon, 80 miles away from London, to provide a disaster recovery location for the London data center. The current highly available deployment of Exchange Server 2013 must be extended to fail over to the disaster recovery data center. A. Datum wants to be able to activate the Exchange Server infrastructure within 30 minutes after the primary site fails, and to be able to access all users' mailboxes as well as to send and receive emails. The secondary site will have the same hardware equipment as the primary site, including storage.

Objectives

Students will be able to design and implement site resilience for an Exchange Server 2013 deployment.

Lab Setup

Estimated Time: 75 minutes

Virtual machines	**20342B-LON-DC1** **20342B-LON-DC2** **20342B-LON-CAS1** **20342B-LON-CAS2** **20342B-LON-MBX1** **20342B-LON-MBX2** **20342B-LON-CL1**
User Name	Adatum\Administrator
Password	Pa$$w0rd

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

1. On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In the Microsoft Hyper-V® Manager, click **20342B-LON-DC1**, and then, in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Log on by using the following credentials:
   - User name: **Adatum\Administrator**
   - Password: **Pa$$w0rd**
5. Repeat steps 2–4 for the following virtual machines: **20342B**-**LON-DC2**, **20342B-LON-CAS1**, **20342B-LON-CAS2**, **20342B-LON-MBX1**, and **20342B-LON-MBX2**.
6. In Hyper-V Manager, click **20342B-LON-CL1**, and then, in the Actions pane, click **Start**.
7. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
8. Log on by using the following credentials:
   - User name: **Adatum\Kim**
   - Password: **Pa$$w0rd**

MCT USE ONLY. STUDENT USE PROHIBITED

Exercise 1: Implement Site Resilience

Scenario

Based on the requirements in the lab scenario, you need to design a solution for site resilience for Exchange Server. Also, you need to plan AD DS components and network infrastructure so that they are available on the disaster recovery site, and so that Exchange Server can work on that site. All clients connect to Exchange over the Internet. The clients currently use webmail.adatum.com for Microsoft Outlook® Web App, Microsoft Outlook Mobile Access, and Exchange Web Services. The Autodiscover namespace is autodiscover.adatum.com. The namespace configuration is the same for the secondary site. A Database Availability member and a Client Access server will be deployed in the secondary site. The servers running Exchange Server 2013 in the London data center are in IP subnet 172.16.0.0/25. The servers in the Swindon data center are in IP subnet 172.16.0.128/25. You need to implement the site resilient solution.

The main tasks for this exercise are as follows:

1. Add DNS Entries for LON-CAS1 and LON-CAS2 to the webmail.adatum.com and autodiscover.adatum.com DNS A Records

2. Configure the Client Access virtual directories

3. Prepare the cluster network object for a database availability group (DAG)

4. Create a DAG and add Mailbox servers to it

5. Add a copy of the Mailbox database on LON-MBX2

6. Verify the successful copying of a database

7. Configure an alternate file share witness and configure Datacenter Activation Mode

▶ Task 1: Add DNS Entries for LON-CAS1 and LON-CAS2 to the webmail.adatum.com and autodiscover.adatum.com DNS A Records

1. On LON-DC1, open DNS and create a host (A) resource record for **webmail.adatum.com**, which points to the IP address of LON-CAS2.
2. Create a host (A) resource record for **Autodiscover.adatum.com**, which points to the IP address of LON-CAS2.
3. Create a host (A) resource record for **webmail.adatum.com**, which points to the IP address of LON-CAS1.
4. Verify that there is a record named **autodiscover**, which is mapped to LON-CAS1.
5. Close the DNS Manager.

▶ Task 2: Configure the Client Access virtual directories

1. On LON-CAS1, open the Windows Internet Explorer® browser, navigate to **https://webmail.adatum.com/ecp**, and then log on as **Adatum\administrator** with the password **Pa$$w0rd**.
2. Configure the external access domain for LON-CAS1 and LON-CAS2 to **webmail.adatum.com**.
3. Modify the Autodiscover Internal Uniform Resource Identifier (URI) for LON-CAS1 and LON-CAS2 using **Set-ClientAccessServer** to be **https://autodiscover.adatum.com/autodiscover/autodiscover.xml**.
4. Configure the internal and external hostnames for Outlook Anywhere, a feature of Microsoft Exchange, on LON-CAS1 and LON-CAS2 to **webmail.adatum.com**.

MCT USE ONLY. STUDENT USE PROHIBITED

▶ Task 3: Prepare the cluster network object for a database availability group (DAG)

1. On LON-DC1, open Server Manager, and then open **Active Directory Users and Computers**.
2. In **Active Directory Users and Computers**, enable **Advanced Features**.
3. In the navigation pane on the left, expand **Adatum.com**, and then create a computer object named **DAG1** in the **Computers** container.
4. Change security settings for DAG1 as follows:
5. Exchange Trusted Subsystem group: **Full control**
6. LON-MBX1 (ADATUM\LON-MBX1$): **Full control**
7. Disable the **DAG1** computer account.

▶ Task 4: Create a DAG and add Mailbox servers to it

1. Log on to LON-CAS1 by using **Adatum\Administrator** with the password **Pa$$w0rd**.
2. Open Internet Explorer, and then type **https://webmail.adatum.com/ecp**, and then sign in as **Adatum\administrator** with the password **Pa$$w0rd**.
3. In the Exchange admin center, create a new DAG by using the following settings:
   - Database availability group name: **DAG1**
   - Witness server: **LON-CAS1**
   - Witness directory: **C:\FSWDAG1**
   - Database availability group IP addresses: **172.16.0.33** and **172.16.0.225**
4. Click **Manage DAG membership** for **DAG1**, and then add the following servers:
   - **LON-MBX1**
   - **LON-MBX2**

▶ Task 5: Add a copy of the Mailbox database on LON-MBX2

1. Add a copy of the mailbox database named **Mailbox Database 1** to **LON-MBX2**.

▶ Task 6: Verify the successful copying of a database

1. In tabs, click **Refresh**, and wait until the details pane shows **Mailbox Database 1\LON-MBX2** as **Passive Healthy**. This might take several minutes and up to several hours depending on the size of the database.
2. View details for **Mailbox Database 1\LON-MBX2** and verify the following:
   - Status: **Healthy**
   - Content index state: **Healthy**.
3. Close Internet Explorer.

▶ Task 7: Configure an alternate file share witness and configure Datacenter Activation Mode

1. Use Exchange Management Shell to set LON-CAS2 as the alternate witness server for DAG1.
2. Use Exchange Management Shell to set the database activation mode to DAGOnly.
3. Restart LON-CAS1 and LON-CAS2.

MCT USE ONLY. STUDENT USE PROHIBITED

Results: After completing this exercise, you will have successfully:

- Added DNS entries for LON-CAS2 to the webmail.adatum.com and autodiscover.adatum.com host (A) resource records.
- Configured virtual directories on LON-CAS2.
- Prestaged the cluster network object for a DAG.
- Created a DAG and added mailbox servers to it.
- Added a copy of the mailbox database on LON-MBX2.
- Configured LON-CAS2 as an alternate witness server.

Exercise 2: Validate Site Resilience

Scenario

After completing the site resilient implementation, the next step is to simulate the failure of the primary data center and then to validate the site failover. You also need to validate the process for restoring functionality in the primary data center.

The main tasks for this exercise are as follows:

1. Verify the location of the active Mailbox copy
2. Verify the normal functionality of Outlook
3. Initiate a failure of the active Mailbox copy on LON-MBX1 and verify Outlook functionality
4. Initiate a failure of the Client Access services on LON-CAS1
5. Initiate a failure of the witness server, and test Outlook functionality
6. Recover the DAG in the secondary site, and verify Outlook functionality
7. To prepare for the next module

▶ Task 1: Verify the location of the active Mailbox copy

1. Verify that **Mailbox Database 1** is active on LON-MBX1.

▶ Task 2: Verify the normal functionality of Outlook

1. Log in to LON-CL1 by using **Adatum\Kim** with the password **Pa$$w0rd**.
2. Open Microsoft Outlook 2013 and create a default profile.
3. Send a test email message to amr@adatum.com.
4. Close Outlook.

▶ Task 3: Initiate a failure of the active Mailbox copy on LON-MBX1 and verify Outlook functionality

1. Shut down LON-MBX1 by using Microsoft Hyper-V® Manager.
2. On LON-CL1, logged on as **Adatum\Kim**, open Outlook 2013.
3. Send a test email message to amr@adatum.com.
4. Close Outlook.

MCT USE ONLY. STUDENT USE PROHIBITED

▶ Task 4: Initiate a failure of the Client Access services on LON-CAS1

1. Stop the Microsoft Exchange Health Manager and the World Wide Web Publishing services on LON-CAS1.
2. On LON-CL1, log on as **Adatum\Kim** with the password **Pa$$w0rd**.
3. Open Outlook 2013.
4. Send a test email message to amr@adatum.com.
5. Verify that the message is not in the Outbox folder.
6. Close Outlook.

▶ Task 5: Initiate a failure of the witness server, and test Outlook functionality

1. Shut down LON-CAS1 by using Hyper-V Manager.
2. On LON-CL1, log on as **Adatum\Kim** with the password **Pa$$w0rd**.
3. Open Outlook 2013.
4. Verify that Outlook does not connect.

▶ Task 6: Recover the DAG in the secondary site, and verify Outlook functionality

1. Log on to LON-MBX2, and then stop the DAG in the primary site by running **Stop-DatabaseAvailabilityGroup DAG1 –ActiveDirectorySite London**.
2. Stop the Cluster service on LON-MBX2.
3. Restore the DAG in the secondary site by running **Restore-DatabaseAvailabilityGroup DAG1 –ActiveDirectorySite Swindon**.
4. Mount Mailbox Database 1.
5. On LON-CL1, log on as **Adatum\Kim** with the password **Pa$$w0rd**.
6. Open Outlook 2013.
7. Send a test email message to amr@adatum.com.

▶ Task 7: To prepare for the next module

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

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20342B-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for all remaining virtual machines.

Results: After completing this exercise, you will have successfully:

- Verified the location of the active mailbox copy.
- Verified the normal functionality of Outlook.
- Initiated a failure of the active mailbox copy on LON-MBX1 and verified continued Outlook functionality.
- Initiated a failure of the Client Access services on LON-CAS1.

MCT USE ONLY. STUDENT USE PROHIBITED

- Initiated a failure of the witness server and tested Outlook functionality.
- Recovered the DAG in the secondary site and verified continued Outlook functionality.

Question: Why does Outlook no longer connect to the mailbox after LON-MBX1 and LON-CAS1 are shut down?

Question: In the lab, how is failover for Outlook Anywhere provided?

MCT USE ONLY. STUDENT USE PROHIBITED

Module Review and Takeaways

Best Practice

Carefully plan a site resilient design by testing and validating the final design.
Never rely solely on the high availability features in Windows Server® or Microsoft Exchange Server. You must be familiar with the proper way to manage the software, and you must properly maintain the environment by installing updates and performing maintenance.

Review Question(s)

Question: You are planning a two site, active/active, site resilient deployment. How can you determine how many database copies you need?

Tools

- Exchange Server Mailbox Server Role Requirements Calculator helps you to identify storage and network requirements for a multisite DAG deployment.

MCT USE ONLY. STUDENT USE PROHIBITED

MCT USE ONLY. STUDENT USE PROHIBITED

Module 2

Planning Virtualization for Microsoft Exchange Server 2013

Contents:

Module Overview 2-1

Lesson 1: Planning a Hyper-V Deployment to Exchange Server 2013 2-2

Lesson 2: Virtualizing Exchange Server 2013 Server Roles 2-12

Lab: Planning the Virtualization of Exchange Server Roles 2-19

Module Review and Takeaways 2-27

Module Overview

Many organizations are exploring ways to decrease the cost of providing an IT infrastructure. Frequently, organizations are finding that many of the servers that they have deployed use only a small percentage of the hardware resources that are available on those servers. For this reason, organizations are exploring the option of virtualizing servers that are running Microsoft® Exchange Server 2013. You can deploy Exchange Server 2013 to virtual machines, but you must plan the deployment carefully to ensure that it meets your organization's requirements. This lesson provides an overview of the server virtualization options that are available in the Windows Server® 2012 operating system, and then it provides details about how to plan the Exchange Server 2013 deployment with virtualization.

Objectives

After completing this module, you will be able to:

- Design a deployment of Hyper-V® for Microsoft Exchange Server 2013.
- Design the virtualization of Exchange Server 2013 server roles.

MCT USE ONLY. STUDENT USE PROHIBITED

Lesson 1
Planning a Hyper-V Deployment to Exchange Server 2013

If you are considering using virtualization for an Exchange Server 2013 deployment, you first need a basic understanding of how Hyper-V in Windows Server 2012 works. This lesson introduces Hyper-V and the hardware requirements for using it with Exchange Server 2013.

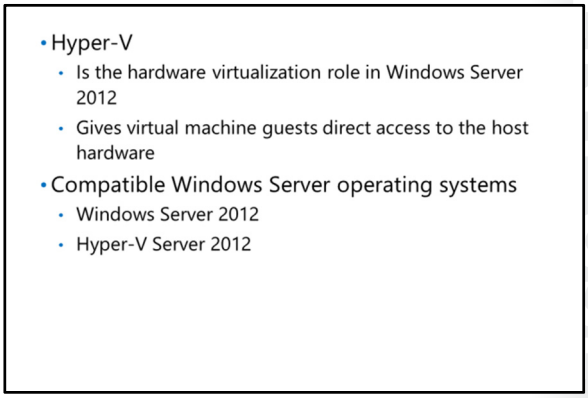
Lesson Objectives

After completing this lesson, you will be able to:

- Describe the purpose and functionality of Hyper-V in Windows Server 2012.
- Describe the new features in Hyper-V for Windows Server 2012.
- Describe the hardware requirements for Hyper-V.
- Describe the hardware best practices for Hyper-V.
- Plan Hyper-V storage.
- Describe the high availability options when you are using Hyper-V technology.
- Identify potential candidates for virtualization.

What Is Hyper-V in Windows Server 2012?

Hyper-V is the hardware virtualization role that is available in Windows Server 2012. Hardware virtualization provides virtual machines direct access to the hardware of the virtualization server. This approach contrasts with software virtualization products, such as Microsoft Virtual Server 2005 R2, that use the operating system of the virtualization server to provide indirect access to the hardware.



You can deploy Hyper-V to a computer that is running Windows Server 2012 by using the Add Roles and Features Wizard. You can configure Windows Server 2012 as a virtualization server by using Hyper-V Manager. Windows Server 2012 can then host virtual machine guests that are running supported operating systems. You can administer virtual machines locally by using the Windows PowerShell® command-line interface, or you can administer them remotely by using the Hyper-V Manager console.

You can install the Hyper-V role both on the Server Core installation of Windows Server 2012 and on Windows Server 2012 in the non–Server Core configuration. Microsoft Hyper-V Server 2012 edition, which includes only the components necessary to host virtual machines, is also available.

Note: In some documentation, the virtualization server, such as a Windows Server 2012 computer that is running Hyper-V, is called the *parent partition*, and a virtual machine that is running on the server is called the *child partition*.

MCT USE ONLY. STUDENT USE PROHIBITED

Hyper-V Features in Windows Server 2012

Windows Server 2012 includes important new features in Hyper-V which enable you to use larger virtual machines and more sophisticated virtualized deployments.

	Resource	Windows Server 2008 R2 Hyper-V	Windows Server 2012 Hyper-V
Host	Logical processors	64	320
	Physical memory	1 TB	4 TB
	Virtual CPUs/Host	512	2048
Virtual machine	Virtual CPUs/Virtual machine	4	64
	Memory/Virtual machine	64 GB	1 TB
	Active virtual machines/Host	384	1024
Cluster	Maximum nodes	16	64
	Maximum virtual machines	1000	8000

The Hyper-V role in Windows Server 2012 provides software infrastructure and basic management tools that you can use to create and manage a virtualized server computing environment. Hyper-V is a hypervisor-based virtualization technology. The hypervisor is a processor-specific virtualization platform that enables multiple isolated operating systems to share a single hardware platform.

The virtualization stack runs within the parent partition, and it manages access to the physical hardware resources. The parent partition creates child partitions, which host the guest operating systems. After the initial Windows Server 2012 installation, the operating system can access the server hardware directly. After you add the Hyper-V role, a thin hypervisor layer is added between the operating system and the hardware resources. The Windows Server 2012 operating system becomes the parent partition, and you can create and manage child partitions from there. Child partitions share access to the hardware resources, and they are presented a virtual view of the resources as virtual devices.

Drivers in the parent partition are used for accessing the server hardware. Child partitions use virtualized devices through virtual server client drivers, which communicate through a virtual machine bus (VMBus) with virtual service providers in the parent partition. Requests to the virtual devices are redirected either through the VMBus or through the hypervisor to the devices in the parent partition.

The VMBus manages the requests. The VMBus is a logical inter-partition communication channel. The parent partition hosts virtual service providers, which communicate over the VMBus to handle device access requests from child partitions. Child partitions host virtual server clients, which redirect device requests to virtual service providers in the parent partition through the VMBus.

Hyper-V Enhancements in Windows Server 2012

Windows Server 2012 provides some important enhancements for Hyper-V. You can use these features to virtualize more applications and to provide higher availability than is possible in previous versions. You should familiarize yourself with these new features, so that you can determine how best to implement them.

Enhanced Storage Features

Improvements to storage include the following:

- Additional Virtual Hard Disk type. This new virtual disk type, called *VHDX*, has a maximum size of 64 terabytes (TB). VHDX helps protect against data loss during storage write failures, and improves performance on 4 kilobyte (KB) disks.
- Offloaded Data Transfer. This feature enables storage related activities, such as copying files between virtual machines, to be offloaded to the storage array. This approach reduces the need for processing within the virtual machines, and improves the performance of storage related tasks.
- Virtual Fibre Channel Adapter. You can use this feature to configure a virtual Fibre Channel host bus adapter (HBA) inside a virtual machine. By using the virtual Fibre Channel HBA, the virtual machine can communicate directly with the storage area network (SAN). This capability is particularly beneficial if you build failover clusters that use virtual machines.

MCT USE ONLY. STUDENT USE PROHIBITED

Enhanced Networking Support

Improvements in networking support include the following:

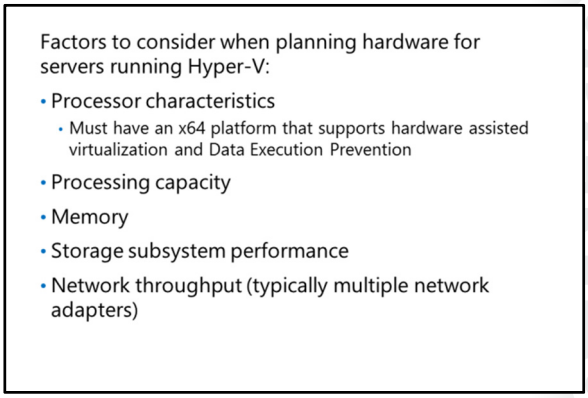
- Port Monitoring and Mirroring. Monitor traffic from ports on Hyper-V virtual switches and mirror the traffic on the port to another virtual port for analysis or reporting.
- Virtual port access control. Isolate networks by using access control lists.
- Trunk mode for virtual network adapters. A single virtual machine network adapter can access multiple virtual local area networks (VLANs).
- Private VLANS. Isolates traffic between virtual machines that are on the same VLAN.

Improved Scalability

Scalability is significantly improved in Windows Server 2012 as well. For example, a single Hyper-V host can have 4 TB of memory and 320 logical processors. Also, virtual machines scale to larger memory sizes, and they can have 64 virtual processors and 1 TB of RAM. Finally, high availability is improved by larger cluster sizes, which can have up to 64 nodes and 8,000 active virtual machines.

Hardware Requirements for Hyper-V

The server that you plan to install the Hyper-V role on must meet the following hardware requirements:

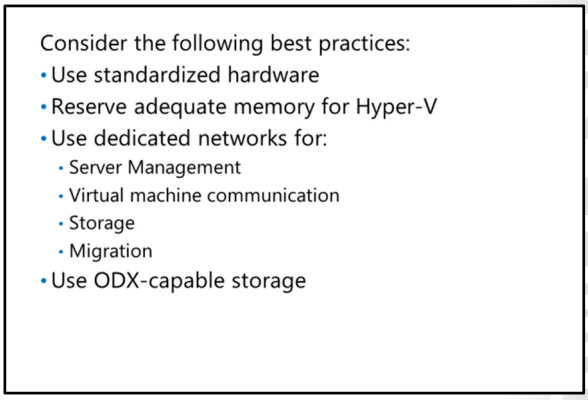


- It must have an x64 platform that supports hardware assisted virtualization and Data Execution Prevention (DEP).
- It must have enough central processing unit (CPU) capacity to meet the requirements of the guest virtual machines.
- A virtual machine that is hosted on Hyper-V in Windows Server 2012 can have up to 64 virtual processors.
- The server must have enough memory to support all of the virtual machines that must run concurrently, plus enough memory to run the host Windows Server 2012 operating system.
- The server must have at least 4 gigabytes (GB) of RAM.
- A virtual machine that is hosted on Hyper-V in Windows Server 2012 can have a maximum of 1 TB of RAM.
- The storage subsystem performance must meet the input/output (I/O) needs of the guest virtual machines. Whether deployed locally or on SANs, you may need to place different virtual machines on separate physical disks, or you may need to deploy a high performance Redundant Array of Independent Disks (RAID), solid-state drives (SSDs), hybrid-SSDs, or a combination of all three.
- The network adapters for the virtualization server must be able to support the network throughput needs of the guest virtual machines. You can improve network performance by installing and using multiple network adapters.

MCT USE ONLY. STUDENT USE PROHIBITED

Hardware Best Practices for Hyper-V

Consider the following best practices for configuring the host computers that run virtual machines:

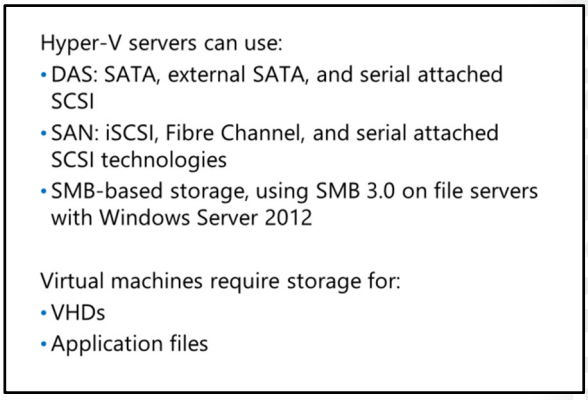


- Simplify and standardize the platform that you are deploying the server virtualization on. If possible, try to develop a standard configuration for all Hyper-V hosts. This approach makes it easier to deploy and manage the environment. For example, if you develop a standard host server build, you can develop consistent deployment practices for the number of virtual machines that are deployed on each host computer.
- Use Windows Server 2012 Server Core or Hyper-V Server 2012 as the virtualization host. The benefits of using Server Core include the following:
- The Server Core installation has fewer components than the full server installation, so there are fewer components to update and less server-maintenance overhead is required. This lower-overhead installation will maintain a higher availability level for hosts and for the virtual machines through fewer software changes, service restarts, and system restarts.
- The Server Core installation provides a smaller surface area for attack, because fewer components are installed.
- A computer with a Server Core installation can be managed from a remote machine using graphical tools or Windows PowerShell, reducing the need to sign on to the server.
- Automate and standardize administration of the virtual server environment. Large organizations may deploy hundreds, or even thousands, of virtual servers over time. The only way to manage them efficiently is to standardize the deployment process and to automate management tasks as much as possible.
- Separate the administration of the host computers and the virtual machines. In most cases, the virtual machine administrators do not need administrative permission to the host computers. For example, if you deploy virtual machines running Exchange Server, the Exchange Server administrators can use Remote Desktop or any of the remote administration tools to manage the Exchange servers. The Exchange Server administrators should not have administrative access to the host computer, because their actions may affect virtual machines other than the intended Exchange servers.
- Reserve adequate memory for the host operating system. The total RAM assigned to running virtual machines should be at least 1 GB less than the total RAM on the server.
- Use dedicated networks for management, live migration, and virtual machine communication. You can create networks by using separate network adapters or by creating separate Quality of Service (QoS) settings if you are using data center bridging connections.
- Configure separate logical unit numbers (LUNs) for the host computer operating system, the virtual machine operating system, and virtual machine storage.
- Use Offloaded Data Transfer-capable storage to improve storage performance of large storage operations.

MCT USE ONLY. STUDENT USE PROHIBITED

Planning Hyper-V Storage

One of the most important components of planning a Hyper-V deployment is planning storage for the virtual environment. Because virtual machines use virtual disks for both the operating system and application files—as well as possibly for data files—planning storage for virtual machines is quite a bit different than planning storage for physical servers. There are three main storage options for virtual machines, each with additional considerations for redundancy, performance, cost, and complexity. The following summarizes these options:



- Virtual hard disks (VHDs). These files can be stored on a local disk, on a SAN, or on a Server Message Block (SMB) 3.0 file share to support Hyper-V failover clustering. Each virtual machine is configured to mount the VHDs. Storage is configured separately.
- Pass-through disks. These can be storage devices attached to the Hyper-V server or to a SAN. If you migrate a virtual machine to a new host, the pass-through disk must be on a SAN. Pass-through disks perform better than a VHD stored on the same disk because there is less overhead.
- Virtual machine presented storage. You can configure virtual machines with virtual Fibre Channel adapters or with network adapters for use with Internet Small Computer System Interface (iSCSI). The storage is presented directly to the virtual machine, and it performs better than a VHD that is stored on the same disk because there is less overhead.

VHD

A VHD is a file format that represents a traditional hard disk drive that you configure with partitions and an operating system. Windows Server 2012 supports booting to a VHD, which means that you can configure a computer to boot into a Windows Server 2012 operating system that is deployed on a VHD or into certain editions of the Windows 8 operating system that are deployed on a VHD. Windows Server 2012 has a new type of VHD, which uses the .vhdx extension. VHDs that have the new format have the following benefits over VHDs that are used in Hyper-V on Windows Server 2008 and on Windows Server 2008 R2:

- VHDs that have the .vhdx format can be as large as 64 TB, but VHDs that have the .vhd format are limited to 2 TB.
- VHDs that have the .vhdx format are less likely to be corrupted if the virtualization server suffers an unexpected power outage.
- The .vhdx format supports better alignment when it is deployed to a large sector disk.
- VHDs that have the .vhdx format can hold larger dynamic and differencing VHDs, which means that the dynamic and differencing VHDs perform better.

A VHD is stored on a file system that is accessible from the Hyper-V server. The underlying storage technology is configured separately.

Physical Storage Options

Hyper-V virtual machines must have access to physical storage that is attached to the host computer. You can use the following types of physical storage with a server running Hyper-V:

- Direct-attached storage (DAS). Storage attached to the management operating system. You can use Serial Advanced Technology Attachment (SATA), external Serial Advanced Technology Attachment,

MCT USE ONLY. STUDENT USE PROHIBITED

Parallel Advanced Technology Attachment, small computer system interface (SCSI), serial attached SCSI, universal serial bus (USB), and FireWire.

- SAN. You can use iSCSI, Fibre Channel, and serial attached SCSI technologies. You can use virtual Fibre Channel adapters for Fibre Channel–based storage or virtual network adapters for iSCSI-based storage within a virtual machine to present storage directly to the guest.
- SMB-based storage. Windows Server 2012 supports VHDs in the .vhdx format that are stored on SMB 3.0 file shares. This option is an alternative to storing .vhdx files on iSCSI or Fibre Channel SAN devices. When you create a virtual machine in Hyper-V on Windows Server 2012, you can specify a network share when you choose the VHD location or when you attach an existing VHD. The file share must support SMB 3.0. To use this option, you must place VHDs on file shares that are hosted on file servers with Windows Server 2012. Previous versions of Windows Server do not support SMB 3.0.

How Virtual Machines Use Storage

Virtual machines use the physical storage that is attached to the servers in the following ways:

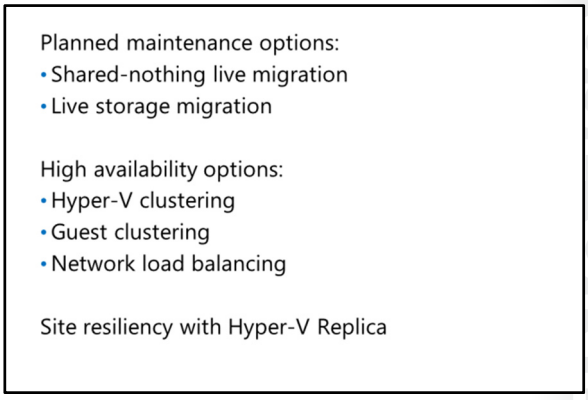
- Storage for virtual hard disk files. All operating system files and application files in a virtual machine are stored in a virtual hard disk file. When you plan the storage for virtual machines, you must consider how much space these files need.
- Storage for snapshots. Each snapshot creates a new, automatic virtual hard disk file. Changes to the operating system or other changes in the virtual machine are written to the file. The file must be stored in the same disk volume as the VHD file. If you intend to use snapshots as part of your virtualization strategy, you must plan for additional space for these files.

 **Note:** Snapshots of Exchange servers are not supported and should never be used in a production deployment of Exchange Server 2013. Using snapshots with Exchange servers may have unpredictable consequences, in part because Exchange Server maintains its own state information across multiple servers or in Active Directory® Domain Services (AD DS).

- Storage for saving virtual machine state. When you save a virtual machine, the virtual machine memory is written to the hard disk. The amount of space required is approximately the same as the virtual machine RAM.
- Storage for failover scenarios. If you are planning to implement high availability for virtual machines by using host failover clustering, you must provide shared storage on a SAN for the virtual machines. All host machines that are part of the failover cluster must be able to access the shared storage.
- Application data storage. Like physical computers, virtual machines are likely to require access to data. When you configure an Exchange server running on physical hardware, one of your design decisions is how to configure the hard drive to store the Exchange Server databases and transaction logs. You must make similar design decisions for the Exchange servers running in virtual machines.

MCT USE ONLY STUDENT USE PROHIBITED

High Availability and Site Resilience in Hyper-V



Most organizations have some applications that are business critical and that must be highly available. To make an application highly available, it must be deployed in an environment that provides redundancy for all of the components that the application requires. Exchange Server 2013 includes features such as database availability groups (DAGs) to help protect from server failures. Exchange Server 2013 also makes it possible for you to load-balance Client Access servers for high availability. In addition, Windows Server 2012 includes features that work with Hyper-V to provide redundancy for applications. Many of the Hyper-V high availability features can be deployed in conjunction with Exchange Server 2013 availability features, as long as the supportability guidance is followed.

When you deploy a server virtualization that uses Hyper-V, you can use or build on the high availability options available in Windows Server. The following options are available to build a highly available virtualized application by using Hyper-V:

- *Host clustering*. With host clustering, you configure a failover cluster by using the Hyper-V host servers. A failover cluster consists of two or more computers, or *cluster nodes*, along with the storage and network infrastructure. When you make an application highly available, you configure the application so that it can be moved from one of the nodes to another. If a node fails, the virtual machine can automatically restart—or *fail over*—to another node. When you configure host clustering for Hyper-V, you configure the virtual machine as highly available. With host clustering, the virtual machine operating system and applications or services running in the virtual machine do not need to be compatible with failover clustering. Because the failover is at the virtual machine level, there are no dependencies on the virtual machine. Exchange Server 2013 can be deployed on clustered Hyper-V servers to provide redundancy for Mailbox and Client Access servers.
- *Guest clustering*. Guest failover clustering works just like physical server failover clustering, except that the cluster nodes are virtual machines rather than physical servers. In this scenario, you create two or more virtual machines, add them to a failover cluster, and then enable an application or service for high availability. If you deploy the virtual machines on separate Hyper-V host computers, you help protect the application or service against the failure of a single host computer. A DAG uses guest clustering to maintain availability across members. Each DAG member should be hosted on separate Hyper-V hosts, to maintain availability if a host becomes unavailable.
- *Network load balancing*. Network Load Balancing (NLB) works with virtual machines the same way that it works with physical hosts. When you configure an NLB cluster, you must install and configure the application on both virtual machines. Then, you either configure the NLB feature in Windows Server or you configure a hardware load balancing solution. Client Access servers can be load balanced this way whether they are deployed as physical servers or virtual servers.

You can use the following features to move virtual machines between hosts with minimal downtime:

- *Live Migration*. You can use this feature to migrate a running virtual machine between hypervisors without the need for shared storage or for other special hardware. You can move a running virtual machine from a hypervisor that needs maintenance or to a hypervisor that has more capacity, a capability that helps increase the availability of the virtual machine.
- *Live Storage Migration*. You can use this feature to move the data stored within VHDs from one storage location to different storage locations without needing to shut down the virtual machine.

MCT USE ONLY. STUDENT USE PROHIBITED

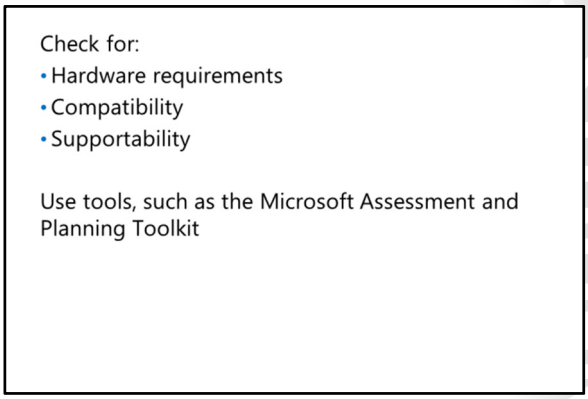
New storage can be presented to the hypervisor, and you can move the VHDs on the virtual machines to the new storage. This way, you can provide more capacity or improve performance without needing to shut down the virtual machine.

You can also use Hyper-V to provide site resilience by using the Hyper-V Replica feature. You can use Hyper-V Replica to replicate a virtual machine from one location to a second location. Hyper-V Replica keeps a transaction log of all storage changes, and it replicates the changes to a Hyper-V server in a secondary site. The passive copy is available to bring online if needed.

Note: Not all features of Hyper-V are supported for use with Exchange Server 2013. Virtualization support for Exchange Server 2013 is discussed in detail in Lesson 2 of this module.

Identifying Virtualization Candidates

You can use virtualization to address many business and IT requirements, but you cannot virtualize all servers and applications. Before you implement virtualization, you need to identify which applications and servers are the best candidates for virtualization.



Note: You can use the Microsoft Assessment and Planning Toolkit (MAP) to identify physical servers that are good candidates for virtualization. By using MAP, you can perform a hardware inventory of the servers in the organization, gather performance information about the servers, and generate reports that provide information about the servers that are the best candidates for virtualization.

To determine whether a server is a good candidate for virtualization, consider the hardware, compatibility, and support requirements.

Hardware Requirements

Typically, a virtual machine requires approximately the same resources as a physical server. For example, if a physical server is currently using 1 GB of RAM, a virtual machine uses the same amount of RAM, assuming it is running the same operating system and applications. When you plan resource utilization on the host computer, remember that the host computer has overhead and requires additional resources. For example, if the virtual machine requires 1 GB of RAM, the hypervisor may require 1.2 GB of RAM. This overhead varies by hypervisor, guest operating system, and application, so you should test the overhead needs in your own environment to get accurate sizing.

In some cases, a server workload may require hardware resources that make it impractical to deploy the workload on a virtual machine. For example, if a server running Exchange Server 2013 requires the same resources as an entire physical server to perform adequately, you should not virtualize the server.

Compatibility

You also must determine whether the application can run in a virtualized environment. Business applications range from simple applications to complex, distributed, multi-tier applications. Consider the requirements for specific application components, such as specific needs for communication with other infrastructure components, and requirements for direct access to the system hardware. Some lightly used

MCT USE ONLY. STUDENT USE PROHIBITED

web servers can be easily virtualized, but back-end components may need to continue to run on dedicated hardware.

Applications and services that have specific hardware or driver requirements are generally not well-suited to virtualization. For example, an application may not be a good candidate for virtualization if it contains low-level drivers that require direct access to the system hardware. This access may not be possible through a virtualization interface, or it may negatively impact performance.

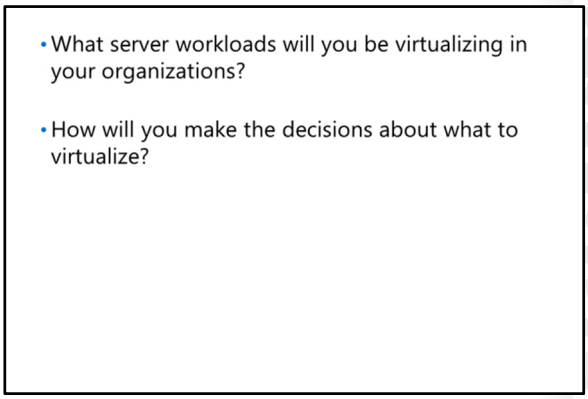
Supportability

Evaluate whether the operating system and the application are supported in a virtualized environment. Verify the support parameters for the application to ensure that it is deployed correctly.

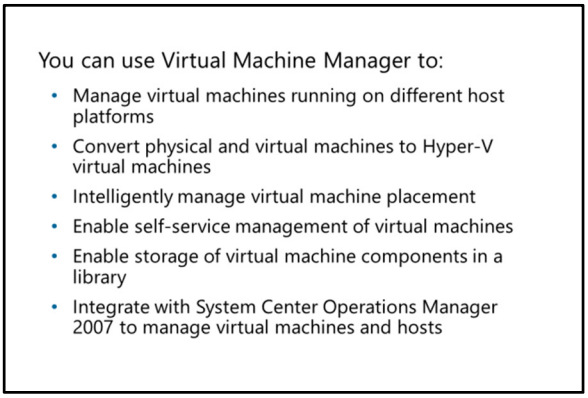
Discussion: Determining Virtualization Requirements in Your Organization

Discuss the following question:

What server workloads will you be virtualizing in your organizations? How will you make the decisions about what to virtualize?



Using Microsoft System Center 2012 to Manage Virtual Environments



Microsoft System Center 2012 Virtual Machine Manager is the primary tool that you use to manage virtual machines, whether they are deployed in Hyper-V or in other virtual environments. Virtual Machine Manager (VMM) is a management tool that you can use to manage multiple physical host computers as well as the virtual machines running on the host computers.

VMM provides the following functionality:

- Enables management of virtual environments running on different host platforms. You can use VMM to manage host computers and virtual machines running Windows Server 2008, Windows Server 2008 R2 with Hyper-V, Hyper-V Server 2008 R2, Windows Server 2012, Hyper-V Server, VMware ESX Server, and Citrix XenServer. You can manage the host server configuration, and you can deploy and manage virtual machines on the host servers by using a single interface.
- Physical and virtual machine conversion. You can use VMM to convert a physical machine to a virtual machine while the physical machine is online. You can also use VMM to convert VMware-based virtual machines to Hyper-V.
- Intelligent virtual machine placement. If you create a new virtual machine or use VMM to move a virtual machine from one host to another, VMM analyzes the available physical hosts, and then it recommends the best location for the virtual machine. You can integrate this process with System Center 2012 - Operations Manager, which enables Intelligent Placement to factor in past performance characteristics to find the best possible match between the virtual machine and its host hardware.

MCT USE ONLY. STUDENT USE PROHIBITED

- VMM Library. VMM provides a centralized library to store various virtual machine components, such as offline machines, templates, and virtual hard disks. You can use the components in the library to deploy virtual machines rapidly by using standardized templates.
- Windows PowerShell integration. VMM is built on the command line and scripting environment provided by Windows PowerShell. VMM provides Windows PowerShell cmdlets that you can use to automate VMM management tasks.
- System Center Operations Manager integration. VMM includes Performance and Resource Optimization (PRO), which you can use to manage virtual resources dynamically by using management packs for System Center Operations Manager. You can use PRO to set rules for moving or configuring virtual machines based on the host server performance.

MCT USE ONLY. STUDENT USE PROHIBITED

Lesson 2
Virtualizing Exchange Server 2013 Server Roles

Exchange Server relies on the hardware it runs on, whether that hardware is physical or virtualized. You should fully understand the support guidelines for virtualizing Exchange Server to help ensure that, no matter what hardware you choose, Exchange Server performs as expected on that hardware. This lesson covers the general support guidelines for using virtualized hardware for Exchange Server 2013. It also covers which questions you should ask to best determine how to use virtualization in your deployment.

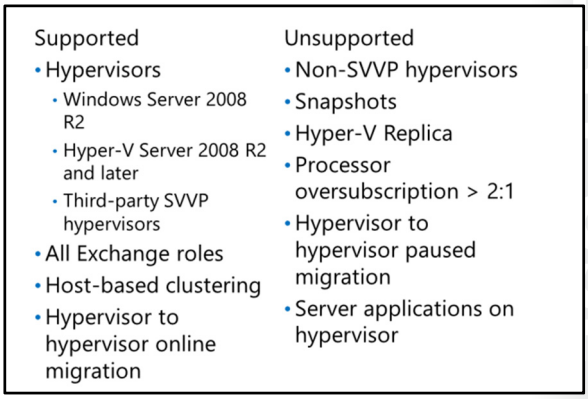
Lesson Objectives

After completing this lesson, you will be able to:

- Describe the Hypervisor requirements for virtualizing Exchange Server.
- Describe the virtual machine requirements for Exchange Server.
- Design a virtualized Mailbox server deployment.
- Explain the considerations for implementing high availability.
- Design a virtualized client access deployment.
- Describe the benefits of using virtualization for Exchange Server 2013/

Hypervisor Requirements for Virtualizing Exchange Server

Exchange Server 2013 has specific requirements for deploying it in a virtual environment. To begin with, you must select virtualization software that is supported. The following hypervisors are supported:



- Windows Server 2008 R2 with Hyper-V technology.
- Hyper-V Server 2008 R2.
- Windows Server 2012.
- Hyper-V Server 2012.
- Other hypervisors that are validated by the Server Virtualization Validation Program (SVVP).

You can deploy all Exchange Server 2013 server roles in a virtual machine. However, there are limitations on how Exchange Server 2013 is configured. The following hypervisor requirements must be met:

- The hypervisor must be dedicated. You cannot use it as an AD DS domain controller, an Exchange server, or for any other server software. You may install management agents for monitoring, antivirus software, backup software, and other management tools on the hypervisor as needed.
- You can use hypervisor-based high availability features, such as failover clustering and migration technology, as long as the virtual machines are not put into a saved state on disk. An example of an unsupported migration option is the Quick Migrate feature in Windows Server 2008 R2 with Hyper-V technology. Planned migrations between hypervisors either require an online migration, or they require that the virtual machine be shut down on the first hypervisor and booted on the second hypervisor. An example of a supported migration technique is the Live Migration feature in Windows Server 2008 R2 with SP1. The implementation of the virtual machine migration is supported by the

MCT USE ONLY. STUDENT USE PROHIBITED

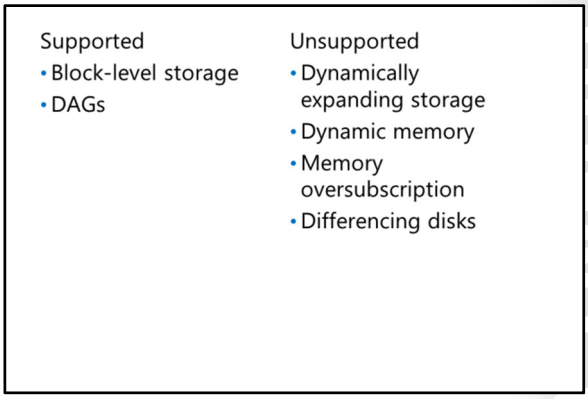
hypervisor vendor. Therefore, you must ensure that your hypervisor vendor supports migration of Exchange virtual machines and that the hypervisor is properly configured for the migration. An unplanned migration occurs if a hypervisor fails. In that case, the guest virtual machines must boot when the virtual machine is activated on the secondary hypervisor.

- The ratio of virtual processor to logical processor can be no more than 2:1, but we recommend a ratio of 1:1. For example, for a hypervisor running on a physical server with 16 logical processors, you should not have more than 32 virtual processors defined in the all guest virtual machines. When you determine the maximum number of virtual processors to run on each hypervisor, you must also account for the amount of processing overhead that the hypervisor needs. If you are using processors that support Intel Hyper-Threading Technology, for the best results, do not use the hyper-threaded cores when you calculate the ratio.

Reference Links: For more information about Exchange Server 2013 Virtualization, see Exchange 2013 Virtualization at http://go.microsoft.com/fwlink/?LinkId=290687

Virtual Machine Requirements for Exchange Server

If you virtualize Exchange Server, you must meet the following storage requirements:



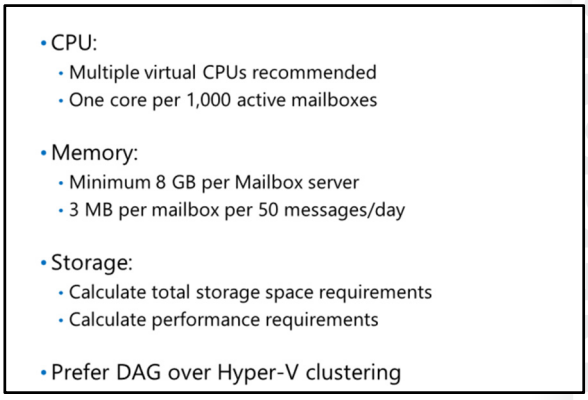
- Virtual machine snapshots and replication are not supported for virtual machines running Exchange Server. You can use snapshots to create a point-in-time copy of the state of the virtual machine on disk, and then be able to revert the virtual machine to the saved state later. Reverting the state of a virtual machine to a snapshot can seriously damage the state information and the database consistency, because Exchange maintains state information separate from the information that a hypervisor snapshot captures. Similarly, Hyper-V Replica is also unsupported, because it replicates changes outside the native Exchange Server processes. Damage to Exchange Server databases and configuration information can occur any time changes are made outside of these processes.
- For Exchange Server virtual machines, the system disk must meet the minimum requirements for the installed operating system and the paging file. For Windows Server 2012, the virtual machines need at least 32 GB plus the size of the allocated virtual memory. This value provides enough storage for the operating system files and paging file disk requirements. For example, if you allocate 16 GB of memory to a virtual machine in Windows Server 2012, the guest operating system disk needs at least 48 GB.
- Memory oversubscription techniques are not supported. Exchange Server maintains and optimizes its own cache by using system memory. The database cache optimizes and tunes performance by using all available memory to reduce the number of disk I/O operations per second. If the dynamic memory feature reduces the amount of memory available to the database cache, Windows memory manager pages the database cache to the Windows paging file of the disk. Also, if the hypervisor exhausts its available memory, the hypervisor will page the oversubscribed virtual machine memory to disk. In either case, if the database cache is paged to disk, it can no longer reduce disk I/O. Dynamic memory and memory oversubscription lead to significant performance degradation, and are therefore not supported.

MCT USE ONLY. STUDENT USE PROHIBITED

- Storage presented to an Exchange Server virtual machine must be block-level storage. The storage can be presented directly to the virtual machine by using virtual Fibre Channel or iSCSI. Storage can also be SCSI pass-through disks or fixed size virtual disks. If you configure iSCSI storage directly to the guest, be sure to configure the entire network to support and optimize iSCSI traffic, such as allowing jumbo frames all the way to the virtual machine.
- Dynamically expanding virtual disks are not supported. Performance during the expansion of the virtual disks is poor, and the expansion requires enough storage on the underlying hardware to finish. If there is not enough storage to expand the disk, the expansion fails, resulting in failures writing to the disk.
- Some hypervisors can be configured to create differencing virtual hard disks. These disks have a static parent VHD file and a separate VHD where all changes are made. Installing Exchange Server on or using differencing disks for storage of any application files is not supported.

Designing a Virtualized Mailbox Server Deployment

Sizing the Exchange Server Mailbox server is important for both physical servers and virtualized servers. The Mailbox server role requires that the processor, memory, and storage configuration be correct for a virtual environment.



To design Mailbox services, you must identify the information required for mailboxes and public folders. Typically, the information you gather helps you determine the size of databases that need to be accommodated and the processing load that those databases will place on the Mailbox servers.

To design mailbox databases, consider the following factors:

- Number of users. A larger number of users typically increases disk utilization.
- Frequency of usage. Higher frequency of usage typically increases disk utilization.
- Size of mailboxes. Larger mailboxes combined with a larger number of users increases overall database size.
- Service level agreements (SLAs). To meet your recovery requirements, you may need to keep databases small in order to reduce restore times.

A best practice in Exchange Server 2013 is to locate multiple databases on a single LUN, because the disk I/O is random. By storing database files and log files on separate volumes or disks, you can replay transaction logs after a database restore if a database is lost due to a failed volume or disk. This is especially useful if you use backups for recovery.

CPU Requirements

Exchange Server 2013 requires a 64-bit processor and a 64-bit operating system. Exchange Server 2013 supports two specific processor architectures: AMD64 and Intel Extended Memory 64 Technology. It does not support Itanium processors. Exchange Server 2013 can take advantage of multicore processors, which can process multiple tasks at the same time.

The number of processor cores required for a Mailbox server varies, depending on the number of mailboxes and how intensely they are used. For average usage, a single processor core can support

MCT USE ONLY. STUDENT USE PROHIBITED

approximately 1,000 active mailboxes. Average usage is defined as a user who sends 10 messages a day and receives 40 messages a day.

Memory Requirements

The memory requirements for Exchange Server 2013 vary, depending on the number of mailboxes and how intensely they are used. The minimum recommended RAM for a Mailbox server is 8 GB. A server that combines multiple roles should have a minimum of 8 GB of RAM.

When calculating the memory required for a Mailbox server, take the minimum memory required, and then add additional memory for each user based on their messaging volume. As a general rule, for every 50 messages per day sent or received, you should allocate 3 megabytes (MB) per user. For example, if the average user in your organization sends and receives 100 messages per day, you should allocate 6 MB per user in addition to the minimum RAM for the Mailbox server configuration. As outlined earlier, memory cannot be oversubscribed, so you must estimate the memory configuration to properly configure the hypervisor and the virtual machines. You must also include storage space for the paging file, which must be at least 10 MB larger than the amount of RAM assigned to the virtual machine.

Storage Considerations

You must determine what type of storage you will use for the virtualized Mailbox servers. You must store the operating system on a VHD. However, you can choose VHD, pass-through disk, or directly presented SAN storage for the Exchange files.

Regardless of where you store the data, you need to calculate accurately the storage requirements for the databases. When you do, you need to consider more than just the raw size of each mailbox in the database. The following factors contribute the total storage requirements:

- Indexes. Each index uses approximately five percent of the mailbox database disk space. This index is placed in the same location as the database. In most cases, you may want to enable indexing on databases to speed up searches.
- Single item recovery. Single item recovery retains deleted messages in a database for a specified period of time. When you enable single item recovery, the database size increases.
- Personal archives. A personal archive is typically used for longer-term retention of mailbox content. If you enable personal archives, the database size may increase. If you plan to use a recovery database, you must have sufficient disk space available to restore the database and transaction logs.

You must also consider storage performance. Whether the storage is presented to the hypervisor or directly to the virtual machines, the storage must be able to perform adequately. If multiple virtual machines are using the same storage, you must ensure that the aggregated performance requirements are met. To estimate the requirements, use the Exchange Mailbox Server Role Requirements Calculator. Then always verify the configuration by using sizing tools such as the Exchange Jetstress tool.

Consider the following:

- Replicated database copies increase the amount of storage space required. If your organization uses DAGs to replicate mailbox databases for high availability, consider the number of database copies when you calculate how much disk space you need and what it costs.
- Regardless of whether they are locally attached or part of a SAN, slower disks have a lower cost per GB than faster disks. Exchange Server 2013 has reduced disk I/O requirements, so large capacity 7,200 RPM disks are suitable for many organizations. You can obtain 7,200 RPM disks of equal size with the SATA or serial attached SCSI interface. Serial attached SCSI disks cost slightly more than SATA disks, but, in testing at Microsoft, serial attached SCSI disks had a 50 percent lower failure rate than SATA disks, so the extra cost might be justified for your organization.
- DAS is significantly less expensive than a SAN. Therefore, DAS is preferable if you use DAGs to create multiple replicated copies of data. You can purchase external drive arrays and use them to connect a

MCT USE ONLY. STUDENT USE PROHIBITED

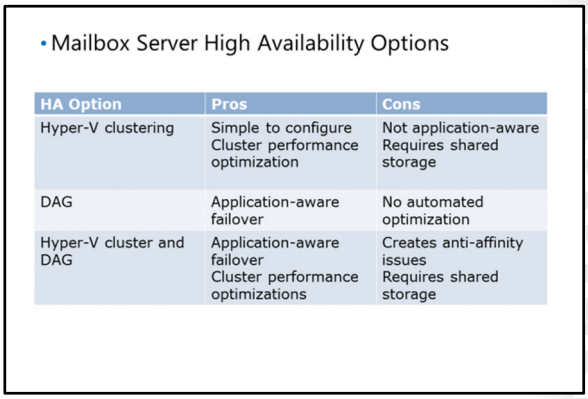
large number of disks to a single server. The lower reliability of DAS is mitigated by the multiple database copies in the DAG. But if you have a SAN that has available space, you might prefer to use the SAN for the higher reliability it provides.

- Some organizations have a significant investment in SANs for all server storage. If you use a SAN, the increased reliability may mean that you choose to implement fewer database copies in a DAG. You can also keep some database copies on a SAN and others on DAS. However, you must consider how this affects your Hyper-V high availability configuration.
- Use the Exchange Server Mailbox Server Role Requirements Calculator to help you plan the storage configuration of Mailbox servers. This spreadsheet contains many calculations to help you accurately estimate the hardware requirements to support a specific number of users with a specific storage configuration. You can download this tool, which is updated regularly, from the Microsoft website.

High Availability Considerations

Both Exchange Server and Hyper-V have options for high availability, and you must decide on which options to use. You can choose to use a DAG, Hyper-V clustering, or both. Hyper-V clusters alone do not provide application-aware failover.

The following table compares the options available for deploying virtualized Mailbox servers.



High availability option	Pros	Cons
Hyper-V clustering	Simple to set up Cluster optimization with System Center 2012	Not application-aware Failover requires a reboot Requires shared storage (SAN or SMB 3.0 file share)
DAG	Application-aware failover	No automated cluster optimization
Hyper-V clustering and DAG	Application-aware failover Cluster optimization with System Center 2012	Creates anti-affinity issues Requires shared storage

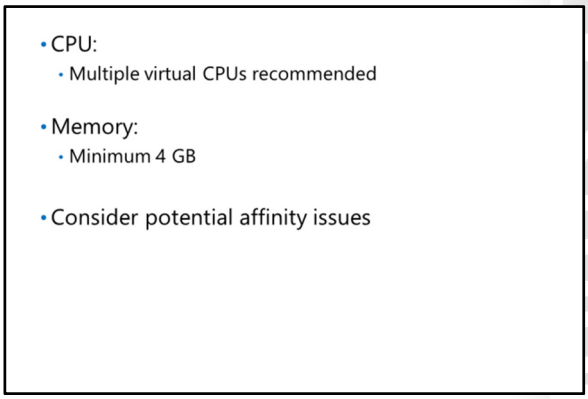
To create a highly available design, you need to analyze all components. For example, if you deploy multiple DAG members on the same SAN storage, the SAN storage is a single point of failure. To eliminate the SAN storage as a single point of failure, either configure multiple SAN storage devices or use another storage option, such as SMB 3.0 file shares or DAS. Another potential problem can occur if virtual machines that are members of the same DAG are running on the same host. If the host fails, multiple DAG members go offline, which can cause a failover with significant losses to an active DAG member. A host

MCT USE ONLY. STUDENT USE PROHIBITED

failure can also cause the DAG to lose quorum if a majority of the DAG members are not available, and therefore to cause all of the databases to go offline. In a non-clustered configuration, you must deploy DAG members to different hosts. In a clustered configuration, either you can deploy DAG members to separate Hyper-V clusters, or you can set the **AntiAffinityClassName** property on each virtual machine in the DAG. If you set this property, the cluster avoids activating multiple virtual machines that have the same **AntiAffinityClassName** value on a single host. You need to have enough Hyper-V hosts to activate all of the virtual machines in the event of a failure. For example, if you have an eight node DAG on a seven node Hyper-V cluster, you do not have enough Hyper-V hosts to satisfy the anti-affinity requirements.

Designing a Virtualized Client Access Deployment

You can virtualize the Client Access server role based on the requirements discussed in the previous topics.



Hardware Requirements

A Client Access server requires at least 4 GB of RAM, and the paging file must be at least 10 MB larger than the amount of RAM assigned to the virtual machine.

High Availability Considerations

The Client Access server role does not have a native high availability feature. It relies on NLB to provide load balancing and failover. You can deploy and load balance multiple Client Access virtual machines just like physical servers. If you are using the Windows Network Load Balancing feature, you must enable spoofing of media access control (MAC) addresses in the virtual machine settings.

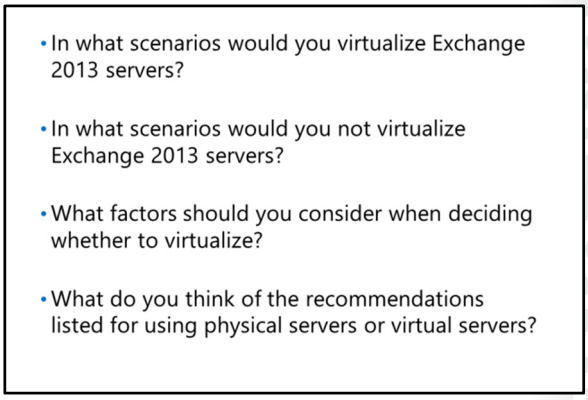
Similar to the considerations for Mailbox servers, when you design for high availability, you need to analyze all components. For example, if you deploy multiple Client Access servers on the same SAN storage, the SAN is a single point of failure. To eliminate the SAN storage as a single point of failure, configure multiple SAN storage devices or use another storage option, such as SMB 3.0 file shares or DAS. Another potential problem can occur if multiple Client Access servers are running on the same host. If the host fails, multiple Client Access servers go offline. If all of the Client Access servers are hosted on the failed host, the Client Access servers may become unavailable. In a non-clustered configuration, you must deploy Client Access servers to different hosts. In a clustered configuration, either you can deploy Client Access servers to separate Hyper-V clusters, or you can set the **AntiAffinityClassName** property on each Client Access virtual machine. If you set this property, the cluster avoids activating multiple virtual machines that have the same **AntiAffinityClassName** value on a single host. Be sure to have enough resources to satisfy the anti-affinity requirements. For example, if you have six Client Access servers deployed on a six node Hyper-V cluster, you do not have enough Hyper-V hosts to satisfy the anti-affinity requirements if one node is offline.

MCT USE ONLY. STUDENT USE PROHIBITED

Discussion: Using Virtualization for Exchange Server 2013

Discuss the following questions:

- In what scenarios would you virtualize Exchange 2013 servers?
- In what scenarios would you not virtualize Exchange 2013 servers?
- What factors should you consider when deciding whether to virtualize?
- What do you think of the following recommendations?



- Use physical servers if:
- Average physical server utilization is greater than 65 percent.
- You are using DAGS and need to provide for one or more server failures.
- Use virtual servers if:
- Average physical server utilization is less than 65 percent.
- You are deploying a single Exchange server.
- You are deploying in a recovery site or branch office where performance is not critical.

MCT USE ONLY. STUDENT USE PROHIBITED

Lab: Planning the Virtualization of Exchange Server Roles

Scenario

A. Datum is considering virtualizing some resources that are currently hosted on physical machines, including Exchange services, using Hyper-V. You need to provide information about virtualization possibilities and the main features of Hyper-V, and you need to propose a design to virtualize Exchange services.

Objectives

After completing this lab, you will be able to :

- Use the Exchange Server 2013 Mailbox Server Role Requirements Calculator to design a virtual environment.
- Design a supportable Exchange Server 2013 deployment for a medium-sized organization.
- Design an Exchange Server 2013 deployment for a .medium-complexity organization.

Lab Setup

Estimated Time: 60 minutes

Virtual machines	**20342B-LON-DC1** **20342B-LON-CL1**
User name	Adatum\Administrator
Password	Pa$$w0rd

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

1. On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In the Microsoft Hyper-V Manager, click **20342B-LON-DC1**, and then in the Actions pane click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Log on by using the following credentials:
    a. User name: **Adatum\Administrator**
    b. Password: **Pa$$w0rd**
5. Repeat steps 2 and 3 for **20342B-LON-CL1**.

Exercise 1: Designing a Microsoft Exchange Server 2013 Deployment for a Large Organization

Scenario

A. Datum has a virtualization infrastructure that is based on Hyper-V in Windows Server 2012. You must use the Exchange Server 2013 Mailbox Server Role Requirements Calculator to create a supported solution for 25,000 mailboxes that have 5 GB mailbox quotas. The mailboxes must be highly available, with at least three copies of the mailbox databases. The Client Access servers must also be highly available. A. Datum has hardware load balancers available to deploy the solution. The virtual environment has eight hypervisors with 128 processor cores and 512 GB of memory. Each Global Catalog server will have four virtual processors and 8 GB of memory assigned. Each Mailbox and Client Access server will have eight

MCT USE ONLY. STUDENT USE PROHIBITED

virtual processors assigned. Client Access servers will have 8 GB of memory assigned. Because Exchange native protection will be used, at least four database copies are required. You must determine the configuration of the Exchange Server deployment and how or whether you will recommend using Hyper-V in the solution.

The main tasks for this exercise are as follows:

1. Use the Exchange Mailbox Calculator to create a configuration
2. Verify the processor configuration generated by the Mailbox Role Calculator
3. Verify the memory configuration recommended by the Mailbox Configuration Calculator
4. Decide whether to deploy Exchange using either virtual or physical servers

▶ Task 1: Use the Exchange Mailbox Calculator to create a configuration

1. Sign in to the **LON-CL1** virtual machine with username of **Adatum\Administrator** and password **Pa$$w0rd**.
2. In Microsoft Excel® 2013, open **C:\Files\E2013Calc.xlsm**.
3. In the corresponding fields, change the following entries to the listed values:
   - Server Multi-Role Configuration (MBX+CAS): **No**
   - Server Role Virtualization: **Yes**
   - Number of Mailbox Servers Hosting Active Mailboxes / DAG: **16**
   - Consider Storage Designs Utilizing JBOD (if applicable): **No**
   - Mailbox Server Guest Machines < Processor Cores / Server: **8**
   - Mailbox Server Guest Machines < SPECint2006 Rate Value: **43**
   - Hypervisor CPU Adjustment Factor Value: **10%**
   - Total Number of Tier-1 User Mailboxes / Environment: **25000**
   - Mailbox Size Limit (MB)**: 5120**

▶ Task 2: Verify the processor configuration generated by the Mailbox Role Calculator

1. Switch to the **Role Requirements** tab in Excel 2013.
2. Document the following information.

Core Requirements	
Recommended Minimum Number of Mailbox server Cores	
Recommended Minimum Number of Client Access server Cores	
Recommended Minimum Number of Global catalog Cores	
Total cores needed	

3. How many Client Access servers are needed?

MCT USE ONLY. STUDENT USE PROHIBITED

4. How many global catalog servers are needed?

▶ Task 3: Verify the memory configuration recommended by the Mailbox Configuration Calculator

1. Calculate the following information.

Memory Requirements	
Total Mailbox server memory	
Total Client Access server memory	
Total global catalog memory	
Total memory needed	

▶ Task 4: Decide whether to deploy Exchange using either virtual or physical servers

1. Answer the following questions:
   - Are there enough processor cores to virtualize this Exchange configuration?
   - Is there enough memory to virtualize this Exchange configuration?
   - Will you deploy physical or virtual servers?
2. Close Excel 2013 and do not save E2013Calc.xlsm.

Results: After completing this exercise, you should have designed a Microsoft® Exchange Server 2013 deployment for a large organization.

Exercise 2: Designing an Exchange Server 2013 Deployment for a Medium-Sized Organization

Scenario

A. Datum has a virtualization infrastructure based on Windows Server 2012 and Hyper-V. You must use the Exchange Server 2013 Mailbox Server Role Requirements Calculator to create a supported solution for a small site that is located in Chicago. The site has 2,500 mailboxes with 5 GB mailbox quotas. The mailbox databases do not need to be highly available. The Client Access servers must be highly available. A. Datum has hardware load balancers available to deploy the solution. You must determine the configuration of the Exchange Server deployment and how or whether you will recommend using Hyper-V in the solution. The virtual environment has eight hypervisors with 128 available processor cores and 512 GB of memory. Each Global Catalog server will have four virtual processors and 8 GB of memory assigned. Each Mailbox and Client Access server will have eight virtual processors assigned. Client Access servers will have 8 GB of memory assigned. Because Exchange native protection will be used, at least four database copies are required.

The main tasks for this exercise are as follows:

1. Input information into the Exchange 2013 Mailbox Server Role Requirements Calculator
2. Verify the Processor Configuration generated by the Mailbox Role Calculator
3. Verify the memory configuration recommended by the Mailbox Configuration Calculator

MCT USE ONLY. STUDENT USE PROHIBITED

4. Decide whether to deploy Exchange using either virtual or physical servers

▶ Task 1: Input information into the Exchange 2013 Mailbox Server Role Requirements Calculator

1. Sign in to **LON-CL1** as the domain administrator.
2. In Excel 2013, open C:\Files\E2013Calc.xlsm.
3. In the corresponding fields, change the following entries to the listed values:
   - Server Multi-Role Configuration (MBX+CAS): **No**
   - Server Role Virtualization: **Yes**
   - High Availability Deployment: **No**
   - Number of Mailbox Servers: **4**
   - Mailbox Server Guest Machines < Processor Cores / Server: **4**
   - Mailbox Server Guest Machines < SPECint2006 Rate Value: **20**
   - Hypervisor CPU Adjustment Factor Value: **10%**
   - Total Number of Tier-1 User Mailboxes / Environment: **2500**
   - Mailbox Size Limit (MB)**: 5120**

▶ Task 2: Verify the Processor Configuration generated by the Mailbox Role Calculator

1. Switch to the **Role Requirements** tab in Excel 2013.
2. Document the following information.

Core Requirements	
Recommended Minimum Number of Mailbox server Cores	
Recommended Minimum Number of Client Access server Cores	
Recommended Minimum Number of Global catalog Cores	
Total cores needed	

3. Planning for high availability, what is the minimum number of Client Access servers needed?
4. Planning for high availability, what is the minimum number of global catalog servers needed?
5. Some of the core requirements are lower than the minimum required for high availability. Calculate the following using minimums required for high availability.

Core Requirements	
Recommended Minimum Number	

MCT USE ONLY. STUDENT USE PROHIBITED

Core Requirements	
of Mailbox server Cores	
Recommended Minimum Number of Client Access server Cores	
Recommended Minimum Number of Global catalog Cores	
Total cores needed	

▶ Task 3: Verify the memory configuration recommended by the Mailbox Configuration Calculator

1. Calculate the following information.

Memory Requirements	
Total Mailbox server memory	
Total Client Access server memory	
Total global catalog memory	
Total memory needed	

▶ Task 4: Decide whether to deploy Exchange using either virtual or physical servers

1. Answer the following questions:
   - Are there enough processor cores to virtualize this Exchange configuration?
   - Is there enough memory to virtualize this Exchange configuration?
   - Will you deploy physical or virtual servers?
2. Close Excel 2013 and do not save E2013Calc.xlsm.

Results: After completing this exercise, you should have designed an Exchange Server 2013 deployment for a medium-sized organization.

MCT USE ONLY. STUDENT USE PROHIBITED

Exercise 3: Designing an Exchange Server 2013 Deployment for a Medium Complexity Organization

Scenario

In each site, A. Datum has a virtualization infrastructure that is based on Windows Server 2012 and Hyper-V. You must use the Exchange Server 2013 Mailbox Server Role Requirements Calculator to create a solution with 5,000 mailboxes with 1 GB mailbox quotas. The mailbox databases must be highly available, with at least three copies. The Client Access servers also must be highly available. A. Datum has hardware load balancers available to deploy the solution. You must determine the configuration of the Exchange Server deployment and how or whether you will recommend using Hyper-V in the solution. The virtual environment has six hypervisors which have 72 available processor cores total and 192 GB of memory. Each Global Catalog server will have four virtual processors and 8 GB of memory assigned. Each Mailbox and Client Access server will have eight virtual processors assigned. Client Access servers will have 8 GB of memory assigned. Because Exchange native protection will be used, you require at least four database copies.

The main tasks for this exercise are as follows:

1. Input information into the Exchange 2013 Mailbox Server Role Requirements Calculator
2. Verify the processor configuration generated by the Mailbox Role Calculator
3. Verify the memory configuration recommended by the Mailbox Configuration Calculator
4. Decide whether to deploy Exchange using either virtual or physical servers
5. To prepare for the next module

► Task 1: Input information into the Exchange 2013 Mailbox Server Role Requirements Calculator

1. Sign in to **LON-CL1** as the domain administrator.
2. In Excel 2013, open C:\Files\E2013Calc.xlsm.
3. In the corresponding fields, change the following entries to the listed values:
   - Server Multi-Role Configuration (MBX+CAS): **No**
   - Server Role Virtualization: **Yes**
   - High Availability Deployment: **Yes**
   - Number of Mailbox Servers Hosting Active Mailboxes / DAG: **5**
   - Total Number of HA Database Copy Instances (Includes Active Copy) within DAG: **3**
   - Consider Storage Designs Utilizing JBOD (if applicable): **No**
   - Mailbox Server Guest Machines < Processor Cores / Server: **8**
   - Mailbox Server Guest Machines < SPECint2006 Rate Value: **46**
   - Hypervisor CPU Adjustment Factor Value: **10%**
   - Total Number of Tier-1 User Mailboxes / Environment: **5000**
   - Mailbox Size Limit (MB)**: 1024**

MCT USE ONLY. STUDENT USE PROHIBITED

▶ Task 2: Verify the processor configuration generated by the Mailbox Role Calculator

1. Switch to the **Role Requirements** tab in Excel 2013.
2. Document the following information.

Core Requirements	
Recommended Minimum Number of Mailbox server Cores	
Recommended Minimum Number of Client Access server Cores	
Recommended Minimum Number of Global catalog Cores	
Total cores needed	

3. How many Client Access servers are needed?
4. How many global catalog servers are needed?

▶ Task 3: Verify the memory configuration recommended by the Mailbox Configuration Calculator

1. Calculate the following information.

Memory Requirements	
Total Mailbox server memory	
Total Client Access server memory	
Total global catalog memory	
Total memory needed	

▶ Task 4: Decide whether to deploy Exchange using either virtual or physical servers

1. Answer the following questions:
   - Are there enough processor cores to virtualize this Exchange configuration?
   - Is there enough memory to virtualize this Exchange configuration?
   - Will you deploy physical or virtual servers?
2. Close Excel 2013 and do not save E2013Calc.xlsm.

▶ Task 5: To prepare for the next module

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

1. On the host computer, start Hyper-V Manager.

MCT USE ONLY. STUDENT USE PROHIBITED

2. In the **Virtual Machines** list, right-click **20342B-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 to 3 for **20342B-LON-CL1.**
5. In Hyper-V Manager, click **20342B-LON-DC1**, and in the Actions pane, click **Start**.
6. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
7. Sign in using the following credentials:
   a. User name: **Adatum\Administrator**
   b. Password: **Pa$$w0rd**
8. Repeat steps 5 to 7 for **20342B-LON-MBX1**. When you have successfully signed in to LON-MBX1, repeat steps 5 to 7 for **20342B-LON-CAS1**, **20342B-LON-CL1**, **20342B-LON-CL2, and 20342B-LON-LY1.**

Results: After completing this exercise, you should have designed an Exchange Server 2013 deployment for a medium-sized organization.

Question: In the first configuration, only eight hypervisors were available. What potential issues does this create?

Question: In the final exercise, did you choose to virtualize or to use physical servers? Why?

MCT USE ONLY. STUDENT USE PROHIBITED

Module Review and Takeaways

Best Practice

Become familiar with the server virtualization guidelines provided. These guidelines change from time to time. These changes mean that you should review any updated documentation to so that you can continue to support your organization's virtual and physical servers.

Review Question(s)

Question: What are the biggest factors for deciding to virtualize Exchange Server 2013?

Tools

You can use the tools in the following table to work with Windows PowerShell.

Tool	Description
Exchange Mailbox Server Role Requirements Calculator	An Excel 2013 spreadsheet used to estimate the storage and processor requirements for Mailbox Servers
Exchange Server 2013 Help: Exchange 2013 Virtualization http://go.microsoft.com/fwlink/?LinkId=290687	The latest guidance for virtualizing Exchange Server 2013

MCT USE ONLY. STUDENT USE PROHIBITED

MCT USE ONLY. STUDENT USE PROHIBITED

Module 3

Overview of Exchange Server 2013 Unified Messaging

Contents:

Module Overview **3-1**

Lesson 1: Overview of Telephony Technologies **3-2**

Lesson 2: Unified Messaging in Exchange Server 2013 **3-9**

Lesson 3: Unified Messaging Components **3-14**

Lab: Unified Messaging Overview **3-21**

Module Review and Takeaways **3-23**

Module Overview

Unified Messaging is the general term for the integration of email, voice mail, voice access, and fax services in a user's mailbox. It enables a user to access all types of messages within the same email client, such as Microsoft® Office Outlook® or other email clients.

New features have been added to Unified Messaging in Exchange 2013 Preview. Moreover, Exchange Server 2013 no longer has a dedicated Unified Messaging server role. Instead, Unified Messaging functionalities are part of the redeveloped Client Access server and Mailbox server roles.

This module provides an overview of the various components in a Unified Messaging infrastructure, such as the telephone systems and gateway. It also explains the basic terminology related to Voice over IP (VoIP).

Objectives

After completing this module, students will be able to:

- Describe telephony technologies.
- Describe Unified Messaging in Exchange Server 2013.
- Describe the Unified Messaging components.

MCT USE ONLY. STUDENT USE PROHIBITED

Lesson 1
Overview of Telephony Technologies

Unified Messaging combines voice and email messaging into one location, accessible from a telephone and a computer or a mobile device. Exchange Server 2013 Unified Messaging integrates Exchange Server with different telephony networks and makes the Unified Messaging features available in the user mailbox. This lesson describes basic telephony concepts that you need to understand before you implement Unified Messaging.

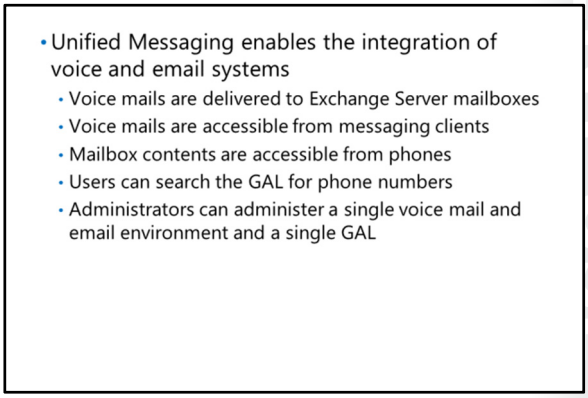
Lesson Objectives

After completing this lesson, you will be able to:

- Describe Unified Messaging.
- Describe the different types of private branch exchange (PBX).
- Provide a summary of telephony terminology.
- Describe the differences between circuit-switched and packet-switched networks.
- Describe VoIP and VoIP protocols.
- Describe the purpose and functionality of a VoIP gateway.

What Is Unified Messaging?

Unified Messaging enables the integration of phone systems and the email functionality in Exchange Server 2013. When you deploy Unified Messaging:



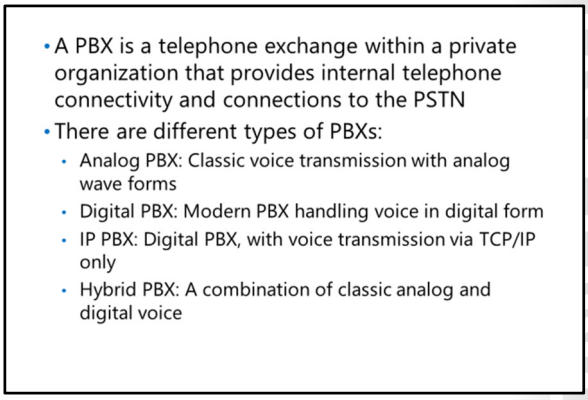
- Unified Messaging delivers voice mails left by telephone callers to user mailboxes on Exchange.
- Users can access the voice mails by using most messaging clients. The Unified Messaging system will transcribe the voice mails into text, or users can listen to recordings of the voice mails.
- Users can access the voice mails, as well as other mailbox contents, from most phones.
- Users can manage the contents of their mailboxes, including meetings and meeting requests, by using a phone.
- Both internal and external users can call the Unified Messaging deployment and search the global address list (GAL) for user phone numbers. When a user locates a necessary phone number, the Exchange 2013 server can place the call automatically to the requested user.
- Administrators can manage a single environment for email and voice messages, and manage a single GAL that is accessible from both messaging and telephone clients.

Since Unified Messaging enables the integration of two disparate systems, there are extra components that you must deploy to implement Unified Messaging. To design and implement these components correctly, you must have some understanding of telephone components.

MCT USE ONLY. STUDENT USE PROHIBITED

What Is a PBX?

A PBX usually is the core device or system that provides telephony and telephone features for residential or organizational use, such as for homes with multiple phones or businesses with thousands of phones. Every PBX connects externally to the public voice network. The following section describes several types of PBXs.



Analog PBX

Analog PBX systems send voice and signaling information, such as the touch tones of dialed phone numbers, as actual analog sound. Analog PBX systems never digitize the sound. To direct the call, the PBX and the phone company's central office listens for the signaling information.

Digital PBX

Digital PBXs encode analog sound into a digital format. Typically, they encode the voice by using a standard industry audio codec, G.711. After digital PBXs encode the sound, they send the digitized voice on a channel by using circuit switching.

The process of circuit switching establishes an end-to-end, open connection, and leaves the channel open for the call's duration and for the call's users only. Some PBX manufacturers have proprietary signaling methods for call setup.

IP PBX

IP PBXs carry voice over data networks. The IP phone contains a Network Interface Card (NIC), and it is part of the network. The phone converts voice into digitized packets, which it then places on the data network. The network sends the voice packets through packet switching, a technique that enables a single network channel to handle multiple calls.

The IP PBX also acts as a gateway between the internal packet-switched network and the external circuit-switched networks that telephone companies use. In this situation, external phone calls arrive at the IP PBX on the normal public phone lines, and the IP PBX converts the phone call to packets sent on the internal IP-based network.

Hybrid PBX

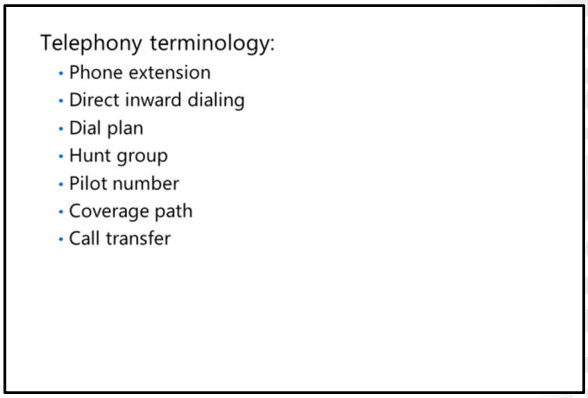
Hybrid PBXs provide both digital and IP PBX capabilities. This hybrid approach enables a customer to run a mixture of digital and IP-based phones. Most modern PBXs are in this hybrid category.

Note:
PABX and EPABX also are terms that refer to PBX. In the telephone industry, PABX is a *private automatic branch exchange*, while EPABX is an *electronic private automatic branch exchange*. These three terms refer to the same type of system.

MCT USE ONLY. STUDENT USE PROHIBITED

Telephony Terminology

Telephony administrators use specialized terminology to describe many of the features and concepts of telephone systems. When deploying Unified Messaging, you need to understand these terms and how they relate to Unified Messaging.



Phone Extension

In most organizations, users have unique extensions for their phone numbers. These extensions are often four or five numbers long. Users within the organization can call other users in the organization by dialing just the user extension rather than the full phone number.

Direct Inward Dialing

A Direct Inward Dialing (DID) phone number is a unique number that an organization assigns to a person. A user with a DID number can receive calls directly from an external phone without having to transfer the call. The DID is a combination of a company-specific phone number and the user's extension. If the organization implements a PBX, the PBX maps DID numbers to internal extensions, so that it can route calls to the correct phone.

Dial Plan

A dial plan consists of the rules that a PBX uses to determine what action to take when it receives a set of dialed numbers. For example, a "9" often triggers call setup to an outside line, so that users can call external phone numbers. When "9" is not the first number, the PBX needs to know how many numbers to collect before taking action. If internal extension numbers are four digits long, it waits for just four numbers before taking action.

Hunt Group

A hunt group is a collection of extensions. In most cases, a hunt group represents a set of identical resources that an application or a group shares. This grouping provides more-efficient access to applications, such as voice mail, an auto attendant, or even a call center. This ensures that callers do not experience a busy signal. Instead, the PBX hunts for an open line to which to connect them.

Pilot Number

A pilot number is the address or label that the PBX uses to identify a hunt group. It is an unused extension, meaning it is not associated with a person or phone.

For example, there may be a specific extension number 3900 for the telesales team, which may be the pilot number for the hunt group of telesales-extension numbers. When a call comes into the 3900 sales number, the PBX recognizes it as a pilot number, and searches for an available line within the sales hunt group. The PBX then delivers the call to an available sales-extension number.

Coverage Path

A PBX uses a set of directions that you configure for each extension, and it tells the PBX where to route unanswered calls and calls that receive busy signals. The set of directions is a *coverage path*. If a DID call arrives at the Unified Messaging server through a user's desktop phone, and the line is busy or not answered within a certain number of rings, the PBX knows to send the call to the pilot number for the hunt group that attaches to the VoIP gateway. The PBX routes the call through the VoIP gateway to the Unified Messaging server, where the caller can record a voice message. The Unified Messaging server sends the voice message to the Unified Messaging user's mailbox.

MCT USE ONLY. STUDENT USE PROHIBITED

Call Transfer

Call transfer is a PBX function that enables users to transfer a phone call to another phone or attendant. Call transfers typically are initiated by using a transfer button on a phone or by using the transfer option on computer-based voice clients. There are two call-transfer modes: announced and unannounced.

During an announced transfer, the call recipient places the caller on hold, and initiates a call to another recipient. If the second recipient accepts the call, the PBX transfers the call. This type of call transfer is also called *attended* or *supervised* transfer. During an unannounced transfer, the PBX transfers the call without checking whether the second recipient will accept the call. This type of call transfer is also called *unattended* or *unsupervised*.

Discussion: Describing Your Telephony Infrastructure

Examine the following questions to describe and discuss the telephony infrastructure in your organization:

Question: What telephony system are you using currently in your environment?

Question: What additional functionality would users in your organization like to see from their telephony system?

- What telephony system are you currently using in your environment?
- What additional functionality would users in your organization like to see from their telephony system?

Circuit-Switched and Packet-Switched Networks

Telephony systems and computer systems usually use different networks to communicate. A telephony system typically uses a circuit-switching network, while a computer system uses a packet-switching network.

- Circuit-switched networks use a dedicated connection between two network devices, such as phones
  - An example of a circuit-switched network is PSTN
- Packet-switched networks divide each data message into packets that are sent to their destination and reassembled
  - The Internet uses packet-switched networks

Circuit-Switched Networks

A circuit-switched network uses a dedicated connection between two network devices. For example, you pick up the telephone receiver and dial a phone number. By answering the call, the recipient completes the circuit. After the two nodes establish a call between them, only these two nodes can use the connection. When one of the nodes ends the call, the connection is removed.

The public switched telephone network (PSTN) uses circuit-switched networks. PSTN connections use the copper medium, as well as fiber-optic cables, microwaves, and satellite links. Although the PSTN can transmit multiple calls across the same transmission medium, each connection is set up as an end-to-end circuit.

MCT USE ONLY. STUDENT USE PROHIBITED

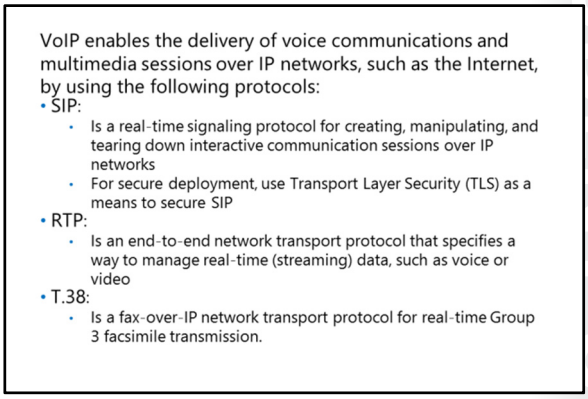
Packet-Switched Networks

Packet switching is a technique that divides a data message into smaller units, or *packets*. The network sends the packets to their destination by the best route available, and then reassembles them at the receiving end.

In packet-switched networks, such as the Internet, hosts route packets to their destination through the most expedient route. A packet-switched network routes packets individually between nodes over data links that other nodes may share. It is possible that not all packets traveling between two hosts travel the same route, even if they are from a single message. This arrangement means that the packets can arrive at different times and out of order. With packet switching, unlike circuit switching, multiple connections to a network's nodes share available bandwidth. This means that the available bandwidth on packet-switched networks may vary from one second to the next. Furthermore, because of the nature of traffic sent on packet-switched networks, these networks are not ideal for time-sensitive traffic, such as voice or video streaming. These types of traffic depend on traffic arriving in order and in a timely manner, but a packet-switched network cannot guarantee this. The protocols that transmit voice and video over packet-switched networks are designed to alleviate these issues, as much as possible, but the underlying network remains critical in determining the quality of service.

What Is VoIP?

VoIP is a technology that enables an IP-based network to act as the transmission medium for telephone calls. It sends voice data in IP packets rather than by circuit-switched telephone lines. Translating a call from a circuit-switched network to a packet-switched network is complicated because the underlying network connections are so different. Packets sent on a packet-switched network may arrive out of order or in bursts, which would result in garbled communication unless the network reassembles the packets in the right order and sends them at the right speed. VoIP real-time protocols help achieve the level of reliability and voice quality that traditional circuit-switched telephone calls provide. These protocols protect against packet loss, delay, or variability in delay.



There are a number of voice-related, IP-based protocols, and a Unified Messaging environment with Exchange Server 2013 uses the following:

- *Session Initiation Protocol (SIP)*. SIP is a real-time signaling protocol that creates, manipulates, and tears down interactive communication sessions on an IP network. You can use SIP in conjunction with Transport Layer Security (TLS) to provide security. Exchange Server Unified Messaging uses SIP mapped over Transmission Control Protocol (TCP), and supports TLS for secured SIP environments.

  SIP clients, such as IP/VoIP gateways and IP/PBXs, can use TCP port 5060 or port 5061 (for Secure SIP) to connect to SIP servers.

- *Real-Time Transport Protocol (RTP).* RTP is for voice transport between the IP gateway and the Unified Messaging server. RTP provides high-quality, real-time, streaming voice delivery. One of the issues with sending voice messages over an IP network is that voice requires real-time transport, with specific quality requirements, to ensure that the voice sounds normal. If the protocol uses large packets, listeners must wait for the entire packet to arrive before they can respond. Any delay in packet delivery can produce undesirable periods of midstream silence, and packet loss can cause voice garbling.

MCT USE ONLY. STUDENT USE PROHIBITED

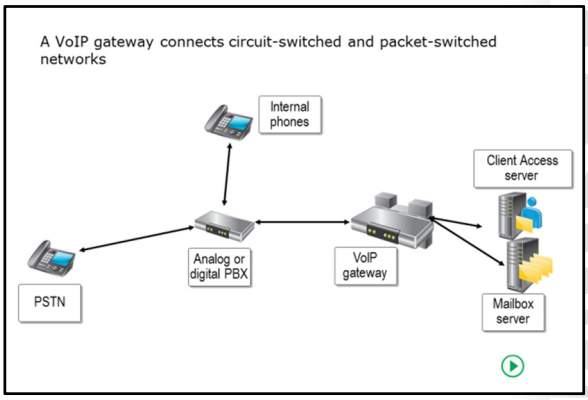
Additional Reading: For more information, refer to Request for Comment (RFC) 3550, which updates RFC 1889, and describes RTP); and RFC 3261, which updates RFC 2543, and describes SIP. You can access all RFCs at http://go.microsoft.com/fwlink/?LinkId=290684.

- *Real-Time Facsimile or T.38.* Real-Time Facsimile or T.38 is an Internet fax-transport protocol. T.38 sets procedures for fax transmission when a portion of the path includes an IP network. The Unified Messaging system uses it to relay a fax that a user sends, through a voice line across an IP network, in real time.

What Is a VoIP Gateway?

A VoIP gateway is a hardware device or product that converts traditional phone-system or circuit-switching protocols into data-networking or packet-switched protocols.

Exchange 2013 servers that are running the Unified Messaging components can connect only to packet-switched data networks. This requirement means that organizations with a traditional PBX must deploy a VoIP gateway to communicate between the PBX and the Exchange 2013 servers. The PBX also provides access to the PSTN for internal phone users and for Unified Messaging.



Options for Implementing VoIP Gateways

There are several different options available for implementing VoIP gateways, depending on your current telephone and PBX deployment. The most common options are:

- Deploying Unified Messaging with an analog or digital PBX. An analog or digital PBX can communicate only on circuit-switched networks, so a VoIP gateway is required when using this type of PBX. The VoIP gateway translates all communication between the Exchange 2013 servers and the PBX. The PBX connects the analog or digital phones, and also provides connectivity to the PSTN.
- Deploying Unified Messaging with an IP or hybrid PBX. An IP or hybrid PBX has one interface that connects to a circuit-switched network, and one interface that connects to a packet switched network. In this case, the PBX operates as a VoIP gateway, so no dedicated VoIP gateway is required. You can use both analog or digital phones, as well as VoIP phones, to connect to the PBX. The circuit switched interface provides connectivity to the PSTN.
- Deploying Unified Messaging with Microsoft Lync® Server. A Lync Server also can operate as a VoIP gateway for the Exchange 2013 servers that are running Unified Messaging. Like Exchange servers, Lync servers can communicate only on packet-switched networks, so no other VoIP gateways are necessary for the Exchange Server. In this deployment, the Lync server provides telephone services for Lync clients and other VoIP phones. The Lync server also must be able to communicate with the PSTN, which may require a VoIP gateway between the Lync Server and the PSTN.

Note: There are several options for connecting a Lync Server to the PSTN. The next module will provide additional details on these options.

MCT USE ONLY. STUDENT USE PROHIBITED

MCT USE ONLY. STUDENT USE PROHIBITED

Lesson 2
Unified Messaging in Exchange Server 2013

Unified Messaging enables users to receive email, voice, and fax services in their Exchange Server Inbox, and allows users to access mailbox contents by phone. This simplifies the experience for users, because they must access and manage only one location for all message types. This also provides more functionality for users, because they can use traditional messaging clients to access voice or fax messages, and they can use telephone technology to access email messages. Unified Messaging also simplifies administrators' workloads because they must manage this data in one location only. This lesson introduces the Exchange 2013 Unified Messaging features.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe Unified Messaging features.
- Describe the Exchange Server 2013 Unified Messaging features.
- Describe the user-configuration options for Unified Messaging.

Unified Messaging Features

Exchange Unified Messaging provides the following core features:

Unified Messaging consolidates voice mail, email, and fax services into one store, and features include:

- Access to voice mail in user mailboxes
- Play on Phone
- Call answering
- Call Answering Rules
- Outlook Voice Access
- Message Waiting Indicator
- Voice Mail Preview
- Missed call and voice mail notification using SMS
- Protected Voice Mail
- Voice mail form
- User configuration

- Access to voice mail in user mailboxes. UM-enabled users can access their voice mail from mobile phones, clients, and through Outlook Web App. Users can play voice-mail messages from the reading pane by using the integrated Windows Media® Player in Outlook or Windows® Internet Explorer®.
- Play on Phone. UM-enabled users also can play their messages by using any normal phone to dial into Exchange 2013 or by using Microsoft Lync 2013. This arrangement also prevents others from listening to confidential voice mails if the computer only has external speakers.
- Call answering. This feature supports playing personal greetings, recording messages, and answering incoming calls on behalf of other users. Users can submit the results of these actions for delivery to user inboxes as an email message.
- Call Answering Rules. UM-enabled users can organize how the phone system handles their incoming calls. This feature is similar to Inbox rules, which users can apply to normal email messages. No call answering rules are activated by default.

If the Exchange 2013 Unified Messaging service answers a call, it prompts the caller to leave a voice message similar to a normal answering machine. With Call Answering Rules, the user can customize the experience for callers when they connect to leave a message.

- Outlook Voice Access. UM-enabled users have two options for Outlook Voice Access: the Telephone User Interface (TUI) and the Voice User Interface (VUI). This feature facilitates internal and external access by using phone systems, and enable users to :

MCT USE ONLY. STUDENT USE PROHIBITED

  - Access voice mail.
  - Listen to, forward, or reply to email messages.
  - Listen to calendar information.
  - Access or dial contacts who are stored in the global address list (GAL) or a group in their Contacts folder.
  - Accept or cancel meeting requests.
  - Set a voice-mail message to let callers know the called party is away.
  - Set user-security preferences and personal options.

- Voice Mail Preview. In Exchange 2013, the Unified Messaging feature uses Automatic Speech Recognition (ASR) on new voice-mail messages. When users receive voice messages, the messages contain both a recording and voice-mail preview text, which the system creates from the voice recording.
- Message Waiting Indicator. The Message Waiting Indicator is any mechanism that indicates the existence of new Unified Messaging messages. Unified Messaging enables you to implement this in several ways, depending on which client you are using. In Outlook, the Messaging Waiting Indicator displays as an unread voice-mail message. Lync clients can display Message Waiting Indicators when a when a new voice mail has been left in the user mailbox. If the client is a phone, the Message Waiting Indicator may be a light on the phone.
- Missed call and voice-mail notifications by using SMS. If users are members of a hosted or consumer dial plan, and they configure their voice-mail settings, including their mobile phone number, with call forwarding, they can receive notification about missed calls and newly arrived voice mail on their cell phones as short message service (SMS) text messages.
- Protected Voice Mail. This extended feature is provided in conjunction with Active Directory® Rights Management Service (AD RMS), and it enables the secure storage of voice-mail messages. This restricts the forwarding, copying, or extracting of voice file from email.
- Voice mail form. The Outlook 2010, Outlook 2013, and Outlook Web App form for voice mail resembles the default email form. Users can perform several actions, such as playing, stopping, or pausing voice messages, playing voice messages on a telephone, and adding and editing notes.
- User configuration. UM-enabled users can perform several voice-mail options by using Outlook Web App. Examples include setting telephone-access numbers or voice-mail *Play on Phone* numbers, or resetting a personal identification number (PIN) for voice-mail access.

Unified Messaging Features in Exchange 2013

Exchange 2013 has added several new features to Unified Messaging.

New features in Unified Messaging Exchange 2013:
- Changes in Voice Architecture
- IPv6 Support
- Improvements in Voice Mail Preview
- Enhanced Caller ID Support
- Group addressing using Outlook Voice Access

Changes to the Server Architecture

In Exchange Server 2007 and Exchange Server 2010, to provide Unified Messaging functionality, you had to deploy a dedicated Unified Messaging server role. In Exchange Server 2013, there is no dedicated Unified Messaging server role. Rather, the Unified Messaging functionality is distributed between the Client Access server role and the

MCT USE ONLY. STUDENT USE PROHIBITED

Mailbox server role. Exchange 2013 has two core components to handle Exchange Unified Messaging functionality:

- Microsoft Exchange Unified Messaging Call Router Service. The Client Access server includes this service, which handles signaling of traffic and forwards processing to the responsible Mailbox Server.
- Microsoft Exchange Unified Messaging Service. The Mailbox server includes this service, which enables Microsoft Exchange Unified Messaging features, and allows Microsoft Exchange to store voice and fax messages and provides users with telephone access to their email.

IPv6 Support

In Exchange Server 2007 and Exchange Server 2010, the Unified Messaging server role only supported IPv4. In Exchange Server 2013, the Unified Messaging architecture now requires the Unified Communications Managed application programming interface (API), known as UCMA v4.0. UMCA v4.0 supports both IPv4 and IPv6, so all Unified Messaging components and services fully support IPv6 networks.

Voice Mail Preview

Exchange Server 2013 includes the following enhancements for Voice Mail Preview:

- Audio normalization. Before Exchange Unified Messaging compresses an audio signal, it normalizes it, which means that Exchange 2013 improves the audio signal processing so that the resulting peak amplitude matches the best target wave form.
- Speech recognition. If you allow sharing of voice-mail speech-recognition results, Unified Messaging can use the results to add words and phrases to the speech engine. You can enable this by setting the **VoiceMailAnalysisEnabled** parameter to **$true** on the **Set-UMMailbox** cmdlet or by setting the **AllowVoiceMailAnalysis** parameter to **$true** on the **Set-UMMailboxPolicy** cmdlet.
- Voice Mail Preview confidence. In Exchange 2013, the confidence calculation is more accurate. It provides a score that represents the accuracy of the transcribed message.
- Filtering. Unified Messaging detects and filters offensive words, and caches and stores the results in the user's mailbox.
- Hiding the text preview. If a confidence score is below a defined threshold, Exchange hides the preview text. The voice-mail preview contains text stating that the confidence of the voice mail was too low for results to be displayed.
- Transcription performance. Speech to text is an intensive central processing unit (CPU) operation, and it requires twice the processing power of standard audio-file processing. If processing of voice mail takes too long, Exchange 2013 CPU throttling stops the preview processing.
- Color schemes. This feature has been removed in Exchange 2013 Preview for Outlook Web App and Outlook, due to confusion that resulted from the former color scheme that was used to indicate low, medium, and high voice-mail confidence.

Enhanced Caller ID Support

In Exchange Server 2007 and Exchange Server 2010, a Unified Messaging server that received a call used caller ID to locate the caller's possible identity in Active Directory Domain Services (AD DS) and in the user's personal contacts folder. Exchange 2013 Unified Messaging can search the user's other Exchange and personal contact folders, and enable contact aggregation from external social networks. This feature means that contacts from external social networks are available in the contact folder of the user's mailbox and in other contact folders.

MCT USE ONLY. STUDENT USE PROHIBITED

Group Addressing Using Outlook Voice Access

Exchange 2013 adds an additional group-based forwarding option. When UM-enabled users sign in to their mailbox by using Outlook Voice Access, they can send email and voice messages to users in a group stored in their personal Contacts list.

Call Answering Rules

Call answering rules, also known as Personal Auto Attendants, allow users to create and customize rules to enhance the experience that callers have when the system answers their calls. For example, the call-answering rules can include features such as special greetings by contact or time of day.

- Call answering rules provide users with more control over their personal call answering system, with features such as special greetings by contact or time of day
- Call answering rules include:
  - Conditions: If it evaluates to true, then run
  - A greeting and menu: Collect caller's menu selection
  - Actions: Transfer, "find me", or leave message
- You can configure call answering rules in Outlook Web App or Outlook 2010

Using call answering rules, the caller can decide to:

- Leave a voice message for the Unified Messaging-enabled user.
- Transfer to an alternate contact of the Unified Messaging-enabled user.
- Transfer to an alternate contact's voice mail.
- Transfer to other phone numbers that the Unified Messaging-enabled user configures.
- Use the Find-Me feature, or locate the Unified Messaging-enabled user through a supervised transfer.

Call-answering rules consist of conditions, a greeting and menu, and actions. You can configure call-answering rules in Outlook Web App or Outlook 2010 or newer.

Conditions

The following conditions are available:

- If the caller is: calling from a phone number, this specific contact, or in my contacts folder.
- If it is during this period: working hours or nonworking hours to a specific time defined.
- If the user's schedule shows a status of: free, tentative, busy, or away.
- If you turn on automatic replies, such as when you turn on an automatic Out of Office message.

Greeting and Menu

Greeting and Menu is the area where the caller can take specific actions that users predefine. For example, after hearing a greeting that you recorded previously, you can provide a prompt so that the caller can dial you at home.

Actions

Actions define the tasks that occur when callers choose specific menu selections. You can select the following actions:

- Find me at the following numbers: Defines a recording text and the number key to press to transfer, and enables you to call two phone numbers for a specific time.
- Transfer the call to: Defines a recording text, the number key to press to transfer, and either a phone number or a contact; or it indicates that the call should transfer directly to voice mail.
- Leave a voice message. Transfers the caller to voice mail directly.

MCT USE ONLY. STUDENT USE PROHIBITED

Demonstration: Configuring Client Options for Unified Messaging

In this demonstration, you will see some of the configuration options available to users when you enable them for Unified Messaging.

Demonstration Steps

1. On LON-CAS1, connect to **https://lon-cas1.adatum.com/owa**.
2. Sign in as **Adatum\Allie** using the password **Pa$$w0rd**.
3. Review the email welcoming Allie to Unified Messaging.
4. In the Outlook Web App options, on the **voice mail** page, review the options other than the call answer rules options.
5. Create a new call answering rule by using the following settings:
    a. Name: **OOF Rule**
    b. Configure the rule to apply when a user reaches Allie's voice mail and Allie is away.
    c. Configure the option to call Allie's mobile number at 12229998888 if the caller selects 1, or transfer to Allie's assistant at 20022 if the caller presses 2.

MCT USE ONLY. STUDENT USE PROHIBITED

Lesson 3
Unified Messaging Components

To configure Unified Messaging in Exchange Server 2013, you first need to understand how Exchange Server 2013 implements Unified Messaging. The Unified Messaging architecture has been changed from previous Exchange Server versions.

This lesson describes the basic Exchange Server 2013 Unified Messaging components and how they interact to provide Unified Messaging services.

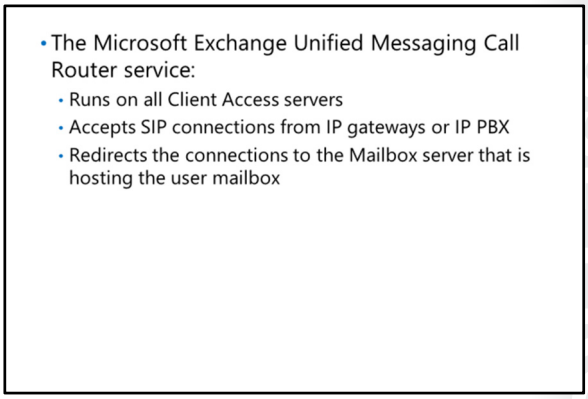
Lesson Objectives

After completing this lesson, you will be able to:

- Describe the functionality of the Microsoft Unified Messaging call router service.
- Describe the functionality of the Microsoft Exchange Unified Messaging service.
- Describe the communications flow in a Unified Messaging environment.
- Identify the ports used by Exchange Server 2013 for Unified Messaging services.
- Describe the components required to configure the integration of Unified Messaging and a telephone system.
- Describe the functionality of a Unified Messaging mailbox policy.
- Describe the Unified Messaging clients.
- Describe the functionality of an Auto Attendant.

Microsoft Unified Messaging Call Router Service

In Exchange Server 2013, the Unified Messaging functionality is distributed across the Client Access and Mailbox server roles. The Microsoft Exchange Unified Messaging Call Router service provides the Unified Messaging functionality on the Client Access server, and it runs on each of the organization's Client Access servers.



The Microsoft Exchange Unified Messaging Call Router service implements routing logic and SIP redirect for Unified Messaging. When a Client Access server receives a SIP INVITE for an incoming call, the Microsoft Exchange Unified Messaging Call Router service redirects the incoming call to the Mailbox server in which the user mailbox is located. Although the Client Access server acts as a SIP redirector to identify the database where the user's mailbox is located, it only handles SIP requests from IP gateways or IP PBXs. It does not receive any media traffic. When you deploy Exchange 2013 and Unified Messaging, you have to configure your VoIP gateways or IP PBXs to point to the Client Access servers, so that incoming calls are routed correctly for Unified Messaging.

MCT USE ONLY. STUDENT USE PROHIBITED

Microsoft Exchange Unified Messaging Service

The Microsoft Exchange Unified Messaging Service runs on all Exchange Server 2013 Mailbox servers, and provides most of the functionality formerly provided by the dedicated Unified Messaging server role. The Microsoft Exchange Unified Messaging Service is the SIP-processing component in Exchange 2013, and it handles media traffic, and processes incoming calls.

The Microsoft Exchange Unified Messaging Service:

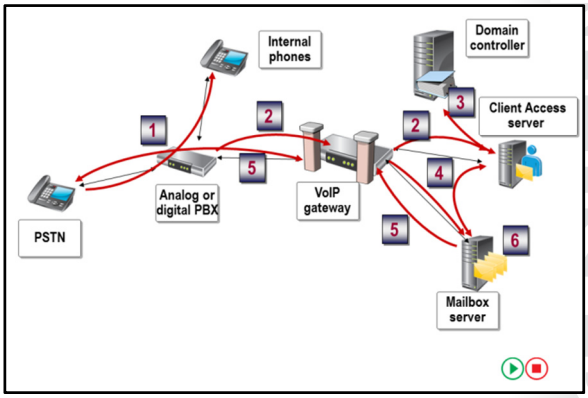
- Runs on all Mailbox servers
- Accepts redirected calls from the Client Access server
- Establishes RTP or SRTP connections with IP gateways or IP-PBX
- Provides call answering and voice mail functionality

When the Microsoft Exchange Unified Messaging Call Router Service redirects a call on a Client Access server to the Mailbox server, it creates an RTP or secure real-time transport protocol (SRTP) connection between the VoIP gateway, or IP PBX and the Mailbox server. Media traffic that uses RTP or SRTP only passes between the Mailbox server and SIP peers such as VoIP gateways or IP PBXs, and not to the Client Access server.

After it establishes the media channel, the Microsoft Exchange Unified Messaging service on the Mailbox server plays the user's voice-mail greeting, processes call-answering rules for the user, and invites the caller to leave a voice message. The Mailbox server then records the voice message, creates a transcription of the message, and deposits it in the user's mailbox.

Overview of Unified Messaging Communications

The following steps describe the communication flow for incoming phone calls when an organization implements Exchange Server 2013 Unified Messaging by using a traditional PBX:



1. A caller dials a user's number in the organization. This caller could be inside or outside the organization. Exchange Server routes the call to the PBX, which uses the call recipient's extension number to route the call to the appropriate desk phone, which then rings. If the recipient does not answer the call, the PBX checks its configuration to see where to route the unanswered call. In this case, the PBX routes the unanswered calls for this phone to the number associated with the VoIP gateway.
2. The VoIP gateway converts the circuit-switched protocols to packet-switched protocols. It uses the information about the Exchange Server Unified Messaging environment, which you configure during the VoIP gateway installation, to route the call to a Client Access server by using SIP. The Client Access server receives the now VoIP-based, packet-switched SIP call.
3. The Client Access server contacts AD DS to retrieve the recipient information. This AD DS lookup occurs by using the combination of dial plan plus extension number, which provides a unique identifier for each mailbox.
4. The Client Access server uses this information to redirect the call to the Mailbox server that is hosting the active mailbox database that contains the user mailbox.

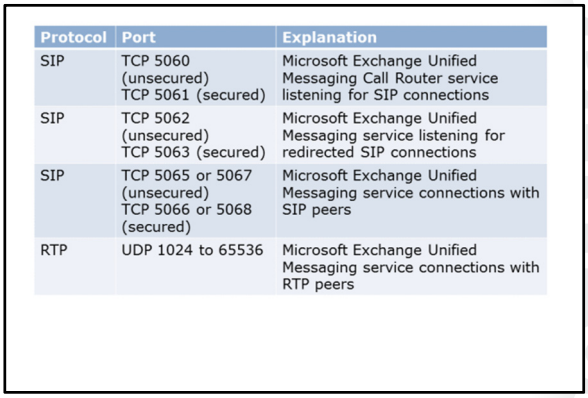
MCT USE ONLY. STUDENT USE PROHIBITED

5. The VoIP gateway connects to the Mailbox server using SIP to set up a call. The Mailbox server retrieves the user's personal greeting, and plays the personal greeting for the caller using the RTP or SRTP.
6. If the caller decides to leave a message, the Mailbox server records the voice mail. The Mailbox server packages the voice mail in to an email message, and sends the message to the user mailbox. The message is accessible to the Unified Messaging subscriber through Outlook Voice Access, Outlook, Outlook Web App, or Exchange ActiveSync®.

These steps describe the communication flow when Exchange Server 2013 Unified Messaging answers a call. The process is similar when you use other systems, such as Outlook Voice Access or auto attendant access. For example, when using Outlook Voice Access, the user calls a number that you configure so that the PBX forwards the call automatically to the VoIP gateway. The gateway then forwards the call to Client Access server, which checks AD DS for the user mailbox location. It then redirects the call to the appropriate Mailbox server, which provides access to the user mailbox through the VoIP gateway. When you use an auto attendant, the PBX forwards the phone number through the VoIP gateway to the Client Access server, which redirects the call to a Mailbox server. The Mailbox server then responds to the call, providing the requested information from the GAL.

Microsoft Exchange Unified Messaging Ports

The Unified Messaging services on the Client Access server and Mailbox servers use specific ports to communicate with other Unified Messaging components. The Microsoft Exchange Unified Messaging Call Router service found on a Client Access server listens for SIP on either TCP or mutual TLS. The Microsoft Exchange Unified Messaging on Mailbox servers listens for redirected SIP connections from the Microsoft Exchange Unified Messaging Call Router service. Furthermore, the Microsoft Exchange Unified Messaging service on the Mailbox server also uses specific ports to connect to SIP peers, and to connect with RTP clients.



Protocol	Port	Explanation
SIP	TCP 5060 (unsecured) TCP 5061 (secured)	Microsoft Exchange Unified Messaging Call Router service listening for SIP connections
SIP	TCP 5062 (unsecured) TCP 5063 (secured)	Microsoft Exchange Unified Messaging service listening for redirected SIP connections
SIP	TCP 5065 or 5067 (unsecured) TCP 5066 or 5068 (secured)	Microsoft Exchange Unified Messaging service connections with SIP peers
RTP	UDP 1024 to 65536	Microsoft Exchange Unified Messaging service connections with RTP peers

The following table summarizes the ports that Exchange 2013 Unified Messaging uses.

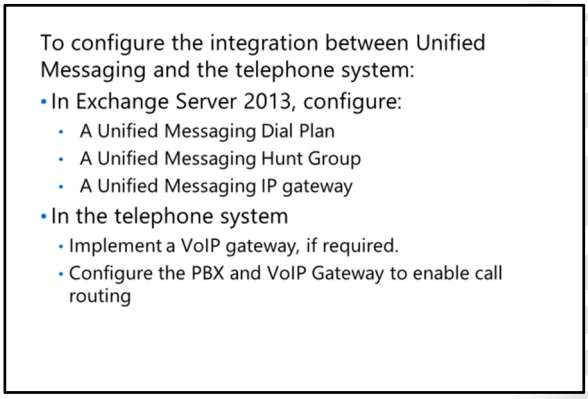
Protocol	Port	Explanation
SIP	TCP 5060 (unsecured) TCP 5061 (secured with TLS)	Microsoft Exchange Unified Messaging Call Router service listens for connections from IP gateways or IP PBXs. You can change ports by using the using the Exchange Management Shell.
SIP	TCP 5062 (unsecured) TCP 5063 (secured with TLS)	Microsoft Exchange Unified Messaging service listens for SIP redirects from Client Access servers.
SIP	TCP 5065 and 5067 (unsecured) TCP 5066 and 5068 (secured with TLS)	Microsoft Exchange Unified Messaging service connections with SIP peers use these ports. If you set the service startup mode to dual, it uses both ports. Only 5065 and 5066 are used if you set the startup mode to TCP or TLS.

MCT USE ONLY. STUDENT USE PROHIBITED

Protocol	Port	Explanation
RTP	UDP ports between 1024 and 65536	Microsoft Exchange Unified Messaging service connections with RTP clients, such as VoIP gateways and IP-PBX, use these ports.

Configuring the Integration Between Unified Messaging and Telephone System

In order to implement Unified Messaging, you must configure the integration between the Exchange 2013 servers that are running Unified Messaging and the telephone system.



Configuring the Unified Messaging Environment

To configure the integration of Unified Messaging and a telephone system, you need to configure the following components:

- Unified Messaging dial plan. The Unified Messaging dial plan is an Active Directory container object that is a logical representation of a telephony PBX dial plan. The Unified Messaging dial-plan settings, such as the telephony extension-numbering plan, must match the telephony dial plan. Within Unified Messaging, the dial plan, plus the extension number, provides the unique identifier for each UM-enabled user. The dial plan also controls the numbering scheme and the outbound dialing plan.
- Unified Messaging Hunt Group. The Unified Messaging hunt group is a logical representation within Exchange Server 2013 of an existing PBX or IP PBX hunt group. Unified Messaging hunt groups locate the PBX hunt group from which the incoming call was received. A pilot number that you specify for a hunt group in the PBX also must be specified within the Unified Messaging hunt group. The pilot number enables the Client Access server to associate the call with the correct dial plan so that it can route the call correctly.
- Unified Messaging IP Gateway. The Unified Messaging IP gateway is an Active Directory container object that logically represents a physical IP gateway device. The Unified Messaging IP gateway can represent either a VoIP gateway or an IP-PBX. The Unified Messaging IP gateway contains one or more Unified Messaging hunt-group objects and other Unified Messaging IP gateway-configuration settings, including the actual IP gateway object. The combination of the IP gateway object and a Unified Messaging hunt-group object establishes a logical link between an IP gateway hardware device and a Unified Messaging dial plan.

Configuring the Telephone Environment

You also must configure components in the telephone environment to integrate with Exchange Server 2013 Unified Messaging, and the configuration options depend on the telephone environment, which can include:

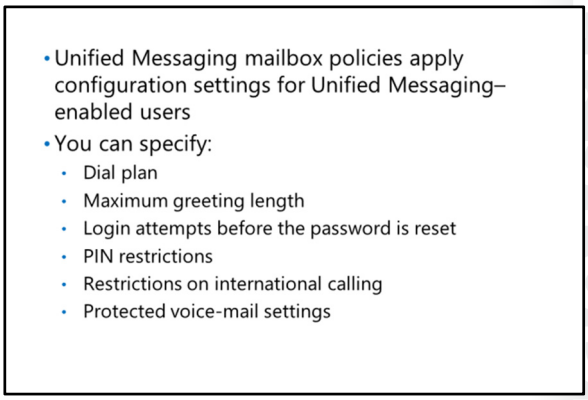
- Traditional PBX. If your organization is using an analog or digital PBX, you will need to deploy a VoIP gateway before you can deploy Unified Messaging. You then must configure the PBX to enable call routing to the VoIP gateway, and configure the VoIP gateway to enable call routing between the PBX and the Exchange 2013 servers.

MCT USE ONLY. STUDENT USE PROHIBITED

- IP-PBX. If your organization is using an IP-PBX, you must configure the IP-PBX to enable call routing to the Exchange 2013 servers.

Note: To ensure compatibility between a telephone system and Unified Messaging, you must ensure that all components are compatible with Exchange 2013 Unified Messaging, and that configuration notes are available for configuring the telephone components. See http://go.microsoft.com/fwlink/?LinkId=290685 to access a list of the Configuration Notes for Supported VoIP Gateways, IP PBXs, and PBXs.

Unified Messaging Mailbox Policies



UM mailbox policies enable you to configure the user experience or security settings to UM-enabled mailboxes. By implementing UM mailbox policies, you can define these settings, and then apply them to all users, or a subset of users, in your organization.

UM mailbox policies apply and standardize Unified Messaging configuration settings for UM-enabled users. You can create Unified Messaging mailbox policies, and then add the policy to UM-enabled mailboxes to apply a common set of policies or security settings. Unified Messaging mailbox policies are required before you can enable users to use Unified Messaging.

When you create a dial plan, a single, default Unified Messaging mailbox policy is created for it. However, you can create additional Unified Messaging mailbox policies based on your organization's needs. When you create a Unified Messaging mailbox policy, you can configure a wide variety of settings, including the following:

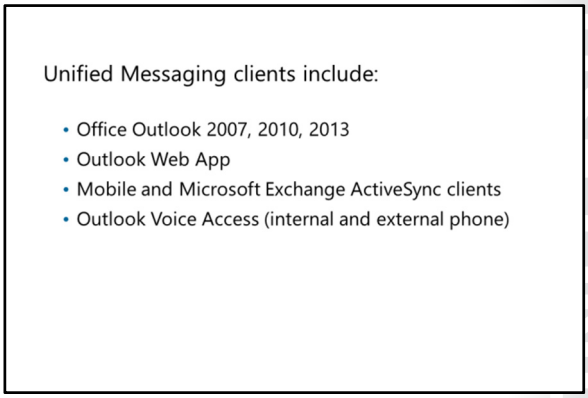
- Dial plan (required)
- Maximum greeting length
- Number of unsuccessful login attempts before it resets the password
- Minimum number of digits that a PIN requires
- Number of days until users must create a new PIN
- Number of previous passwords that it does not allow
- Restrictions on in-country/region or international calling
- Protected voice-mail settings

Each Unified Messaging-enabled user's mailbox must link only to one Unified Messaging mailbox policy.

MCT USE ONLY. STUDENT USE PROHIBITED

Unified Messaging Clients

You can use several different types of clients to access Unified Messaging features in Exchange 2013.

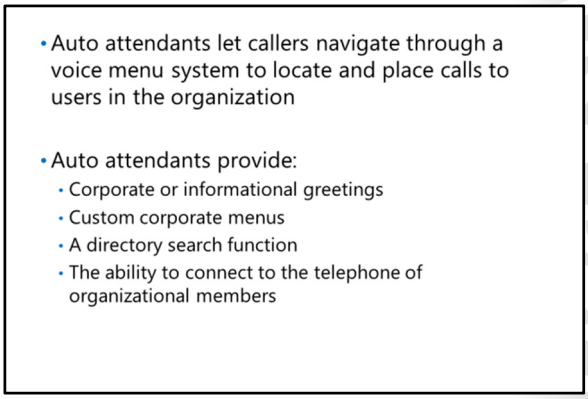


In addition to standard Microsoft mail clients, such as Microsoft Outlook, users also can access Unified Messaging features by using Outlook Web App, on a mobile phone, and directly by using a traditional desk-top phone:

- Microsoft Outlook client. When you configure Unified Messaging, Exchange Server delivers voice messages to the user's mailbox. The user can read a transcript of that message in the Outlook client, and also play an audio recording of the message. Users can play voice messages through media-player integration or download them.
- Outlook Web App. Users also can use Outlook Web App to view voice messages, play them through media-player integration, or download them.
- Mobile devices and Microsoft Exchange ActiveSync clients. Users can play voice messages through media-player integration for mobile phones.
- Outlook Voice Access. Users can use any phone to access their mailbox through Outlook Voice Access. Users can dial the Outlook Voice Access number, and then enter their PIN to access their mailbox. Users can listen to voice messages in their mailbox, and listen and respond to emails and meeting requests.

Unified Messaging Auto Attendants

A Unified Messaging auto attendant is an optional component in a Unified Messaging deployment. It creates a voice-menu system that enables external and internal callers to navigate through voice menus to locate and place, or transfer, calls to company users or organizational departments.



When anonymous or unauthenticated users call an external business telephone number, or when internal callers call a specified extension number, voice prompts help them place a call to a user, or locate and call a user.

The Unified Messaging auto attendant uses a series of WAV files that callers hear instead of a human operator. The Unified Messaging auto attendant lets callers navigate the menu system, place calls, or locate users by using dual-tone multi-frequency (DTMF) or voice inputs.

A Unified Messaging auto attendant provides:

- Corporate or informational greetings, such as business hours or directions to a location.
- Custom corporate menus that you can customize to have more than one level.
- A directory search function that enables callers to search the organization's name directory.

MCT USE ONLY. STUDENT USE PROHIBITED

- The ability for callers to connect to the telephone of, or leave a message for, organizational members.

MCT USE ONLY. STUDENT USE PROHIBITED

Lab: Unified Messaging Overview

Scenario

A. Datum Ltd. is planning to deploy Exchange Server 2013 Unified Messaging to integrate with the Lync Server 2013 deployment. Before starting the design and deployment, it is important that you understand the Unified Messaging terminology and how the different Unified Messaging components relate to each other.

Objectives

After completing this lab, you will be able to identify Unified Messaging components and their interrelationships.

Estimated Time: 20 min

No virtual machines are required for this lab.

Exercise 1: Identifying Unified Messaging Components

Scenario

Review the following blank diagram, which depicts a Unified Messaging deployment. The diagram includes blank labels for the different Unified Messaging components and labels that identify relationships between components. You need to identify the different Unified Messaging components, and describe the relationship between the different components.

The main tasks for this exercise are as follows:

1. Matching the Unified Messaging terminology

▶ Task 1: Matching the Unified Messaging terminology

- Review the following list of terms, and then match them to the correct definition or description.

Term	Description
A. VoIP gateway	_____ Defines user extensions.
B. Dial plan	_____ Can be configured by each user.
C. Auto Attendant	_____ A device that can be used with VoIP phones
D. Call answering	_____ One call to this number could reach many phones
E. Hunt group	_____ A device that would be deployed between an analog PBX and an Exchange 2013 server.
F. Unified Messaging IP Gateway	_____ Identifies the first hop when Exchange 2013 servers need to communicate with the PSTN
G. IP-PBX	_____ A device that cannot communicate directly with Exchange 2013 servers.
H. UM Mailbox Policy	_____ Identifies a group of telephone users

MCT USE ONLY. STUDENT USE PROHIBITED

Term	Description
I. Pilot number	_____ Provides a searchable phone list
J. Analog PBX	_____ Defines some of the user experience with Unified Messaging.

Results: After completing this exercise, you should be able to identify the main Unified Messaging components.

Question: Name all of the new Unified Messaging Voice mail preview features in Exchange Server 2013.

Question: Describe the changes in Exchange Server Unified Messaging architecture.

MCT USE ONLY. STUDENT USE PROHIBITED

Module Review and Takeaways

Review Question(s)

Question: If your company has implemented Lync Server 2013 and connected Lync Server to the PSTN, do you need an additional IP PBX for Exchange Server 2013 Unified Messaging?

Question: You want to provide outside callers with an automated system for searching your organization's directory for user phone numbers. What Unified Messaging component do you need to implement?

MCT USE ONLY. STUDENT USE PROHIBITED

MCT USE ONLY. STUDENT USE PROHIBITED

Module 4

Designing and Implementing Exchange Server 2013 Unified Messaging

Contents:

Module Overview 4-1

Lesson 1: Designing a Unified Messaging Deployment 4-2

Lesson 2: Deploying and Configuring Unified Messaging Components 4-13

Lesson 3: Designing and Implementing Exchange Server 2013 UM Integration with Lync Server 2013 4-22

Lab: Designing and Implementing Exchange Server 2013 Unified Messaging 4-29

Module Review and Takeaways 4-39

Module Overview

Unified Messaging in Exchange Server 2013 combines email messaging and voice messaging into a single infrastructure. Users can access the email and voice-messaging pieces of Unified Messaging (either from within a network or remotely. After you fully deploy Unified Messaging in Exchange Server 2013, your users can access their messages easily by using Outlook® Voice Access from any device. Users also can use their mobile device, Microsoft® Lync® client, or Lync Phone Edition.

This module provides an overview of the entire design and deployment process for the Unified Messaging-related components in Exchange Server 2013, as well as the associated components, such as the telephone systems and Voice over IP (VoIP) gateways. This module also explains how you can integrate Exchange Server 2013 with Lync Server 2013.

Objectives

After completing this module, students will be able to:

- Design a Unified Messaging deployment.
- Deploy and configure Unified Messaging in Exchange Server 2013.
- Design and implement Exchange Server 2013 Unified Messaging integration with Lync Server 2013.

MCT USE ONLY. STUDENT USE PROHIBITED

Lesson 1
Designing a Unified Messaging Deployment

Before you implement Unified Messaging in Exchange Server 2013, you need to design your deployment. Unified Messaging deployments can vary significantly depending on an organization's business requirements and infrastructure. Depending on your organization, you may need to design a Unified Messaging deployment that is scalable and highly available.

This lesson provides an overview of the planning process for a Unified Messaging deployment, and details on the types of information that you will need to collect before beginning your design phase. Furthermore, this lesson explains your options for designing scalable, highly available, and secure Unified Messaging deployments.

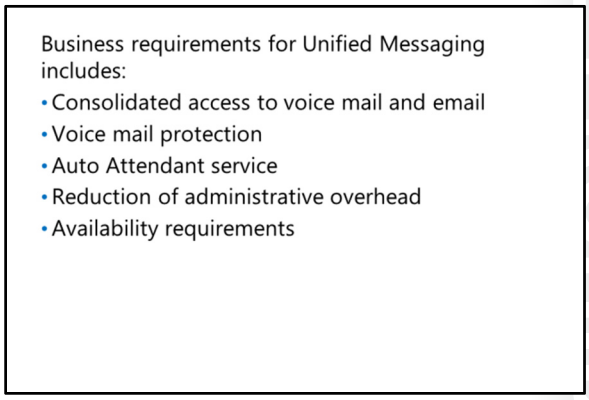
Lesson Objectives

After completing this module, you will be able to:

- Describe the business requirements and other organizational considerations that relate to a Unified Messaging deployment.
- Describe the types of information that you should collect at the beginning of a Unified Messaging deployment
- Design the infrastructure requirements for Unified Messaging.
- Describe the planning considerations for VoIP gateways.
- Design a scalable Unified Messaging infrastructure.
- Design a highly available Unified Messaging infrastructure.
- Design the integration of on-premises and hosted Unified Messaging.
- Design security for a Unified Messaging deployment.
- Describe considerations for implementing Unified Messaging codecs and file formats.

Business Requirements for Unified Messaging

When most organizations implement Unified Messaging into their business processes, their primary goal is to provide voicemail integration with email. Exchange Server 2013 offers more than just voice mail. For example, it also provides Auto Attendant and voice access to user mailboxes. When you are planning your deployment, you need to identify what your business requirements are so that you can ensure that your deployment addresses your business goals.



Consolidated Access to Voice Mail and Email

One of the common requirements for organizations that are implementing Unified Messaging is to ensure that users can access their voice mail and email from almost anywhere and by using almost any type of client. Exchange Server 2013 provides easy access to email by using clients such as Outlook, Outlook Web Access, and Exchange ActiveSync® mobile clients. Users typically access their voice mails by using phones.

MCT USE ONLY. STUDENT USE PROHIBITED

However, Unified Messaging combines the email and voice-mail infrastructures to provide access to voice mail to traditional email clients. This provides the ability for users to access their email by using traditional phones.

Voice Mail Protection

In most phone environments, it is difficult or impossible to protect personal voice messages from being forwarded to other users, who potentially should not have access to that information. Exchange 2013 offers a feature called Protected Voice Mail, which you can use to define rules for restricting the forwarding of voice messages. By integrating the features available in Active Directory® Rights Management Service (AD RMS) and Exchange Server 2013, you can ensure the protection of voice messages.

Auto Attendant Service

The second common requirement for organizations that are implementing Unified Messaging is to provide an automated call answering system. Organizations no longer want to use receptionists or operators as the primary means for directing calls, and would rather provide an automated means to assist users in searching the organization's phone list and for connecting callers to the right department or user. Exchange 2013 Unified Messaging provides an optional Auto Attendant component that enables callers to search for users, based on user name, departments, or extensions. You can configure the Auto Attendant with a wide variety of options, enabling users to access the right people efficiently and quickly.

Reduction of Administrative Overhead

Most large organizations have separate groups of administrators who manage telephony and email services. By combining a significant part of the telephony service with the email service, you can simplify your organization's administrative processes and reduce separation of services.

Availability Requirements

In most organizations, the telephone system is the most critical and sensitive core infrastructure application. Users are accustomed to telephone service always being available. Because Unified Messaging integrates with that system, you need ensure that Unified Messaging provides the same level of service that the telephone system provides. This means that you need to carefully consider the scalability, site deployments, high availability, and security required to design, build and operate the new infrastructure.

Collecting Information for a Unified Messaging Design

There are a number of questions that you must answer when you are planning a Unified Messaging deployment, so that you can ensure that your organization's business requirements are met.

- Identify the number of phone and voice mail users and the anticipated call volumes
- Document the current telephony environment
- Identify the storage requirements for voice mail
- Identify the network requirements for Unified Messaging

Determining the Overall Volume of Calls

When determining the overall number of clients and volume of calls, view call logs and monitor your network for voice sessions. It is typical for a customer to underestimate both the volume and duration of calls. A thorough analysis of customer voice traffic is necessary to ensure that there are enough servers and supporting equipment to handle the voice and other traffic.

If your organization uses a private branch exchange (PBX), ask your telephony department for monthly usage records.

MCT USE ONLY. STUDENT USE PROHIBITED

Determining the Number of Supported Users

The number of servers that you need to deploy is predicated on the total number of expected supported users you will serve. A single Client Access server and Mailbox server that provide Unified Messaging and regular email functionality easily can support a small. However, medium and large companies will have larger numbers of expected users and call volumes, so you need to ensure that your design will meet performance requirements. When calculating the number of users, make sure that you include any expected or potential growth.

Documenting the Current Telephony Environment

As part of the information-gathering process, identify the current telephone system. Does your company have an on-premises PBX? If it does, what type of PBX is it and does it support integration with Unified Messaging? How do users access their voice mail currently? Is there a traditional voice-mail system that connects to a PBX in place, where the users can listen to voice mail via phone only? Or does the organization make use of other IP PBX systems, where a high integration of voice mail is possible, and which users can access either from the email system or telephone soft client?

You also should document the placement of IP/ VoIP gateways, telephony equipment, and the Exchange 2013 servers. In most design scenarios, you must minimize the number of hops that packets must make between the PBX, the VoIP gateway, and telephony equipment. A good design consideration is to place VoIP gateways and Exchange servers on the same network, or within the same physical site.

Identifying Storage Requirements

After you collect all information about users call behavior and volume, you can consider the necessary additional storage that you need. Depending on the number of voice messages that you expect to handle, it might be necessary to increase the user maximum mailbox size and the maximum planned mailbox database size, and possibly decrease the user count for each database.

Identifying Network Requirements

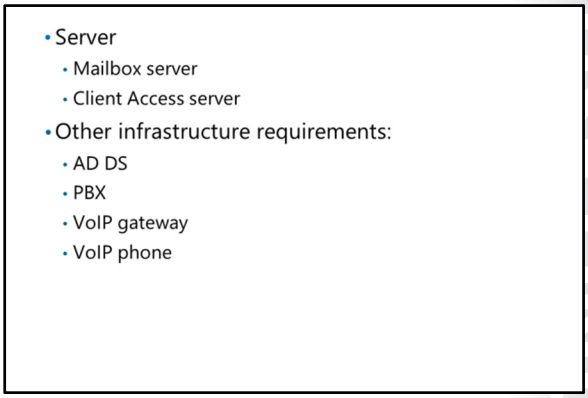
Adding Unified Messaging to your network will consume additional bandwidth. Unified Messaging requires a reliable and low latency network. When collecting information about your network, you need to understand the physical network topology, as well as whether the organization has multiple sites and how they connect. As part of the design process, ensure that the network is sized correctly and able to handle the additional voice traffic for Session Initiation Protocol (SIP) communication. This is particularly important if some users are in external offices with slower network connections to the Unified Messaging environment.

We recommend that network latency should be less than 20 milliseconds (ms) between the IP-PBX or VoIP Gateway and the Exchange 2013 servers. The total amount of required bandwidth depends on the codec that the dial plan uses and concurrent use of voice mail.

If you cannot guarantee network quality between the IP-PBX or VoIP Gateway and Exchange servers, your users might not be able to understand voice messages because of network latency or outages.

MCT USE ONLY. STUDENT USE PROHIBITED

Designing Infrastructure Requirements for Unified Messaging



An Exchange 2013 Unified Messaging deployment requires that you deploy several infrastructure components. There are several optional components, as well, which may enhance the user experience with Unified Messaging. The following sections detail the required infrastructure components.

Mailbox Server

The Mailbox server role provides most of the Unified Messaging services, including call answer, voice-mail recording, and auto-attendant services. When planning the Mailbox server role for Unified Messaging, ensure that you have 500 megabytes (MB) of additional disk space per Unified Messaging language pack on the operating system drive and approximately 250 kilobytes (KB) per voice message stored in the user's mailbox.

The Mailbox server role also is responsible for transcribing voice mail messages if you enable the Voice Mail Preview feature. The capability for voice mail speech recognition that this processor requires is processor intensive. Therefore, we recommend at least 12 central processing unit (CPU) cores on the Mailbox server for an average installation of 1,000 users, and a minimum of 8 gigabytes (GB) RAM.

Client Access Server

The Client Access server role accepts Unified Messaging connections from different sources, such as IP-PBX, IP Gateway, or Lync Server 2013. The incoming Session Initiation Protocol (SIP) traffic is redirected to the user's associated Mailbox server. You then configure a SIP or Real-Time Protocol (RTP) connection between the Mailbox server and the call source, without any additional involvement from the Client Access server.

Because the Client Access server only accepts and redirects the SIP connections, implementing Unified Messaging will not change the hardware requirements significantly for the Client Access server.

Active Directory Domain Services

Like other Exchange Server 2013 components, Unified Messaging depends on Active Directory Domain Services (AD DS). AD DS stores all Unified Messaging objects and their configuration settings. Some Unified Messaging AD DS objects logically represent a telephony hardware device, such as the IP address for a VoIP gateway. Other Unified Messaging AD DS objects are used to store configuration settings, such as a dial plan object or mailbox policy object.

AD DS domain controllers and global catalog servers must be available and provide adequate performance in order to ensure that all Exchange 2013 services are not affected. Adding Unified Messaging to the environment does not significantly change the performance requirements for an AD DS domain controller.

PBX

Exchange Server 2013 Unified Messaging does not provide a telephony system, so you still must deploy some type of telephone system in the organization. Most medium to large organizations have deployed an on-site PBX to provide the internal telephone system and the connection to the external public switched telephone network (PSTN). Part of your planning process for a Unified Messaging deployment should include verification that your PBX supports integration with Unified Messaging and that there are PBX configuration notes, which contain configuration and other settings required to deploy a PBX with Unified Messaging.

MCT USE ONLY. STUDENT USE PROHIBITED

VoIP Gateway

If the PBX does not support IP networking, you will need to deploy a VoIP gateway between the Exchange 2013 servers and the PBX. The VoIP gateway translates between the circuit-based network that the PBX uses and the packet-based network that the Exchange Servers use.

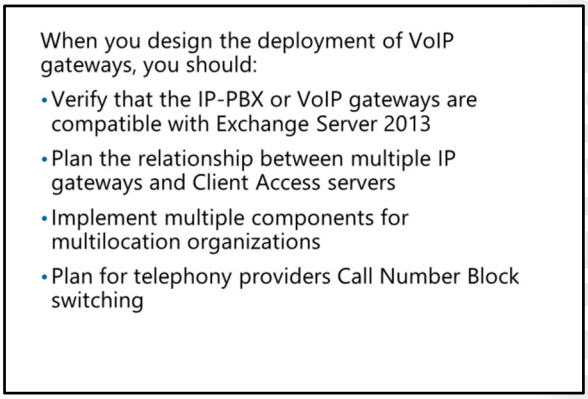
Like PBXs, it is important to verify that the VoIP gateways that you deploy in your organization are compatible with Exchange Server 2013 Unified Messaging and that integration configuration notes exist.

VoIP Phone

Organizations that have deployed a VoIP telephone system also have deployed VoIP phones. There are two types of VoIP phones available: software-based and hardware-based. A software-based phone, such as the Microsoft Lync system, is a communications program that runs from a computer. A hardware-based phone is similar to the phones found currently on desktops, except that they have added functionality. The Lync Phone Edition is one such phone, but there are many other phones available.

Design Considerations for VoIP Gateways

Properly designing and deploying IP/VoIP gateways for your organization is a critical step in your Unified Messaging deployment. Microsoft maintains a website that lists compatible IP gateways, and the required configuration notes and files. You must have these configuration notes to correctly deploy your organization's IP/VoIP gateways to work with Exchange Server 2013 Unified Messaging. It is equally important to match the number of IP/VoIP gateways that you have in your environment to the number of Exchange 2013 servers that are available.



Note: The list of supported PBX and VoIP gateways is revised frequently. For a list of devices supported for Exchange Server 2013, see the Telephony Advisor for Exchange 2013 page at http://go.microsoft.com/fwlink/?LinkId=290686

Unified Messaging and IP/VoIP

IP/VoIP gateways are available from multiple manufacturers, in sizes and models that range from 1 to 960 channels. You can deploy as many IP/VoIP gateways as necessary to provide for capacity and fault tolerance. If the number of calls or ports that are required is larger than the number of calls or ports that a single IP/VoIP gateway supports, you can increase the number of ports or the number of calls that can be accepted by installing and configuring additional IP/VoIP gateways, creating the Unified Messaging IP gateway object, and configuring the appropriate hunt groups to support your environment.

Multiple IP Gateways

You can configure IP gateways that Unified Messaging supports to route calls to Exchange 2013 servers in a round-robin manner, which is a load-balancing mechanism that Domain Name System (DNS) servers use to share and distribute network resource loads. To enable an IP gateway, you must configure each IP gateway with the IP address (or addresses) of your Client Access servers that answer calls from the IP gateway. These are the Client Access servers that are associated with the same dial plan as the Unified Messaging IP gateway object, which logically represents the IP/VoIP gateway. This enables all Unified Messaging IP gateways to forward incoming calls to the Client Access servers that are associated with the

MCT USE ONLY. STUDENT USE PROHIBITED

same dial plan. Then, if an IP gateway fails, the PBX will send the call to another IP gateway that can answer the call. The IP gateway, in turn, forwards the call to a Client Access server within the same dial plan. If the call is sent to a Client Access server that is not available, the IP gateway tries a second time to contact the server. If it is unsuccessful, it then uses the next Client Access server in the list to answer the call.

Multiple Locations

A company with multiple physical locations will frequently have an equal number of Active Directory sites. If this is the case, then each of those sites would have their own PBX and IP gateway. Each site also would have to configure one or more Unified Messaging Dial Plans, Unified Messaging Mailbox policies, and hunt groups.

If a location does not have a local Exchange Server 2013 implementation, you must consider the network links between the location and the closest office with Exchange 2013 Unified Messaging deployed. Ensure that the links have the necessary bandwidth to support the required network traffic and the increased Unified Messaging traffic. This is of special concern if you deploy Lync Server 2013 to support voice calls and conferencing.

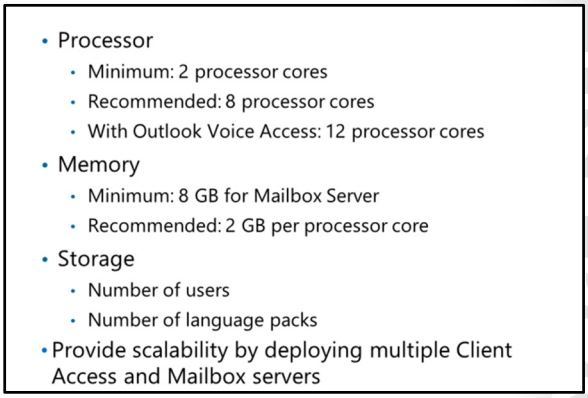
Call Number Block Switching by Telephony Providers

In several countries, telephony providers offer a highly available phone line, known as line failover. In this scenario, each provided phone line always has a unique phone-number block assigned to it, which makes it possible by using the second, redundant line as an individual PSTN connection.

In case of a failure or outage of the primary phone line, the telephony provider will switch the primary call-number block to the second PSTN line. This requires a special design, and you need to consider connecting theses phone lines to the same gateways or to two different gateway.

Designing for Server Scalability

The first option for implementing server scalability is to correctly design the Exchange server hardware and ensure that the required number of concurrent callers is supported. There are many factors to consider when selecting hardware for Exchange Server 2013. Three of the most critical factors to consider include choice of processor, amount of memory, and selection of storage.



Processors

You can deploy Exchange Server 2013 only on servers that are running on an X64-bit processor. Extensive testing on multi-core processors shows that Exchange Server benefits significantly when using multi-core processor technology, with four cores being optimal.

The Mailbox server role performs most of the processor intensive work in a Unified Messaging deployment. You should configure the Mailbox servers in a Unified Messaging deployment with eight processor cores. If Outlook Voice Access is enabled in a medium or large organization, the minimum processor cores should be increased to 12 cores.

MCT USE ONLY. STUDENT USE PROHIBITED

Memory

Having enough memory on the Mailbox servers is critical to ensure the best performance for Unified Messaging. The minimum memory requirement for Mailbox servers is 8 gigabytes (GB). For medium to large deployments, you should add a minimum of 4GB additionally to the Mailbox server RAM requirements.

Hard Drive Requirements for Storage

To install Exchange Server 2013 Unified Messaging, you need an additional 500 MB for each Unified Messaging language pack that you install. Another consideration is the duration of recorded messages in Unified Messaging. By default, there is a 20-minute maximum. You can modify this setting to between 5 and 100 minutes. Using the Windows® Media Audio (WMA) codec, a five-minute voice mail is approximately 250 KB in size.

You also must consider the number of voice-enabled users, and the size of their mailboxes, to help determine hard-drive space. Users who receive large numbers of voice mails may fill their mailboxes quickly if you implement mailbox quotas.

Implementing Scalability by Deploying Additional Servers

The second option for implementing scalability for the Unified Messaging deployment is to deploy multiple Exchange Server 2013 servers. Because the Unified Messaging services now run on all Client Access and Mailbox servers, you can use the same options for providing scalability for Unified Messaging services as you would for other Exchange services. This means that you can:

- Deploy multiple Client Access servers and configure Network Load Balancing (NLB) or a hardware load balancer to distribute the calls across multiple Client Access servers.
- Deploy multiple Mailbox servers and distribute the UM-enabled mailboxes across the available Mailbox servers. You can deploy the Mailbox servers in a Database Availability Group (DAG).

Designing Highly Available Deployments

When planning for a highly available Unified Messaging deployment, you need to include several components in your design.

When designing a highly available Unified Messaging deployment:

- Plan for multiple PSTN connections or multiple SIP trunks
- Plan for multiple VoIP gateways connected to redundant PSTN connections
- Implement multiple Client Access servers and Mailbox servers in a highly available configuration

PSTN

As described before, PSTN connections can be made highly available within a single location or across data centers, within the limitations of geographical dependencies of call-number assignments. If you are within a geographical region, either within a single location or across a county, it is possible to make use of the local PSTN providers. Most telephony providers can supply redundant PSTN lines with a possibility of call number-blocks switching.

VoIP Gateways

VoIP gateway positioning depends on whether the gateway connects directly to PSTN or whether it is behind a PBX. If the VoIP gateway is behind a PBX, you need to ensure the PBX is able to provide at least two connections that will connect to two identically configured gateways. You need to configure both gateways in Exchange Server 2013 Unified Messaging.

MCT USE ONLY. STUDENT USE PROHIBITED

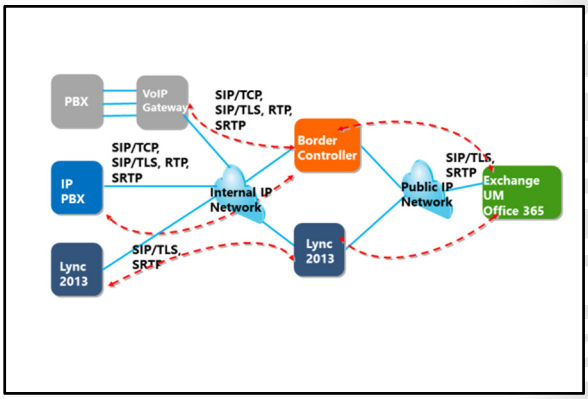
If the VoIP gateways connect directly to the PSTN, you will need to implement two PSTN lines terminated on two gateways. The configuration of each gateway should be identical.

Some organizations use SIP trunking to configure the connection to the PSTN. SIP trunking provides a packet-switched network connection to an external provider, which then provides the connection to the PST. To provide high availability in this scenario, ensure that you have multiple network connections to the SIP trunking provider.

Exchange 2013

Implementing redundancy for the Unified Messaging components in Exchange Server 2013 is straightforward. You only need to deploy multiple Client Access servers and Mailbox servers, and then use the normal Exchange Server 2013 options to ensure high availability. It is helpful to ensure that all Exchange 2013 servers within the same location and dial plan have the same configuration.

Designing for Office 365 Unified Messaging



You can integrate Exchange Server 2013 Unified Messaging with Microsoft Office 365™ by deploying a hybrid solution. In this deployment, the organization has a PBX or VoIP phone system deployed on-premises, but the Exchange 2013 Client Access server and Mailbox server are located on Office 365.

Since the Unified Messaging services are located in a Microsoft data center, the VoIP traffic must cross the public Internet to reach the Exchange servers. To implement this, you must place a Session Border Controller (SBC) and a Microsoft Lync Edge Server at the edge of your internal network. All traffic from the VoIP Gateway and IP PBX to and from Office 365 passes through the SBC. All traffic from the Lync passes through the Lync Edge Server.

Note: Lync Server is not required to integrate an on-premises telephone system with Unified Messaging on Office 365. Implementing Lync Server provides more options for a unified user experience.

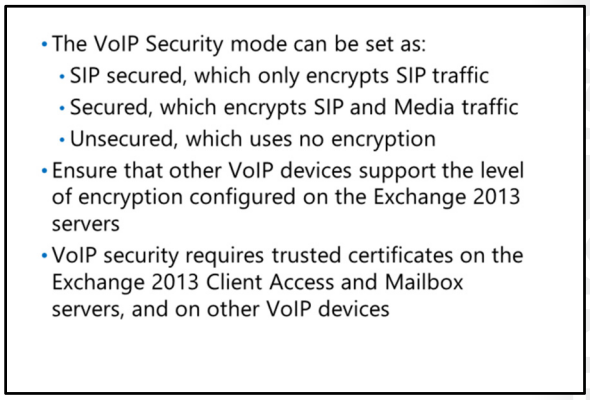
The purpose of the SBC is to protect the customer's private network against attack and intrusion. It is for use at a network's edge, and controls the flow of VoIP traffic to and from the private network to the public network (Internet). The SBC rewrites addressing information in headers when SIP messages pass from one network interface to the other. It secures the signaling and media data between itself and Office 365.

All communication between the on-premises phone deployment and Office 365 must use VoIP. This means that if the organization is using PBX that does not support VoIP, the organization must purchase and configure a VoIP gateway to connect the PBX to Exchange Unified Messaging.

For VoIP signaling and media between the gateway or IP PBX and Exchange Unified Messaging, the customer can choose to use unsecured protocols such as SIP/Transmission Control Protocol (TCP) or RTP, or secured protocols, such as SIP/Transport Layer Security (TLS) and secure real-time transport protocol (SRTP) protocols. Communication between the Microsoft Lync server and Exchange Unified Messaging must use secure protocols at all times.

MCT USE ONLY. STUDENT USE PROHIBITED

Designing for Unified Messaging Security



Client Access servers that are running the Microsoft Exchange Unified Messaging Call Router service (Signaling) and Mailbox servers that are running the Microsoft Exchange Unified Messaging service communicate with VoIP gateways, IP PBXs, and other Exchange computers in either Unsecured, SIP secured, or Secured mode, depending on how the Unified Messaging dial plan is configured. Client Access and Mailbox servers can operate in any mode configured on a dial plan, because you can configure the servers to operate in dual mode. This means that they will listen on TCP port 5060 for Unsecured requests and TCP port 5061 for Secured requests at the same time.

You can configure the VoIP security mode either when you are creating a new dial plan or after you have created a dial plan, by using the Exchange Administration Center (EAC) or Exchange Management Shell. You have three options when configuring the VoIP security mode:

- SIP secured. The SIP Secured setting means that only SIP traffic is encrypted by using TLS while RTP traffic is transmitted over TCP.
- Secured. The Secured traffic means that both SIP traffic and streaming media sent by RTP traffic are encrypted by using TLS. If you are using a Lync Server as the VoIP gateway, this is the option that you must select.
- Unsecured. All traffic is sent unencrypted. This is the default selection when you create a dial plan in Exchange Server 2013.

When you configure the Unified Messaging dial plan to use SIP secured or Secured mode, Client Access and Mailbox servers will try to encrypt the SIP signaling traffic or the RTP media channels, or both. However, to send encrypted data to and from Client Access and Mailbox servers, you must configure the Unified Messaging dial plan correctly, and VoIP devices, such as VoIP gateways, IP PBXs, and SBCs, must support mutual TLS.

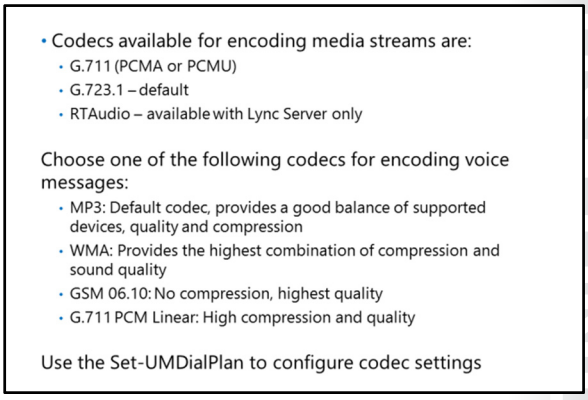
If you want to use mutual TLS to encrypt the VoIP traffic, you must have a certificate installed on the Client Access and Mailbox servers, and the other VoIP devices must trust the certificate. If you deploy an internal certification authority (CA) in the organization, you can use certificates from this CA if you can configure the VoIP devices to trust it. For example, if you are using Lync Server 2013 as the VoIP gateway, you should obtain certificates from the internal CA for both the Exchange 2013 servers and for the Lync 2013 servers. You also must configure the certificate for use by the Unified Messaging service on Mailbox servers and by Unified Messaging Call Router service on Client Access servers.

MCT USE ONLY. STUDENT USE PROHIBITED

Consideration of Codecs and File Formats

A *codec* is a software program that transforms—or codes—digital data into an audio file format or audio streaming format. It then converts, or *decodes*, the audio file back to the digital format. Codecs can vary in sound quality, the amount of bandwidth required to use them, and the system requirements necessary to do the encoding.

Exchange Server 2013 Unified Messaging uses codecs for encoding media streams between the IP/VoIP gateways and the Exchange servers, and for encoding and storing voice messages.



Choosing a Codec for Encoding Media Streams

Exchange Server 2013 can use the G.711 (Pulse Code Modulation A-law (PCMA), which is used in Europe and other countries, and Pulse Code Modulation μ-law (PCMU), which is used in North America and Japan) and the G.723.1 codecs to encode media streams. By default, the Exchange Server 2013 servers use the G.723.1 codec. This codec is widely supported on VoIP gateways.

If you use a Lync Server as the VoIP gateway, you have an additional option for providing higher quality voice recordings. When you configure a dial plan with a Lync Server as the Unified Messaging IP gateway, you have to configure the dial plan as a SIP uniform resource indicator (URI) dial plan. When you do this, the Exchange servers will use RTAudio wideband or high-fidelity audio for recording voice messages. RTAudio provides a higher sampling rate, so the quality of the voice recording will be better.

When the RTAudio codec is used, the voice message will be recorded in high fidelity and stored as an audio file that has a .wma extension. When the voice message is played back to the user in Office Outlook or Outlook Web Access, they will hear the voice message in high-fidelity audio. If users connect to their mailboxes by phone, the outbound media stream will be negotiated by using either the G.711 or G.723.1 codec. This means that callers will always hear lower fidelity audio over the telephone.

Choosing a Codec for Encoding Voice Messages

Exchange Server 2013 supports four codecs for encoding voice messages:

- MP3. This is the default format.
- WMA
- Group System Mobile (GSM) 06.10
- G.711 PCM Linear

To choose the right codec for encoding voice messages, you need to consider the types of clients that will be used to access the voice messages, the storage requirements for each codec and the network bandwidth available for replaying voice messages. The codec options provide the following benefits:

- MP3. The MP3 codec stores files in the .mp3 format, which means that it is compatible with the broadest range of mobile phones and devices and different computer operating systems. MP3 also provides very good compression of voice messages. A 30-second message recorded in an RTAudio codec will use about 120 KBs of storage, while a 30-second message recorded from a call using the G.723.1 codec will use about 60 KBs of storage.
- WMA. WMA provides the highest level of compression of any of the codecs. Since the .wma file format has a much larger header section than the .wav file format, the file size difference is most

MCT USE ONLY. STUDENT USE PROHIBITED

noticeable for messages longer than 15 seconds in length. A 30-second message recorded in an RTAudio codec will use about 70 KBs of storage, while a 30-second message recorded from a call using the G.723.1 codec will use about 40 KBs of storage. Therefore, for the smallest, but highest quality, audio files, use the WMA audio codec.

- G.711 PCM Linear. The G.711 PCM Linear audio codec creates uncompressed .wav audio files. Therefore, the voice-message recordings will require the most storage space. A 30-second message will consume about 240 KBs of storage. Because the files are not compressed, G.711 PCM Linear .wav audio files have the highest audio quality of the audio codecs that Unified Messaging uses. In most cases, the codecs that provide compression also provide acceptable sound quality, so we do not recommend the use of the G.711 PCM Linear audio codec in most cases.
- GSM. The GSM audio codec creates .wav audio files that are compressed. A 30-second message will consume about 50 KBs, which is slightly larger than the audio file that the WMA audio codec creates.

Managing Codec and Voice Recording Settings

To configure the codec and voice recording settings, you will use the **Set-UMDialPlan** cmdlet. The following parameters apply the codec settings:

- **AudioCodec**. Used to set the codec used in Exchange Server 2013 to record voice messages. The default is MP3.
- **MaxRecordingDuration**. Used to set the maximum length of time that messages can be recorded. The default is 20 minutes, but you can change the value to a number from 1 through 100. You may need to modify this number to balance storage requirements with the time necessary to leave meaningful messages.

Discussion: Designing a Unified Messaging Deployment

A. Datum has a consolidated data center that is running a single Exchange 2013 server on which both the Mailbox Server and Client Access server roles are deployed. The Exchange 2013 server and an IP-PBX are deployed in London. A. Datum has opened branch offices in Munich, Germany, and Singapore, and has deployed standalone PBXs in each branch office.

The London location has about 500 users, the office in Munich has about 100 users, and the Singapore office has approximately 10 users.

- What changes will you need to make to the infrastructure to address these requirements?
- What new server and hardware components will you need to deploy?
- What potential issues do you forsee with the deployment?

A. Datum has asked you enhance their Exchange 2013 infrastructure and fulfill the following requirements:

- Implement Exchange 2013 Unified Messaging
- Deploy Exchange 2013 servers only in the London data center.
- Ensure that the Unified Messaging system is highly available
- Ensure that users in all locations can use Unified Messaging
- Make an auto-attendant available for the London and Munich offices

MCT USE ONLY. STUDENT USE PROHIBITED

Lesson 2

Deploying and Configuring Unified Messaging Components

Planning and deploying Exchange Server 2013 Unified Messaging requires coordination between telephony, network, and Exchange Server administrators. During the deployment, you will need to configure connectivity between the Exchange Servers and the telephone system, across the organization's internal network. This lesson discusses how to deploy and configure Exchange Server 2013 Unified Messaging for your organization.

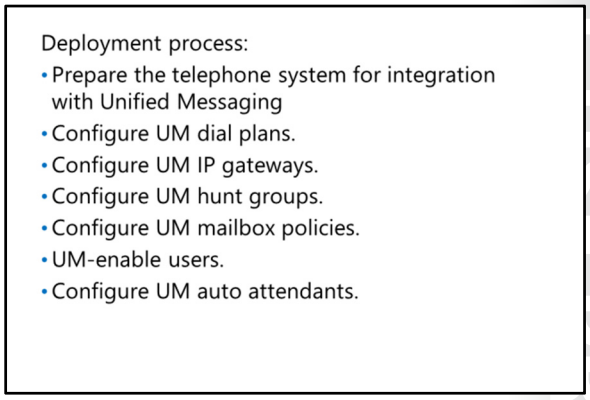
Lesson Objectives

After completing this lesson, you will be able to:

- Describe the overall process of deploying Unified Messaging.
- Configure the UM dial-plan object.
- Configure the UM IP gateway object.
- Configure the UM hunt-group object.
- Configure the UM mailbox-policy objects.
- Configure the UM auto attendants
- Configure protected voice mail

Overview of the Unified Messaging Deployment Process

During the Unified Messaging deployment, you need to configure several components, some which are required, and others that are optional. To complete the Unified Messaging deployment, complete the following steps:



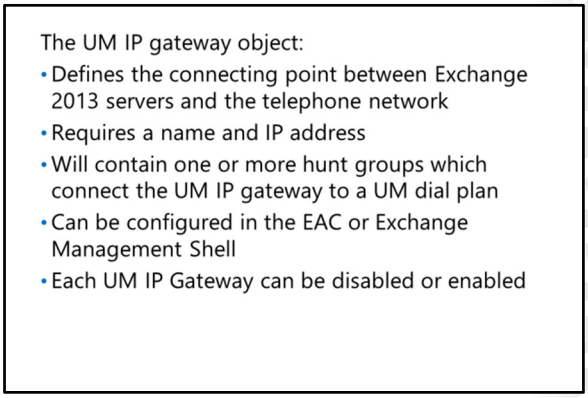
1. Prepare the telephone system for Unified Messaging integration. You may need to configure current telephone system to integrate with Unified Messaging. The exact steps will vary depending on the current system. For example, if your organization has deployed a traditional PBX, you will need to configure a VoIP gateway that can enable communications between the Exchange 2013 servers and the telephone system. After configuring the VoIP gateway, you will need to configure call routing rules and hunt groups on the PBX so that the calls can be directed to the Exchange 2013 servers.
2. Configure UM dial plans. The UM dial plan will mirror the dial plans configured on the telephone system, and you will use them to define phone extensions. UM dial plans are required before you can UM-enable users.
3. Configure UM IP gateways. The UM IP gateway objects define the connection point to the telephone system for the Exchange 2013 servers.

MCT USE ONLY. STUDENT USE PROHIBITED

4. Configure UM hunt groups. UM hunt groups route messages, and create the connection between the UM dial plans and the UM IP gateways. When you create a UM IP gateway, a default UM hunt group is created automatically. You can configure additional hunt groups optionally.
5. Configure UM mailbox policies. UM mailbox policies define the user experience with Unified Messaging. When you create a UM dial plan, a default UM mailbox policy is created automatically. You can modify the default policy, and create additional policies if required.
6. UM-enable users. You must UM-enable users before they can start using Unified Messaging. When you UM-enable users, you assign an extension and UM mailbox policy to the user.
7. Configure UM auto attendants. UM auto attendants are an optional object. You can configure the UM auto attendant to answer and direct calls within the organization.

Configure a UM IP Gateway Object

In Unified Messaging, the UM IP gateway object defines the connection point between the Exchange 2013 servers that are running the Unified Messaging services and the telephone network. Exchange Server 2013 uses the UM IP gateway object to accept calls from the telephone network and to route calls to the telephone network.



The UM IP gateway object references a physical VoIP gateway, IP-PBX, SPC or a Lync Server. A UM IP gateway has organization-wide scope, and each UM IP gateway can reference only a single physical IP gateway. When you configure the UM IP gateway object in Exchange 2013, you must configure the target IP address and an object name.

The UM IP gateway contains one or more UM hunt groups and configuration settings. UM hunt groups link a UM IP gateway to a UM dial plan. By creating multiple UM hunt groups, you can associate a single UM IP gateway with multiple UM dial plans.

After you create a UM IP gateway, the Mailbox servers linked to the UM IP gateway send a SIP OPTIONS request to the IP address configured in the UM IP gateway object, to ensure that the device is responsive. If the device does not respond to the request, the Mailbox server logs an event with ID 1400 stating that the request failed. If this happens, make sure that the VoIP gateway, IP PBX, or SBC is available, and online, and that the Unified Messaging configuration is correct.

You can enable or disable the Unified Messaging IP gateway. If you disable a Unified Messaging IP gateway, it can be in one of two disabled modes. The first disabled mode forces all associated Exchange 2013 servers to drop existing calls. The second disabled mode forces the Exchange 2013 server associated with the Unified Messaging IP gateway to stop handling any new calls that the IP gateway presents. You would use this option if you need to restart the Exchange Server, but do not want to disrupt ongoing calls.

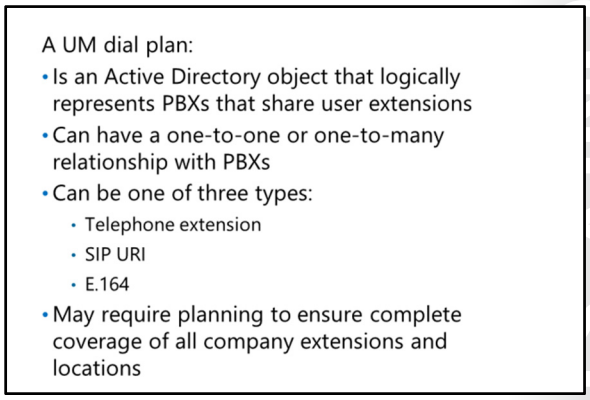
To create and configure an UM IP Gateway, you can use the EAC or Exchange Management Shell. The cmdlets that you can use to manage the UM IP gateway are:

- **New-UMIPGateway**
- **Set-UMIPGateway**
- **Get-UMIPGateway**

MCT USE ONLY. STUDENT USE PROHIBITED

- **Enable-UMIPGateway**
- **Disable-UMIPGateway**
- **Remove-UMIPGateway**

Configuring UM Dial Plan Objects



A *dial plan object* is a container object in AD DS that represents a set or grouping of PBXs logically that share common user-extension numbers. In practical terms, users' extensions hosted on PBXs share a common extension number. Users can dial one another's telephone extensions without appending a special number to the extension or dialing a full telephone number. A UM dial plan is a logical representation of a telephony dial plan. All users within a dial plan have a unique extension number, and the combination of dial plan and the user extension uniquely identifies each UM-enabled user. When you create the UM dial plan, you need to associate it with a Unified Messaging server.

A UM dial plan mirrors a telephony dial plan. You configure a telephony dial plan on PBXs or Lync Server.

In Unified Messaging, the following UM dial-plan topologies can exist:

- A single dial plan that represents a subset of extensions or all extensions for an organization with one PBX or IP PBX. Use this configuration in small customer environments.
- A single dial plan that represents a subset of extensions or all extensions for an organization with multiple PBXs or IP PBXs. Use this option in organizations that have deployed multiple PBXs, but a single set of extensions.
- Multiple dial plans that represent a subset of extensions or all extensions for an organization with one PBX or IP PBX. Use this in complex PBX environments for larger organizations.
- Multiple dial plans that represent a subset of extensions or all extensions for an organization with multiple PBXs or IP PBXs. Use this topology is your organization has many geographically disparate locations.

A dial plan can exist in three different configurations:

- Telephone Extension. This is the most common type of UM dial plan, and you use it with PBXs and IP gateways that support the telephone extension (TelExtn) URI type.
- SIP URI. This is the dial plan that you use when integrating Exchange Server 2013 and Lync Server 2013. The SIP URI resembles an email address, and is sip:*<user name>@<domain or IP address>* format.
- E.164. E.164 is the standard numbering format that you use for the international public-telecommunication numbering plan on the PSTN and some data networks. E.164 numbers can have a maximum of 15 digits, and typically are written with a plus sign before the telephone number. Use an E.164 dial plan type when the IP-PBX or VoIP gateway only support this type.

MCT USE ONLY. STUDENT USE PROHIBITED

Determining an Effective Numbering Plan

One of the components of designing the dial plans is planning the phone number and extension configuration. In many cases, this will have occurred during the telephone system installation, but you may need to work together with the telephone administrators to create this design. Determining an effective numbering plan is based on several factors:

- Does the numbering plan denote the physical sites or departments? One option is to have a different numbering plan for each physical location.
- What is the number of users, and is growth factored into the numbering plan? Basing a dial plan on the current number of users may make your ability to expand the plan more difficult in the future.
- Several factors determine numbering plans: the number of employees in your organization, the departments, and their physical structure. You may use a numbering plan that denotes not only the extension, but also the geographical location of the extension or the department.
- How are international sites numbered? You likely are limited in your ability to have a standardized numbering plan with overseas offices.

You can use the following Exchange Management Shell cmdlets to manage UM dial plans:

- **Set-UMDialPlan**
- **New-UMDialPlan**
- **Get-UMDialPlan**
- **Remove-UMDialPlan**
- **Set-UMCallRouterSettings**
- **Get-UMCallRouterSettings**

Demonstration: Configuring UM Dial Plans and UM IP Gateways

In this demonstration, you will see how to create and configure UM dial plans and UM IP Gateways.

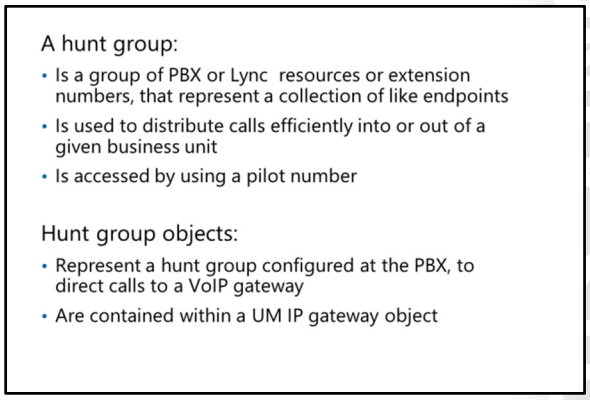
Demonstration Steps

1. On LON-CAS1, in Internet Explorer®, connect to https://lon-cas1.adatum.com/ecp. Sign in as **Adatum/administrator** with the password **Pa$$w0rd**.
2. On the unified messaging feature pane, create a new UM dial plan with the following configuration:
   a. Name: **UM-DIALPLAN**
   b. Extension length: **5**
   c. Dial plan type: **Telephone extension**
   d. VoIP security mode: **Unsecured**
   e. Country code: **1**
3. Edit the UM-Dialplan by changing the following settings:
   - Outside line access code: **9**
   - International access code: **00**.
   - Outlook Voice Access numbers: **12224445555**
4. Review the other default settings.

MCT USE ONLY. STUDENT USE PROHIBITED

5. Create a new UM IP gateway with the following settings:
   - Name: **UM-Gateway**
   - Address: **172.16.0.40**
   - UM dial plan: **UM-Dialplan**

Configuring Hunt-Group Objects



A *hunt group* is an extension that is defined as a group of telephone numbers that are treated as one in some situations. Hunt groups often are used to identify a group of telephone extensions, such as help-desk or call-center personnel. When users call the phone number associated with the hunt group, the call is forwarded to any extension available in the hunt group. In most cases, a hunt group represents a set of identical resources that an application or a group shares. This provides more efficient access to applications such as voice mail or an auto attendant, so callers will not experience a busy signal. Instead, the PBX hunts for an open line to which to connect them.

Pilot Number

A pilot number is the way in which the PBX identifies a hunt group. In other words, a *pilot number* is the address or label for the hunt group. It is a dummy extension, and does not have a person or phone associated with it. It is the number to which a coverage path routes a call.

When you use a PBX with Exchange Server Unified Messaging, it uses a pilot number to target a diverted ring, no answer, or busy call to Exchange Server Unified Messaging, so a message can be taken. Subscribers can use this same pilot number--or a different number--to access the messages in their Exchange Server mailbox. You also can use a pilot number for top-level access to an Exchange Server UM auto attendant.

Implementing UM Hunt Groups

The UM *hunt group object* is a logical representation of an existing PBX or Lync hunt group. The UM hunt group object locates the PBX or IP PBX hunt group from which an incoming call is received. The pilot number that you define for a hunt group in the PBX or IP PBX must be the same as the pilot number assigned to the UM hunt group. The pilot number matches the information about the incoming calls through the SIP signaling message information on the message. The pilot number enables the Client Access server to interpret the call together with the correct dial plan, so that the call can be routed correctly. If the pilot number from the call does not match a pilot number defined on the UM hunt group, the Exchange 2013 servers cannot route incoming call.

UM hunt group objects act as a connection or link between the UM IP gateway and the UM dial plan. Therefore, you must associate a single UM hunt group with at least one UM IP gateway and one UM dial plan.

When you create a new hunt group object, you enable the Exchange 2013 servers in the specified dial plan to communicate with the UM IP gateway object. When creating a new UM hunt group object, you need to specify the dial plan, and the pilot identifier or pilot number, to be used with the new UM hunt group.

MCT USE ONLY. STUDENT USE PROHIBITED

You can have multiple Exchange 2013 servers associated with a single hunt group. You can configure a single Exchange 2013 server to support up to 200 simultaneous calls. If you estimate having more than that, you would need to have multiple UM servers.

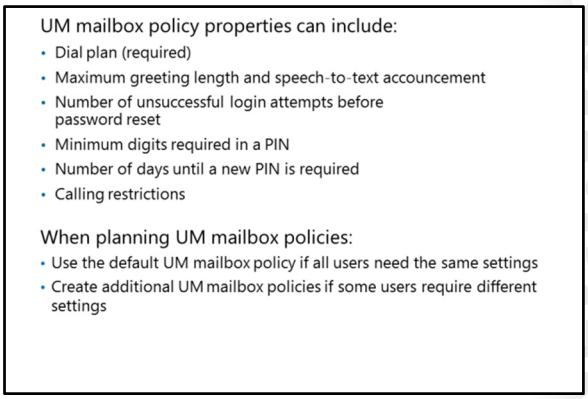
When you create a UM IP gateway, and associate the gateway with a UM dial plan, a default UM hunt group is created. You can associate additional UM hunt groups with the same UM IP gateway.

You can use the following Exchange Management Shell cmdlets to manage UM hunt groups:

- **New-UMHuntGroup**
- **Get-UMHuntGroup**
- **Remove-UMHuntGroup**

Configuring UM Mailbox Policies

UM mailbox policies are required when you enable users for Unified Messaging. You must link the mailbox of each UM-enabled user to a single UM mailbox policy. When you enable Unified Messaging, a default UM mailbox policy is created, and all UM-enabled users are assigned to the policy automatically.



You can use UM mailbox policies to set Unified Messaging user settings, such as:

- Personal Identification Number (PIN) policies
- Dialing restrictions
- Other settings such as maximum greeting duration

Planning UM Mailbox Policies

When planning UM Mailbox Policies, use the following approach:

- Begin by identifying the requirements for configuring the Unified Messaging experience for users. If your organization wants the same settings applied to all users, you can edit the default UM mailbox policy to meet your requirements, and the settings will be applied to all users.
- In some organizations, you may want to set different policies for different users. For example, some users may be working with more confidential information, so you may want to set a stricter PIN policy for these users. You can configure the policy to require a longer PIN and to lock out the user more quickly if he or she enters an incorrect PIN.
- If you require different settings for different groups of users, identify the groups of users that have the same requirements. Create a UM mailbox policy that matches each set of requirements, and assign the appropriate users to each policy.
- Create the UM mailbox policies before you UM-enable mailboxes. In this way, you can assign the UM mailbox policy when you UM-enable the mailbox.
- UM mailbox policies are defined on a UM dial plan. The UM mailbox policy is specific to the dial plan, which means that you cannot use one dial plan with multiple dial plans.

You can use the following Exchange Management Shell cmdlets to manage UM mailbox policies:

- **Get-UMMailboxPolicy**

MCT USE ONLY. STUDENT USE PROHIBITED

- **Set-UMMailboxPolicy**
- **New-UMMailboxPolicy**
- **Remove-UMMailboxPolicy**
- **Set-UMMailbox. –UMMailboxPolicy**. This command assigns a UM mailbox policy to a user account.

Demonstration: Configuring Mailbox Policy Objects, and Enabling Mailboxes for Unified Messaging

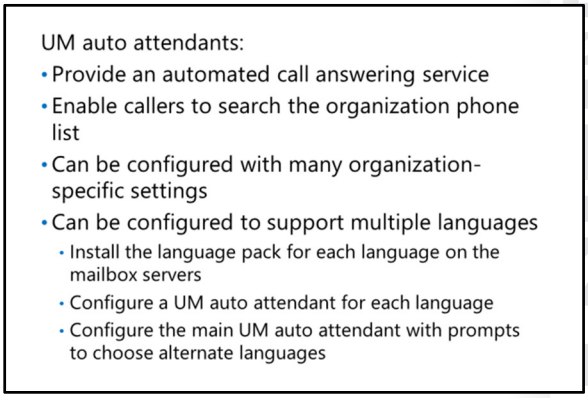
In this demonstration, you will see how to configure the Mailbox Policy that was created automatically during the dial-plan configuration. You also will see how to enable a user mailbox for UM.

Demonstration Steps

1. On LON-CAS1, in the EAC, open the UM-Dialplan properties.
2. Open the UM-Dialplan Default Policy that was created when you created the dial plan.
3. Review the options available in a dial plan.
4. Access Amr Zaki's mailbox properties, and then enable Amr's mailbox for UM by using the extension **22222** and a PIN of **135792**.

Implementing UM Auto Attendants

In telephony or Unified Messaging environments, an auto attendant is used to answer telephone calls and help callers search the internal phone system for an intended recipient, or it transfers callers to the extension of a user or department without a receptionist or operator having to intervene. The auto attendant can provide a simple service, such as enabling callers to search for and connect to simple extensions, or complex services, such as speech recognition and a vast menu system that enables callers to sort and search the internal telephone system.



Exchange Server 2013 Unified Messaging enables you to create one or more UM auto attendants. An auto attendant provides the menu system that lets internal and external users navigate through configured options and place calls to desired recipients. You can present announcements through a .wav file or speech-to-text, so that the caller can navigate through the menu options quickly and easily, enabling them to locate and call the person with whom the user wants to speak.. For navigation, the caller can use dual tone multi-frequency (DTMF) or voice inputs.

You can configure the UM auto attendant with a large set of options, including:

- Customized .wav files for greetings or other voice responses.
- Business hours
- Holiday schedules
- Different responses to user calls during or out of business hours, or on holidays
- Preferred language

MCT USE ONLY. STUDENT USE PROHIBITED

- Key mappings, which enable users to navigate menu items by pressing numbers or through voice prompts

You can configure the UM auto attendant with one or more pilot identifiers. When users call the pilot identifier number, the phone system connects them automatically to the UM auto attendant. If your organizations needs more than one UM auto attendant, you can configure multiple UM auto attendants, and provide a different pilot identifier for each one.

Multiple Language Support with UM Auto Attendants

Many organizations need the capability to provide services in multiple languages. In these scenarios, you can configure the UM auto attendants to support more than one language. Each UM auto attendant is configured with a default prompt language. This setting defines the default language that the caller will hear when the auto attendant answers the incoming call. This applies only if you are using the default prompts that Exchange Server 2013 provides. This language setting affects only the default system, but does not affect custom prompts that you configure on an auto attendant. The language selected as the default for the auto attendant is based on the installed version of Exchange 2013.

When you install the U.S. English version of Exchange Server 2013, U.S. English is the only language available. If you install a localized version of Exchange 2010, you can configure the auto attendant that you create to use the localized language or U.S. English as the default language.

To provide multiple language support for the UM auto attendant, you need to install additional Unified Messaging language packs on the Exchange 2013 Mailbox server. Then, you need to configure multiple UM auto attendants--one for each language. Please note that a UM auto attendant can have one language only configured at a time. So if you need to support multiple languages, you need to configure the main auto attendant with one of the languages, and then configure the appropriate key mappings to access the other auto attendants that use the other languages. Users can then select the alternate languages for the UM auto attendant by using their phone or voice prompts.

You can use the following Exchange Management Shell cmdlets to manage UM auto attendants:

- **Get-UMAutoAttendant**
- **New-UMAutoAttendant**
- **Set-UMAutoAttendant**
- **Remove-UMAutoAttendant**
- **Disable-UMAutoAttendant**
- **Enable-UMAutoAttendant**

Demonstration: Configuring UM Auto Attendants

In this demonstration, you will see how to create and configure an UM auto attendant.

Demonstration Steps

1. On LON-CAS1, in the EAC, open the UM-Dialplan properties.
2. Create a new UM auto attendant with the following configuration:
   a. Name: **UMAutoAttendant**
   b. Create this auto attendant as enabled
   c. Set the auto attendant to respond to voice commands
   d. Access number: **12224443333**.

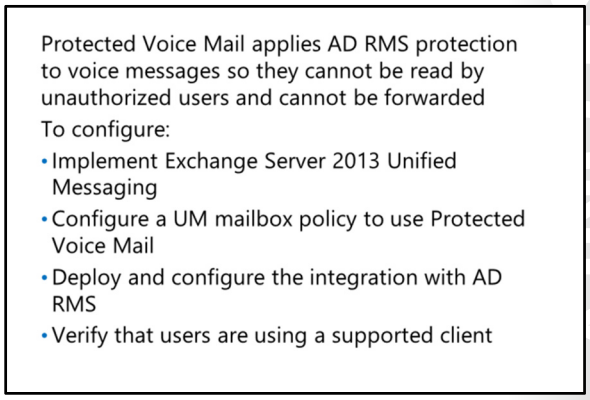
MCT USE ONLY. STUDENT USE PROHIBITED

3. Access the auto attendant properties, and then review the options available in a dial plan.

Configuring Protected Voice Mail

Protected Voice Mail provides security for recorded voice messages by encrypting and protecting voice mail. To deploy Protected Voice Mail, you first must implement AD RMS, and then configure the integration between the Exchange server environment and the AD RMS servers.

When a protected voice mail is accessed from clients such as Outlook 2013, Outlook Web App, and Outlook Voice Access, the Exchange servers ensure that only the intended recipient or recipients of the message can access its content, and users are blocked from forwarding the message.



To enable Protected Voice Mail, you need to:

- Configure Unified Messaging by configuring the UM dial plan and UM IP gateway.
- Configure UM Mailbox Policy that requires Protected Voice Mail. When configuring your UM Mailbox Policy to require Protected Voice Mail, configure the following parameters:
    - **ProtectAuthenticatedVoiceMail**. This parameter specifies whether the Exchange 2013 servers create protected voice mail messages for UM-enabled users. If the value is set to Private, only messages marked as private are protected. If the value is set to All, every voice mail message is protected. The default is none, which means that no protection is applied to voice-mail messages.
    - **ProtectUnAuthenticatedVoiceMail**. This parameter is the same as the previous parameter, but also applies to scenarios where automated messages are sent from the Unified Messaging system to the user mailbox.
    - **RequireProtectedPlayOnPhone**. This parameter specifies whether users can utilize Play on Phone to listen to protected voice-mail messages or whether they can use multimedia software to play protected messages. The default is $false, which means that users can use both means to listen to protected messages.
- Install and configure AD RMS and configure the integration of AD RMS and Exchange Server 2013. You will use the **Set-IRMConfiguration** cmdlet to configure the integration.

Note: The integration of Exchange Server 2013 and AD RMS is covered in more detail in Module 5.

- Ensure that the users are using supported clients to access the protected voice messages. Users need to use Outlook 2010 or Outlook 2013, Outlook Web App on Exchange Server 2010 or Exchange Server 2013 or Outlook Voice Access to access their protected voice messages. ActiveSync clients are not supported.

MCT USE ONLY. STUDENT USE PROHIBITED

Lesson 3

Designing and Implementing Exchange Server 2013 UM Integration with Lync Server 2013

You can configure the integration of Exchange Server 2013 Unified Messaging with Lync Server 2013 Enterprise Voice to provide a complete voice infrastructure. In this configuration, Lync Server 2013 provides the voice functionality that the PBX or IP-PBX provides, while Unified Messaging provides the voice messaging functionality and auto attendant services.

This lesson provides an overview of the Enterprise Voice features that Lync Server 2013 provides, and then describes how to plan and implement the integration of Exchange Server 2013 and Lync Server 2013.

Note: This lesson describes how to integrate Exchange Server 2013 and Lync Server 2013. You also can configure Exchange Server 2013 to integrate with Lync Server 2010 by using the same procedures as this lesson describes.

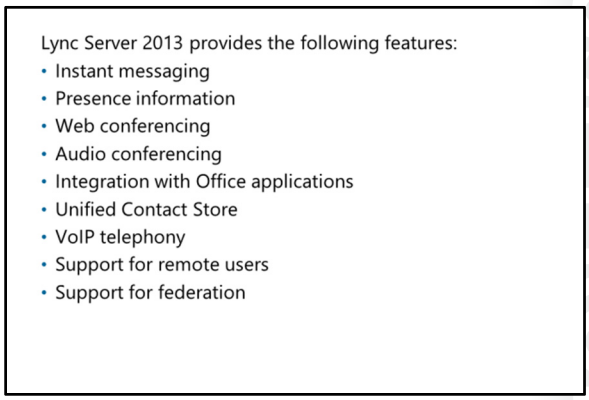
Lesson Objectives

After completing this lesson, students will be able to:

- Describe Lync Server 2013 features.
- Describe the Enterprise Voice features available in Lync Server 2013.
- Provide an overview of the integration components between Exchange Server 2013 and Lync Server 2013.
- Configure the integration of Exchange Server 2013 and Lync Server 2013.

What Is Lync Server 2013?

Exchange Server 2013 and Lync Server 2013 are designed to integrate and work together to provide a complete email and voice system. Exchange 2013 provides an email-messaging system, while Lync 2013 provides a telephony system when you configure it for Enterprise Voice. Unified Messaging can use Lync 2013 to provide the telephony component it needs, while Lync 2013 can use Unified Messaging to provide voice-mail functionality.



When you configure the integration of Exchange Server 2013 and Lync Server 2013, Exchange Server 2013 will use the Lync Server as its IP PBX. On Exchange Server 2013, you will configure an IP Gateway that references the Lync 2013 server.

Lync 2013 also provides other features that integrate with Unified Messaging, such as instant messaging, presence information, Web conferencing, and VoIP telephony:

- Instant messaging. The Lync 2013 client provides instant messaging (IM) functionality that the Lync hosts. The solution provides IM features, such as group IM, and extends the internal IM infrastructure to external IM providers.

MCT USE ONLY. STUDENT USE PROHIBITED

- Presence information. Lync 2013 tracks presence information for all Lync users, and it provides this information to the Lync 2013 client and other applications, such as Outlook 2013.
- Web conferencing. Lync 2013 can host on-premise conferences, which you can schedule or reschedule, and they can include IM, audio, video, application sharing, slide presentations, and other forms of data collaboration.
- Audio conferencing. Users can join Lync 2013-based audio conferences by using any desk or mobile phone. When connecting to an audio conference by using a Web browser, users can provide a telephone number that the audio-conferencing services calls.
- Integration with Office applications. When you implement Lync Server 2013, Exchange Server 2013, Microsoft SharePoint Server® 2013, and Microsoft Office 2013, you can provide a seamless user experience between all of the applications. For example, if you receive an email from another user, you can see the user presence information when you read the email. When a user sets an out-of-office response in Outlook, you will see that same response in your Lync client when viewing the user's presence information.
- Unified Contact Store. The Unified Contact Store feature enables users to store all contact information in there Exchange Server 2013 mailbox, so that the contact list is available in Lync, Outlook, and Outlook Web Access. The Unified Contact Store is enabled by default in Lync.
- VoIP telephony. Enterprise Voice enables Lync 2013 users to place calls from their computers by clicking an Outlook or Lync contact. Users receive calls simultaneously on all of their registered user endpoints, which may be a VoIP phone, mobile phone, or Lync 2013 client. The Lync 2013 Attendant is an integrated call-management client application that enables a user, such as a receptionist, to manage many conversations simultaneously.
- Support for remote users. Lync Server 2013 has an Edge Server role that enables remote users to use all Lync Server features without a virtual private network (VPN) connection.
- Support for federation. You can configure federation with other organizations that are running Lync Server or Microsoft Office Communications Server, and provide full Lync functionality for users between the two organizations.

With Lync, users can keep track of their contacts' availability (Presence); conduct an Instant Messaging (IM) session; make calls via VoIP; initiate or join an audio, video, or web conference; or make a phone call within the Lync organization, with federated partners or to phones on the PSTN. The Microsoft Lync 2013 desktop client is available for Windows and for the Macintosh operating system, and mobile versions are available for Windows® Phone, iPhone iPad, and Android devices.

Note: The Lync client makes use of autodiscover information, much like the Outlook client does. The Lync 2013 queries the Exchange AutoDiscover service for connection information for both internal and external connectivity to Exchange 2013, for the location of the user's Mailbox server, and for the URLS for Outlook features, such as free/busy information. If the autodiscover records are not configured correctly, Lync clients may not be able to display information about the user, such as their Exchange free/busy status or out-of-office messages.

Demonstration: Exploring Lync 2013 Features

In this demonstration, you will see how to:

- Enable user accounts for Lync Server 2013
- Use Instant Messaging and Desktop Sharing in Lync 2013

MCT USE ONLY. STUDENT USE PROHIBITED

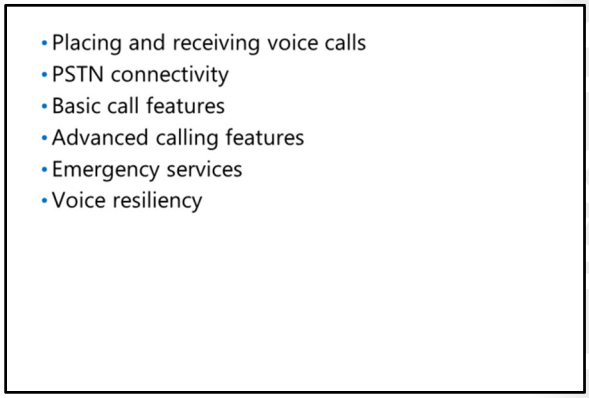
- Create and join a Lync Meeting

Demonstration Steps

1. On LON-LY1, open the Lync Server Control Panel, and sign in as **Administrator** using the password **Pa$$w0rd**.
2. Enable both April Reagan and Brad Sutton for Lync, assigning both to the LON-LY1.ADATUM.COM pool.
3. On LON-CL1, sign in as **April** using the password **Pa$$w0rd**.
4. Open Outlook 2013, and then configure the Outlook profile.
5. Open Lync 2013, and then verify that April is signed into Lync automatically.
6. On LON-CL2, sign in as **Brad** using the password **Pa$$w0rd**.
7. Open Outlook 2013, and then configure the Outlook profile.
8. Open Lync 2013, and then verify that Brad is signed into Lync automatically.
9. Test Instant Messaging functionality between the two users.
10. Test Desktop Sharing between the two users.
11. On LON-CL2, create a Lync Meeting request starting in the next 30 minutes, and then send it to April.
12. On both client computers, verify that you can join the Lync meeting.

Enterprise Voice Components in Lync Server 2013

The Enterprise Voice component in Lync Server 2013 provides a full featured VoIP solution that you can use to enhance or replace traditional PBX telephone systems. The Enterprise Voice component provides the functionality that the following sections describe.



Placing and Receiving Voice Calls

Enterprise Voice enable users can initiate calls from a Lync client by typing a name or phone number on their keyboard, or using a dial pad displayed on their screen. Users also can utilize VoIP Phone Editions or mobile devices to make voice calls via the Lync Server infrastructure. These devices can be active simultaneously.

Users are alerted to incoming calls on all of their devices simultaneously, with customizable ringtones on IP phone devices and a notification similar to an instant message on their computers.

PSTN Connectivity

A Lync Server 2013 Enterprise Voice deployment supports calls to and from the PSTN. Connecting Enterprise Voice to the PSTN requires one or more of the following:

- A SIP trunk to an Internet Telephony Service Provider (ITSP)
- An IP-PBX connected to the PSTN
- A PSTN gateway connected to the PSTN
- A Survivable Branch Appliance (SBA) or Survivable Branch Server connected to the PSTN

MCT USE ONLY. STUDENT USE PROHIBITED

Basic Call Features

Enterprise Voice provides all of the basic features that a traditional PBX provides. For example, while Lync users are on a call, they can answer additional incoming calls or initiate outgoing calls, and the existing active call is put on hold automatically. Users can transfer calls from one user to another, either directly or after the first user speaks privately with the second user. Users also can transfer calls to another device. For example, they could transfer an active call to their mobile phone.

Advanced Calling Features

Enterprise Voice includes several advanced calling features as well, such as:

- Call Parking, which enables users to put a call on hold, and then retrieve it from another phone. When a user parks a call, the original answering phone becomes free for another call.
- Delegation, which enables users to assign call handling to one or more assistants, such as a Personal Assistant or a Colleague. The delegate can perform multiple calling tasks on behalf of the user who initiated the delegation, including screening calls, placing calls, and initiating conferences.
- Team calling, which enables a user to have incoming calls simultaneously ring the phones of teammates, for functions such as group-call pickup and department calling.
- Response Groups, which you can configure for queuing and routing calls intelligently to designated agents. You typically would use this for groups such as your information technology (IT) helpdesks, an accounting hotline, and other internal contact centers.

Emergency Services

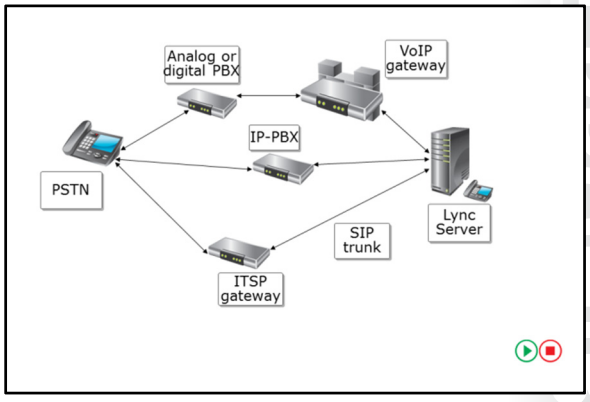
Lync Server 2013 supports enhanced 9-1-1 (E9-1-1) for North America. This feature provides additional location information to dispatchers of emergency services.

Voice Resiliency

The new voice resiliency capability allows a site with an SBA or Survivable Branch Server to continue to provide users with the ability to make and receive Enterprise Voice calls if the wide area network (WAN) that connects the branch and central sites is down. You also can configure it to provide resiliency between central sites.

Enterprise Voice Options for Connecting Lync Servers to the PSTN

When you deploy Enterprise Voice on Lync Server, you have several options for connecting the Lync Server to the PSTN. Like the Exchange 2013 servers that provide Unified Messaging services, Lync servers can communicate only on packet-switched networks. Therefore, some type of gateway is required between the Lync server and the PSTN.



Note: To connect a Lync Server to the PSTN, you must deploy the Mediation Server role in the Lync Server environment. You can collocate this role on Lync front-end servers, or you deploy dedicated servers that are running the Mediation server role only.

At a high level, there are three options for connecting a Lync Server deployment to the PSTN, including:

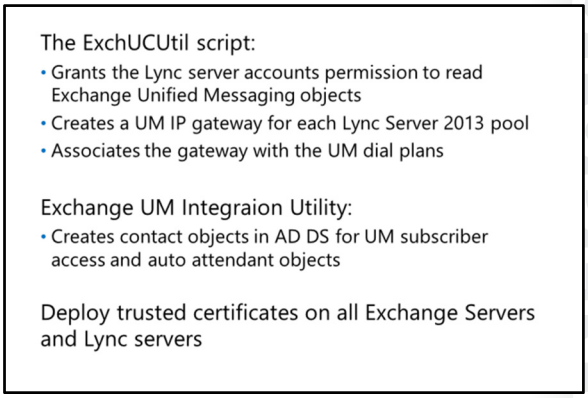
MCT USE ONLY. STUDENT USE PROHIBITED

- Connecting through a VoIP gateway and traditional PBX. This scenario is common in organizations that have deployed an analog or digital PBX, and which want to retain the PBX for their telephone systems. In this deployment, the VoIP gateway provides protocol conversion between the packet-based network where Lync Server is deployed, and the PBX, which is connected to the PSTN.
- Connecting through an IP-PBX. This scenario is common for organizations that have deployed an IP-PBX, and which want to retain PBX for all or part of their telephone system. In this deployment, the IP-PBX provides protocol conversion between the PSTN and the packet-based network where Lync Server is deployed.

Connecting through a SIP trunk to an ITSP. In this scenario, all PBXs have been removed from the company, and all telephone services are provided by Lync Server. To connect to the PSTN without a PBX, you can implement SIP trunking, which provides a packet-based network connection to an ITSP. At the ITSP, a VoIP gateway translates the packet-based traffic to circuit-based traffic, and connects to the PSTN.

Overview of an Exchange Server 2013 and Lync Server 2013 Integration

Before configuring the integration of Exchange Server 2013 and Lync Server 2013, you need to understand how the integration works. When you configure the integration, you are configuring Exchange Server 2013 Unified Messaging components that will enable Exchange Server 2013 to communicate with Lync Server 2013. Furthermore, you are configuring Lync Server 2013 components that enable Lync to communicate with Exchange. Two tools are provided to configure the required objects.



ExchUCUtil Script

To configure the integration of Lync and Exchange Unified Messaging, you must first run the ExchUCUtil script (ExchUCUtil.ps1) to configure the Exchange Server environment. The script does three things:

- It grants the Lync server accounts permission to read Exchange Unified Messaging AD DS objects, so that it can create contact objects for each auto attendant and subscriber access.
- It creates a UM IP gateway object for each Lync Server 2013 pool, and then associates the gateways to the UM SIP dial plans that you define for Lync Server 2013.
- It creates an UM hunt group for each UM IP gateway. The hunt group pilot identifier will be the name of the dial plan associated with the UM IP gateway.

You will run the ExchUCUtil script when you configure Exchange Unified Messaging integration initially with Lync Server 2013. You should run the script again whenever you create Exchange UM SIP dial plans that you will use to integrate with Lync Server, and whenever you add a new Lync 2013 server to the environment.

Exchange UM Integration Utility

The Exchange UM Integration Utility (OcsUMUtil.exe) is a tool that you will run on the Lync 2013 server. When you run the tool, you will create contact objects in Active Directory that Lync Server 2013 uses to link to the Exchange UM Auto Attendant and Subscriber Access numbers. When you run the tool, it will read the Exchange Unified Messaging configuration. You then can create the contact items in an existing organizational unit or you can create a new organizational unit. The contact items are assigned a SIP address and a phone number. By default, the subscriber access contact is assigned the phone number that

MCT USE ONLY. STUDENT USE PROHIBITED

you configured for Outlook Voice Access on the UM dial plan, and the auto attendant contact is assigned the phone number assigned to the UM auto attendant.

Note: In Exchange Server 2007 and Exchange Server 2010, the Exchange UM Integration Utility also verified that the Exchange UM SIP dial plan names match the corresponding Lync Server 2010 dial plans names. In Exchange Server 2010 SP1 and newer versions, the dial plan names do not need to match.

You will run the Exchange UM Integration Utility when you initially configure Exchange Unified Messaging integration with Lync Server 2013. You should run the tool again when you install new Exchange UM dial plans for Lync Server 2013, or when additional subscriber access and auto attendants are added to an existing Lync Server 2013–related dial plan. You can run this utility at any time to troubleshoot the contact item configuration.

Certificate Requirements

The SIP dial plan that you configure on the Exchange Servers must use mutual TLS encryption for all traffic. This means that you must install a certificate on all Exchange 2013 servers that will communicate with the Lync 2013 servers, as well as on the Lync 2013 servers. The certificates that you deploy on both sets of servers must be trusted by the other set of servers. You can configure certificates in several ways:

- Obtain certificates from a trusted public CA for both sets of servers. This will eliminate any trust issues.
- If you deploy an internal CA, you can obtain certificates for both sets of servers from the internal CA.
- If you are using self-signed certificates, you must import the certificates to the trusted root certification authority node on all other servers. We do not recommend this approach.

Note: Lync Server and Exchange Server integration do not support the use of wild card certificates.

Implementing Exchange Server UM Integration with Lync Server

To configure the integration of Exchange Server 2013 and Lync Server 2013, complete the following steps:

1. Install Lync Server in the same location as the Exchange 2013 Client Access servers and Mailbox servers. A fast LAN connection should connect the servers.
2. Configure the Enterprise Voice components on the Lync servers, including:
   a. PSTN connectivity. To provide full telephone functionality, the Lync servers must be able to send and receive calls from PSTN telephones.
   b. Dial plans. You will need to create dial plans for all internal users.
   c. Call routing rules. These rules define how calls are routed within the organization or to the PSTN.

Process for integrating Exchange Server Unified Messaging with Lync Server:

1. Install Lync Server 2013
2. Configure Enterprise Voice in Lync Server
3. Verify certificates
4. Configure a secure SIP URL dial plans
5. Add all Client Access and Mailbox servers to the SIP dial plan
6. Set the Unified Messaging services to dual mode
7. Run the ExchUCUtil.ps1 script on an Exchange Server
8. Run OcsUMUtil.exe on a Lync server
9. Enable users for Unified Messaging and Enterprise Voice

MCT USE ONLY. STUDENT USE PROHIBITED

   d. Normalization rules. These rules define how Lync will handle specific types of calls. For example, if you want users to be able to dial a five-digit extension to reach other internal users, you will need to create a normalization rule that translates the five-digit extension into the full phone number.

3. Verify that the infrastructure's servers trust the certificates installed on the Exchange and Lync servers.
4. Create and configure a SIP URI dial plan in Exchange 2013. You must configure the dial plan to use the SIP Secured or Secured setting to enforce mutual TLS.
5. Add all Client Access and Mailbox servers to the SIP dial plan. This will enable all Exchange servers to answer incoming calls from Lync Server.
6. Set the startup mode for the Unified Messaging services to Dual, and then restart the Microsoft Exchange Unified Messaging service on each Mailbox server, and the Microsoft Exchange Unified Messaging Call Router service on each Client Access server.
7. Run the ExchUCUtil.ps1 script from the <Exchange Installation folder>\Exchange Server\Script folder on any Exchange Server.
8. Run OcsUMUtil.exe from the %CommonProgramFiles%\Microsoft Lync Server 2013\Support folder on a Lync Server.
9. Enable your users for UM and Enterprise Voice. When you enable users for voice mail, create a SIP address for the users who will use Enterprise Voice. In most cases, this SIP address will be the same SIP address that will be used when a user is enabled for Enterprise Voice.

MCT USE ONLY. STUDENT USE PROHIBITED

Lab: Designing and Implementing Exchange Server 2013 Unified Messaging

Scenario

A. Datum is an engineering and manufacturing company. The organization is based in London, England, but is rapidly expanding the London location as well as into other countries. A. Datum has deployed the core functionality available in Exchange Server 2013, but is planning to implement additional features in Exchange Server. A. Datum is deploying Lync 2013 as a VOIP solution. To provide a full VOIP solution, A. Datum wants to deploy Exchange Server 2013 Unified Messaging.

Objectives

After completing this lab, you will be able to:

- Plan the Unified Messaging solution
- Install and configure Unified Messaging features
- Configure the integration of Unified Messaging and Lync Server 2013

Lab Setup

Estimated Time: 105 minutes

Virtual machines	**20342B-LON-DC1** **20342B-LON-CAS1** **20342B-LON-MBX1** **20342B-LON-CL1** **20342B-LON-CL2** **20342B-LON-LY1**
User Name	Adatum\Administrator
Password	Pa$$w0rd

Do not log on to LON-CL1 and LON-CL2 until directed to do so.

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

1. On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In Microsoft Hyper-V® Manager, click **20342B-LON-DC1**, and then in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Sign in by using the following credentials:
   - User name: **Adatum\Administrator**
   - Password: **Pa$$w0rd**
5. Repeat steps 2, 3, and 4 for 20342B-LON-CAS1, 20342B-LON-MBX1, and 20342B-LON-LY1.
6. For 20342B-LON-CL1 and 20342B-LON-CL2, repeat steps 1 through 3. Do not log on until directed to do so.

 **Note:** In some cases, messages sent in this lab may not be delivered immediately. You may notice that when you send messages, the messages stay in the Drafts folder in Outlook Web App.

MCT USE ONLY. STUDENT USE PROHIBITED

Use the following steps to troubleshoot mail flow if you experience this issue in this lab or in any other labs:

1. On LON-MBX1, open the Exchange Management Shell.
2. Type **Test-ServiceHealth**, and press Enter. Verify that all required services are running. If the services are not running, start them.
3. Type **Restart-Service MSExchangeSubmission**, and press Enter.
4. Type **Restart-Service MSExchangeDelivery**, and press Enter. Check to see if the message has been delivered.
5. If not, type **Restart-Service MSExchangeTransport**, and press Enter. Check to see if the message has been delivered.
6. If the messages are still not being delivered, restart the Microsoft Exchange Active Directory Topology service from the Services console. Restart all dependent services. Verify that all services set to automatic start are started. Check to see if the message has been delivered.

Exercise 1: Designing the Unified Messaging Implementation

Scenario

A. Datum is planning to deploy Unified Messaging on the Exchange 2013 servers in its London data center. The Lync Server 2013 deployment project already has deployed a Lync Server 2013 Front-End server and has configured the Enterprise Voice components.

You need to create a design for the Unified Messaging deployment. The project team has collected the following information and requirements:

- A single Lync 2013 server is being deployed. The Lync Server has an IP address of 172.16.0.25.
- Users in the London, Toronto, and Paris offices will be configured as Lync 2013 Enterprise Voice and Unified Messaging users. The phone numbers assigned to each location are:
  - London: +44 171 4444-xxxx with internal extensions configured as 1xxxx
  - Toronto: +1 416 5555-xxxx with internal extensions configured as 2xxxx.
  - Paris: +33 1 6666-xxxx with internal extensions configured as 3xxxx.
- The Lync Server deployment team has implemented dial plans that includes five-digit extensions for all offices.
- The Lync Server deployment team has implemented SBAs in the Paris and Toronto offices. These devices provide local VoIP connectivity in each office, as well as local PSTN connectivity. The devices connect the Lync 2013 server in London across a WAN.
- The last four digits in the extension must match the last four digits in the direct inward dialing (DID) telephone number.
- External users should be able to call a local number, such as (44) (171) 4444-9999, (1) (416) 5555-9999, (33) (1) 6666-9999, to reach the company phone directory. Users should be able to request service in English and French, and search for users in each of the three offices. Internally, the organization's phone directory should be accessible from all locations with the extension x9999, where x is the local extension prefix for each office.
- A. Datum employees should be able to call a local number (+44 171 4444-1111, +1 416 5555-1111, +33 1 6666-1111) when they are out of the office to check their emails and voice mails. They should be able to call x1111 from within each office to check their emails and voice mails.

MCT USE ONLY. STUDENT USE PROHIBITED

- All employees other than managers and members of the legal department are required to use a PIN of at least six characters, but the PINs do not need to be complex. All managers and members of the legal department must use complex eight-character PINs.

The main tasks for this exercise are as follows:

1. Create a design for the Exchange Server 2013 Unified Messaging components

2. Discuss your design with the rest of the class

▶ Task 1: Create a design for the Exchange Server 2013 Unified Messaging components

Review the information in the Exercise Scenario and answer the following questions:

1. What Exchange Server 2013 dial plans will you need to configure? How will you configure the dial plans?
2. How will you meet the requirement enabling external and internal users to reach the organization's telephone directory by dialing local or internal numbers?
3. How will you meet the requirement that users should be to get service in English or French?
4. How will you meet the requirement that users should be able to search for recipients in each office?
5. How will you meet the requirement enabling employees to access their email and voice mail by phone by dialing a local number or internal extension?
6. How will you meet the requirement for the different personal identification number (PIN) settings for different groups of employees?

▶ Task 2: Discuss your design with the rest of the class

As a group, discuss the proposed solutions from the students, to find the solution that best fits the requirements of A. Datum.

Results: After completing this exercise, you will have designed an Exchange Unified Messaging deployment.

Exercise 2: Configuring Unified Messaging Features

Scenario

Now that you have completed your design, the next step is to configure the Unified Messaging components on the Exchange 2013 servers, based on your design.

The main tasks for this exercise are as follows:

1. Configure Unified Messaging dial plans

2. Configure the UM IP Gateway

3. Review the default Unified Messaging hunt group

4. Configure Unified Messaging mailbox policies

5. Configure a Unified Messaging Auto Attendant

6. Configure Unified Messaging users

▶ Task 1: Configure Unified Messaging dial plans

1. On LON-CAS1, open Internet Explorer, and then connect to **https://LON-CAS1.adatum.com/ecp**.
2. Sign in as **Adatum\administrator** using the password **Pa$$w0rd**.

MCT USE ONLY. STUDENT USE PROHIBITED

3. Create a new UM dial plan with the following configuration:
   a. Name: **Lync-Dialplan**
   b. Extension length: **5**
   c. Dial plan type: **SIP URI**
   d. VoIP security mode: **Secured**
   e. Country code: **44**
4. Edit the Lync-Dialplan by changing the following settings:
   a. Outside line access code: **9**
   b. International access code: **00**.
   c. Outlook Voice Access numbers: +**4417144442000**
   d. dialing authorization: **Allow calls to any extension**
   e. transfer & search: **In the entire organization**

▶ Task 2: Configure the UM IP Gateway

1. On LON-CAS1, in the EAC, create a new UM IP gateway with the following settings:
   a. Name: **LON-UM-Gateway**
   b. Address: **172.16.0.40**
   c. UM dial plan: **Lync-Dialplan**

▶ Task 3: Review the default Unified Messaging hunt group

1. On LON-CAS1, open the Exchange Management Shell.
2. Use the **Get-UMHuntGroup** cmdlet to view all hunt groups.
3. View the detailed properties of the default hunt group.

▶ Task 4: Configure Unified Messaging mailbox policies

1. On LON-CAS1, use the **Get-UMMailboxPolicy** cmdlet to view all UM mailbox policies.
2. View the detailed properties of the default mailbox policy.
3. In the EAC, create a new UM Mailbox policy with the following settings:
   a. Name: **Managers-UMMailboxPolicy**
   b. Message text when a user is enabled for UM: **Your mailbox has been enabled for Unified Messaging**
   c. Minimum PIN length: **8**
   d. PIN recycle count: **8**
   e. Enforce PIN lifetime: **30**

▶ Task 5: Configure a Unified Messaging Auto Attendant

1. On LON-CAS1, in the EAC, create a new UM Auto Attendant with the following settings:
   a. Name: **Adatum-AutoAttendant**
   b. **Create this auto attendant as enabled**
   c. **Set this auto attendant to respond to voice commands**

MCT USE ONLY. STUDENT USE PROHIBITED

d. Access numbers: +**4417144449999**

▶ Task 6: Configure Unified Messaging users

1. On LON-CAS1, in the EAC, enable Kelly Rollin's mailbox for UM with the following settings:
   a. UM Mailbox policy: **Lync-Dialplan Default Policy**
   b. SIP address: **Kelly@adatum.com**.
   c. Extension number: **11006**
2. Enable Benno Kurmann's mailbox for UM with the following settings:
   a. UM Mailbox policy: **Managers-UMMailboxPolicy**
   b. SIP address: **Benno@adatum.com**.
   c. Extension number: **11005**.

Exercise 3: Configuring Unified Messaging Integration with Lync Server 2013

Scenario

A. Datum has deployed Lync Server 2013 as its VoIP solution. You now must configure the integration of Unified Messaging and Lync Server 2013, and then verify that the integration is configured correctly.

Lync is deployed in London data center, and it has only a London PSTN connection established via SIP Trunk. It is representing U.K., U.S., and French phone numbers.

If a customer place a call to the AutoAttendant, they must be able to connect a call to all three regions by using either the employee name or the given five-digit extension.

You need to configure the entire system integration between Exchange Server 2013 and Lync Server 2013.

The main tasks for this exercise are as follows:

1. Install a Certificate on the Mailbox Server
2. Configure the Unified Messaging Services To Use TLS
3. Configure the Autodiscover Service URI
4. Configure Exchange Server for Lync Server integration
5. Prepare Lync Server for Exchange Integration

▶ Task 1: Install a Certificate on the Mailbox Server

1. On LON-CAS1, create a folder named **E:\Certs**, and then share it with the default permissions.
2. In the EAC, request a new certificate for LON-MBX1.adatum.com, adding LON-MBX1.adatum.com, LON-MBX2, and LON-MBX2.adatum.com to the names assigned to the certificate.
3. Fill in the following information:
   a. Organization Name: **A. Datum**
   b. Department Name: **Messaging**
   c. Country/Region name: **United Kingdom**
   d. City/Locality: **London**
   e. State/Province: **EN**
4. Store the certificate request as **\\lon-CAS1\certs\MBXcertrequest.req**.

MCT USE ONLY. STUDENT USE PROHIBITED

5. Open the certificate request file in Notepad, and then copy the contents into the clipboard.
6. In Internet Explorer, connect to **http://lon-dc1.adatum.com/certsrv**.
7. Request a new certificate by using the advance certificate request and the certificate request file. Choose the Adatum Web certificate template.

Note: If you receive an error message that the certificate request was denied, restart the Active Directory Certificate Services service on LON-DC1, and then try the request again.

8. Save the certificate to **\\lon-CAS1\certs**.
9. On Internet Explorer, return to the EAC, and then complete the certificate request by using the file **\\lon-cas1\certs\certnew.cer**.

▶ Task 2: Configure the Unified Messaging Services To Use TLS

1. On LON-CAS1, open the Exchange Management Shell, and then use the Get-UMService to review the UM configuration for the Mailbox servers.
2. Use the following command to assign the Lync-Dialplan to both Mailbox servers.

```
Get-MailboxServer | Set-UMService –DialPlans Lync-Dialplan –UMStartupMode TLS
```

3. Use the following command to assign the Lync-Dialplan to both Client Access servers.

```
Get-ClientAccessServer | Set-UMCallRouterSettings –DialPlans Lync-Dialplan –
UMStartupMode TLS
```

4. To view the default UM call router settings, run the **Get-UMCallRouterSettings –Server lon-cas1.adatum.com** cmdlet.
5. On LON-CAS1, in the EAC, assign the LON-MBX1.adatum.com certificate to the UM service on LON-MBX1
6. Assign the Webmail.adatum.com certificate to the UM call router service on LON-CAS1.
7. On LON-CAS1, restart the Microsoft Exchange Unified Messaging Call Router service.
8. On LON-MBX1, restart the Microsoft Exchange Unified Messaging service.

Note: If you get an error message indicating that the service cannot be started, ignore this error for now.

▶ Task 3: Configure the Autodiscover Service URI

1. On LON-CAS1, in the Exchange Management Shell, use the following command to configure the internal uniform resource locator (URL) for the Autodiscover service.

```
Get-ClientAccessServer | Set-ClientAccessServer -AutoDiscoverServiceInternalUri
https://autodiscover.adatum.com/autodiscover/autodiscover.xml
```

2. Use NSLookup to verify that the autodiscover.adatum.com alias references 172.16.0.20.

▶ Task 4: Configure Exchange Server for Lync Server integration

1. On LON-CAS1, in the Exchange Management Shell, run the **ExchUCUtil.ps1** script from **C:\Program Files\Microsoft\Exchange Server\V15\Scripts**.

MCT USE ONLY. STUDENT USE PROHIBITED

2. Use the **Get-UMDialPlan** cmdlet to verify that a UM IP Gateway has been created, named LON-LY1, and associated with the dial plan Lync-Dialplan.

Note: If the Microsoft Exchange Unified Messaging service did not start previously, on LON-MBX1, in the Exchange Management Shell, type **Get-service msexchangeUM**, and then press Enter. If the service still shows as stopped, type **Start-service msexchangeUM**, and then press Enter. If the service still does not start, wait a few minutes, and then try starting the service again. It can take several minutes for the service to start.

▶ Task 5: Prepare Lync Server for Exchange Integration

1. On LON-LY1, at a command prompt, run the **ocsumutil.exe** command from **C:\Program Files\Common Files\ Microsoft Lync Server 2013\Support**.
2. Load the data from the Active Directory forest.
3. Add a new contact to the Lync-Dialplan with the following settings:
   a. Create a new OU named UMIntegration.
   b. Name: **Lync Subscriber Access**
   c. SIP Address: **sip:Lync-SA**
   d. Phone number: default
   e. Contact type: **Subscriber Access**
4. Add another contact to the dial plan with the following settings:
   a. Name: **Lync-Autoattendant**
   b. SIP Address: **sip:Adatum-AA**
   c. Phone number: default
   d. Contact type: **Auto-Attendant**

Note: The previous two tasks create two contact items in the organizational unit (OU) that you specified. The first contact routes messages to Outlook Voice Access, and the second contact routes messages to the auto-attendant.

Results: After you have configured the Exchange 2013 Unified Messaging integration with Lync 2013, you will be able to leave voice messages for UM- enabled Exchange users and use the AutoAttendant via Lync 2013 to connect a SIP call to Lync Enterprise Voice-enabled users.

MCT USE ONLY. STUDENT USE PROHIBITED

Exercise 4: Verify Unified Messaging Functionality

Scenario

Now that you have configured Exchange 2013 Unified Messaging and configured the integration with Lync Server 2013, you need to verify the deployment. You will enable two users in Lync or Enterprise Voice, and then verify that the users can call each other using the Lync client. You then will verify that users can use the Exchange Server 2013 Unified Messaging features.

Note: This lab exercise requires an audio headset for each student. If a headset is not available, you will not be able to complete this exercise. If you have a headset available, plug the headset in now.

The main tasks for this exercise are as follows:

1. Enable Enterprise Voice for Lync users
2. Verify Enterprise Voice functionality
3. Verify Unified Messaging integration
4. To prepare for the next module

▶ Task 1: Enable Enterprise Voice for Lync users

1. On LON-LY1, open the Lync Server 2013 Control Panel, and then sign in as **Administrator** using the password **Pa$$w0rd**
2. Activate Benno Kurmann for Lync, and then for Enterprise Voice using the following settings:
   a. Pool: **LON-LY1.ADATUM.COM**.
   b. SIP URI: **Use user's email address**.
   c. Telephony: **Enterprise Voice**.
   d. Line URI: **tel:+4417144441005;ext=11005**.
3. Activate Kelly Rollin for Lync and for Enterprise Voice using the following settings:
   a. Pool: **LON-LY1.ADATUM.COM**.
   b. SIP URI: **Use user's email address**.
   c. Telephony: **Enterprise Voice**.
   d. Line URI: **tel:+4417144441006;ext=11006**.

▶ Task 2: Verify Enterprise Voice functionality

1. On LON-LY1, open the **Lync Server Management Shell**.
2. Use the **Test-CsExUMConnectivity –TargetFqdn lon-ly1.adatum.com –UserSipAddress kelly@adatum.com** command to verify that Kelly's account is enabled for Lync and UM.
3. Use the same command to check Benno's account.

Note: If you get an error message when you run the **Test-CsExUMConnectivity** command, type **Update-CsAddressBook** at the command prompt, and then press Enter. Wait a few minutes, and then run the **Test-CsExUMConnectivity** commands again.

4. On the host machine, open **D:\Program Files\Microsoft Learning\20342\Drives\LON-CL1.rdp**.

MCT USE ONLY. STUDENT USE PROHIBITED

5. Sign in as **Adatum\Benno** using the password **Pa$$w0rd**.
6. Open Outlook 2013, and then configure the user profile.
7. Open Lync 2013, and verify that Benno is signed into Lync automatically.
8. On the host machine, open **D:\Program Files\Microsoft Learning\20342\Drives\LON-CL2.rdp**.
9. Sign in as **Adatum\Kelly** using the password **Pa$$w0rd**.
10. Open Outlook 2013, and then configure the user profile.
11. Open Lync 2013, and then verify that Kelly is signed into Lync automatically.
12. Verify that the users can communicate with each other by using instant messaging.
13. On LON-CL1, use Lync to call the phone number **+4417144441006**.
14. Verify that the call is connected, and that the users can talk to each other.
15. On LON-CL1, use Lync to call the extension **11006**.
16. Verify that the call is connected, and that the users can talk to each other.

▶ Task 3: Verify Unified Messaging integration

1. In the LON-CL1 Remote Desktop Connection window, in the Lync client, call Kelly's extension (11006).
2. Do not answer the call on LON-CL2.
3. Wait for the call to go to voice mail, and then leave a message for Kelly.
4. Verify that Kelly receives an email with the voice mail. Verify the message transcription, and then verify that you can play the message.
5. In the LON-CL1 Remote Desktop Connection window, in the Lync client, call auto attendant extension (19999).
6. When prompted, request to call Kelly Rollin.
7. Verify that the phone rings for Kelly.
8. In the LON-CL1 Remote Desktop Connection window, in Outlook, create a new Meeting request with Kelly for later today.
9. In the Lync client, call the Outlook Voice Access extension (12000).
10. Provide the PIN assigned to the mailbox, and then follow the voice prompts to change the PIN.
11. Follow the voice prompts to record your name and to record a personal greeting.
12. When the mailbox is prepared, say "calendar", and then listen to the Calendar options.
13. Follow the voice prompts to return to the main menu.
14. Access the mailbox by saying "email".
15. Listen to the reading of the email message, and then hang up..

▶ Task 4: To prepare for the next module

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

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20342B-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.

MCT USE ONLY. STUDENT USE PROHIBITED

4. Repeat steps 2 and 3 for **20342B-LON-CAS1**, **20342B-LON-MBX1, 20342B-LON-LY1, 20342B-LON-CL1**, and **20342B-LON-CL2**.
5. In Hyper-V Manager, click **20342B-LON-DC1**, and in the Actions pane, click **Start**.
6. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
7. Sign in using the following credentials:
   a. User name: **Adatum\Administrator**
   b. Password: **Pa$$w0rd**
8. Repeat steps 5 to 7 for **20342B-LON-MBX1**, **20342B-LON-CAS1**, and **20342B-LON-CL1**.

Results: After completing this exercise, you will have configured two users for Enterprise Voice in Lync 2013, verified the Enterprise Voice functionality, and verified the integration between Exchange 2013 Unified Messaging and Lync 2013.

Question: Why did you need to install certificates on the Exchange 2013 Mailbox servers in this lab?

Question: How would the lab have changed if you would have implemented the full design that you created in Exercise 1?

MCT USE ONLY. STUDENT USE PROHIBITED

Module Review and Takeaways

Best Practices

- Before implementing Unified Messaging, verify the state of your network. If your network has limited available bandwidth, or high latency, the user experience with Unified Messaging will not be positive.
- For the broadest compatibility, leave Exchange Unified Messaging configured to record voice mail in MP3 format. For maximum compression, configure Exchange UM to record voice mail in WMA format.

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
The Lync 2013 client displays warnings that connectivity to Exchange is limited.	
Calls to Exchange Unified Messaging cannot be completed.	

Review Question(s)

Question: What are the various levels of VoIP security available in Exchange Unified Messaging? If Lync Server 2013 is encrypting both signaling and media, what is the appropriate setting for VoIP security in Exchange Unified Messaging?

MCT USE ONLY. STUDENT USE PROHIBITED

MCT USE ONLY. STUDENT USE PROHIBITED

Module 5

Designing and Implementing Message Transport Security

Contents:

Module Overview 5-1

Lesson 1: Overview of Messaging Policy and Compliance Requirements 5-2

Lesson 2: Designing and Implementing Transport Compliance 5-5

Lesson 3: Designing and Implementing AD RMS Integration with Exchange Server 2013 5-15

Lab: Designing and Implementing Message Transport Security 5-25

Module Review and Takeaways 5-31

Module Overview

Microsoft® Exchange Server 2013 provides a wide range of messaging compliance features that you can use for more than just simple messaging and calendaring. You can also use messaging compliance features to control message transport, to implement journaling, to manage messages, and to apply Active Directory® Rights Management Services (AD RMS).

Objectives

After completing this module, you will be able to:

- Describe messaging policy and compliance requirements.
- Design and implement transport compliance.
- Design and implement AD RMS integrations with Exchange Server 2013.

MCT USE ONLY. STUDENT USE PROHIBITED

Lesson 1
Overview of Messaging Policy and Compliance Requirements

Email has become a reliable and ubiquitous communication medium for employees of organizations of all sizes. Messaging stores and mailboxes have become repositories of valuable data. Organizations need to formulate messaging policies that dictate the fair use of their messaging systems, provide user guidelines for how to act on the policies, and, where required, provide details about the types of communication that may not be allowed. Organizations must also create policies to manage the email life cycle. This includes retaining messages for an appropriate length of time based on business, legal, and regulatory requirements, preserving email records for litigation and investigation purposes, and being prepared to search for and provide the required email records to fulfill eDiscovery requests.

Leakage of sensitive information such as intellectual property, trade secrets, business plans, and personally identifiable information collected or handled by your organization must also be prevented. This lesson provides an overview of the options available in Exchange Server 2013 that helps you to comply with your organization's messaging policy and compliance requirements.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe legal messaging compliance requirements.
- Describe corporate messaging compliance requirements.
- Describe options for enforcing messaging policy and compliance.

Identifying Legal and Corporate Compliance Requirements

The most important step for an organization to take when making email communications legally compliant is to create a corporate email policy. Compliance or risk officers should create corporate messaging policies that address compliance requirements based on any regulations and laws regarding their industry.

- Ensure compliance with legal requirements
- Add a general disclaimer to all external messages
- Create an ethics firewall
- Block messages with confidential content
- Restrict user action on messages
- Block messages to specified recipients

In addition to compliance with legislation, organizations also need to reduce corporate liability and financial loss resulting from improper email usage or lack of retention policies for corporate communications. Organizations need to create and enforce corporate email and messaging policies that address areas of potential liability, such as disclosure or transfer of intellectual property, discrimination, harassment, and client/attorney privilege. A strong corporate email policy that is designed from best practices, legal guidelines, and due diligence can help reduce corporate civil and criminal liability and financial loss.

Organizations that consider compliance when they plan their information technology infrastructure, including their email infrastructure, can supply the required documentation on demand with less effort. Organizations that plan their information technology infrastructure with regulatory compliance in mind can also comply with other regulatory requirements more easily.

MCT USE ONLY. STUDENT USE PROHIBITED

The following list provides an overview of possible compliance requirements of an organization:

- Add disclaimers to messages. Many organizations require Exchange Server to add specific, pre-written text to all messages that are sent from the organization to external recipients. Instead of relying on individual users to add the disclaimer, you can centrally implement and enforce the use of disclaimers by using transport rules.
- Restrict users from sending messages to other recipients. You can use transport rules and moderated recipients to control which users can send messages to other recipients. For example, a transport rule can prevent a user from sending messages outside the organization. Alternatively, you can restrict which messages are sent to a distribution group by implementing moderated recipients.
- Block or retain messages with specific content. You can use transport rules to block or retain messages that have specific content. For example, you can create a transport rule that deletes all messages that have the text string "guaranteed return," or you can forward all those messages to a mailbox so that a designated user can review them.
- Restrict what recipients can do with a message. You can use AD RMS to limit what recipients can do with a message. For example, if a message is intended for the company lawyer, you can prevent that message from being forwarded to other recipients.
- Block messages to a specific email domain. You can use transport rules to block messages that are addressed to a specific email domain. For example, you can use a transport rule to delete all messages that are addressed to the contoso.com domain, or, if there are multiple recipients, you can remove all recipients who are in the contoso.com domain.

  Question: How does your organization meet its compliance requirements today?

Options for Enforcing Messaging Policy and Compliance

Many organizations today are required by law, regulatory requirements, or company policies to apply messaging policies that limit the interaction between senders and recipients from inside and outside the organization. These requirements include the following:

- Transport rules
- Message classifications
- Moderated recipients
- AD RMS integration

Additional Message and Compliance Features:
- Messaging records management
- In-Place Archiving
- In-Place Hold
- In-Place eDiscovery

- *Transport rules*. You can use transport rules to apply policies to messages in the transport pipeline. The transport rules can apply actions such as redirecting, adding recipients, rejecting, silently deleting, or rights-protecting messages. The rules act on messages that match the conditions and that do not meet any of the exceptions defined in the transport rules.
- *Message classifications*. Message classifications add metadata to a message. The metadata typically describes how the message should be used and who should have access to it. After you classify a message, you can use transport rules to manage it in a specific way.
- *Moderated recipients*. With moderated recipients, you can require that messages sent to specific recipients be approved by one or more moderators. You can configure any type of recipient as a moderated recipient.

MCT USE ONLY STUDENT USE PROHIBITED

- *AD RMS integration*. You can use Information Rights Management (IRM) in AD RMS to prevent an authorized recipient of an IRM-protected email from forwarding, modifying, printing, or saving the content.

Additional Message and Compliance Features of Exchange Server 2013

There are additional options that you can use to comply with business, legal, and regulatory requirements. The following options can be implemented in your organization. They are not part of this module, but they are described in detail in the next module of this course.

- Messaging records management (MRM). In Exchange Server 2013, MRM is accomplished by using retention tags and retention policies.
- In-Place Archiving. In Exchange Server 2013, In-Place Archiving provides users an alternate storage location for historical messaging data.
- In-Place Hold. In Exchange Server 2013, In-Place Hold places user mailboxes on hold and preserves mailbox items for the period of time specified by the hold.
- In-Place eDiscovery. In-Place eDiscovery in Exchange Server 2013 can help you perform discovery searches for relevant content within mailboxes.

MCT USE ONLY. STUDENT USE PROHIBITED

Lesson 2
Designing and Implementing Transport Compliance

You can use transport compliance to control messages as they are transported through the Exchange server organization. You can control which users can send messages, which users can receive messages, and whether messages are modified as they are transported. To implement transport compliance, you can use transport rules, message classifications, and message moderation.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the options for implementing transport compliance.
- Describe the purpose and functionality of transport rules.
- Plan the implementation of transport rules.
- Configure transport rules.
- Plan message classifications.
- Describe message moderation.
- Explain the purpose and functionality of journaling.
- Configure message moderation and journaling.

Options for Implementing Transport Compliance

The following scenario highlights how Exchange Server 2013 can help organizations comply comprehensively and cost-effectively with legislation and company policies that address email and messaging communications.

Exchange Server 2013 increases capabilities for enforcing data privacy in e-mail messages, with features like:
- Transport Rules
- Information Rights Management

Healthcare: Secure Delivery and HIPAA Compliance

The Health Insurance Portability and Accountability Act (HIPAA), which was passed into law in 1996, mandates how private information must be protected to prevent unauthorized disclosure. Healthcare providers and companies that handle health data for patients or employees are required to implement strictly enforced policies regarding data management, including data in email messages.

Exchange Server 2013 makes it easier for healthcare and related companies to comply because it increases capabilities for enforcing data privacy in email messages. Healthcare organizations can use the features of Exchange Server 2013 to enforce company email policies automatically that help prevent the unauthorized disclosure of private data by using transport rules that apply handling instructions, encrypting email content, and auditing configuration changes.

For example, by using transport rules, an email administrator can create a rule that searches the subject and content of every email message that is sent, looking for social security numbers (SSNs). If a user creates an email message that includes an SSN and tries to send it to a recipient who is either outside of

MCT USE ONLY. STUDENT USE PROHIBITED

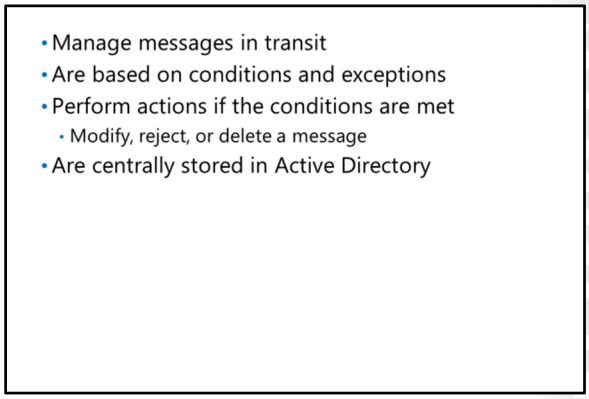
the organization or who is not authorized to receive confidential information, the email is not sent and the sender receives a pre-configured error message.

Enforcement of Data Loss Protection (DLP) policy is another feature of Exchange Server 2013. DLP policies are simple packages that contain sets of conditions, which are made up of transport rules, actions, and exceptions, that you create in the Exchange Administration Center (EAC) and then activate to filter email messages.

Also, healthcare companies can use the IRM features in AD RMS to help automatically protect email messages that contain confidential information and that are sent internally. This approach helps protect sensitive information that is in transit without requiring any client software or end user training. Companies that need to comply with HIPAA can now use the built-in capabilities of Exchange Server 2013 to comply with the security requirements of HIPAA more easily.

Understanding Transport Rules

With transport rules, you can help ensure that every message sent within an organization is checked for specific conditions, and that, if these conditions are met, a defined action is taken on a message before it reaches the recipients. You can use transport rules to apply messaging policies to email messages, to help secure messages, to help protect messaging systems, and to help prevent information leakage.



Transport rules are similar to Inbox rules in the Microsoft Office Outlook® messaging client. The difference between transport rules and Inbox rules is that transport rules are executed when the message is in transit from the sender to the recipients, and Inbox rules are set up in Outlook. Additionally, transport rules have more conditions, exceptions, and actions, so they are more flexible.

Components of Transport Rules

A transport rule consists of the following components:

- *Conditions*. Transport rule conditions specify the characteristics of messages that you want to apply a transport rule action to. Conditions consist of one or more *predicates*, which specify the parts of a message that are examined. Some predicates examine message fields or email headers, such as the name and address of the sender and recipient. Other predicates examine message characteristics, such as the subject, body, attachments, size, and classification. Most predicates require that you specify a comparison operator—such as equals, doesn't equal, or contains—and a value.
- *Exceptions*. Exceptions are based on the same predicates that are used to build transport rule conditions. However, unlike conditions, exceptions identify messages to which transport rule actions are not applied. Exceptions override conditions, preventing actions from being applied to an email message, even if the message matches all of the conditions.
- *Actions*. Actions are applied to messages that match the conditions and that do not match any exception defined in the transport rule. Transport rules have many actions available, including rejecting, deleting, or redirecting messages, adding additional recipients, adding prefixes in the message subject, and inserting disclaimers in the message body.

MCT USE ONLY. STUDENT USE PROHIBITED

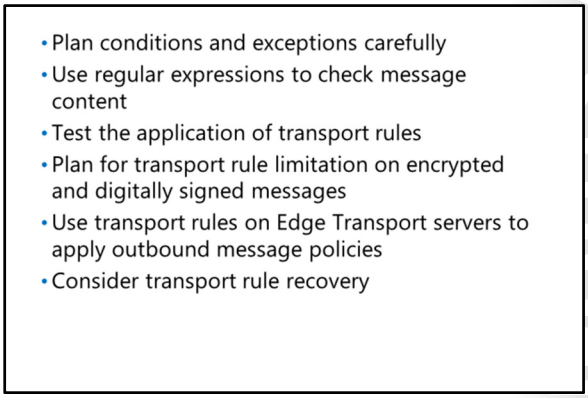
How Transport Rules Work

The workflow for transport rules follows this general pattern:

1. In either the EAC the Exchange Management Shell, you create transport rules that fit your needs. These rules are stored in Active Directory Domain Services (AD DS), so you need to create them only once.
2. While a message goes through the transport pipeline, the transport agent is invoked. The transport rules are processed by a special transport agent, called the transport rule agent.
3. The message is scanned by the transport rule agent, and if the message fits the conditions specified in the transport rule, the specified action is taken on that message.

Planning Transport Rules

Transport rules provide you with a large amount of control over messaging in an Exchange Server organization. Transport rules should be carefully planned to ensure that they behave as intended. Otherwise, you might accidentally delete messages or deliver them to unintended recipients.



You should consider the following when planning transport rules:

- *Plan conditions and exceptions carefully*. Transport rule conditions and exceptions define which messages are affected by the transport rule. If you implement the rules incorrectly, you may unintentionally modify or delete messages.
- *Use regular expressions to check message contents*. Use regular expressions to simplify the list of terms if you include a text string in a condition. You can use one regular expression, rather than a list of variations on the same word.
- *Test the application of transport rules*. Test new transport rules to ensure that they behave as intended. This is important because a new transport rule might conflict with existing transport rules.
- *Plan for transport rule limitations on encrypted and digitally signed messages*. You can use the AD RMS integration with Exchange Server 2013 to implement transport rules and messaging policies if you are using AD RMS Information Rights Management encryption to protect messages. Encryption through other mechanisms may prevent the application of transport rules or records management. For example, Exchange Server may not be able to scan encrypted messages for the text string that is specified in a transport rule. Additionally, antivirus scanners cannot scan messages that have encrypted attachments.
- Use transport rules on Microsoft Exchange Edge Transport servers to apply outbound message policies for delivery to external recipients. The Transport service on the Mailbox Server applies transport rules, which results in unnecessary processing for outbound messages. You can offload this processing to Edge Transport servers instead. Additionally, in some cases, messages from external organizations may be relayed through Microsoft Exchange Edge Transport servers directly to another messaging organization, and not be processed by the Transport service.
- *Consider transport rule recovery*. Deleted transport rules are not easily recoverable. Transport rules are stored in Active Directory Domain Services (AD DS), and restoring rules from AD DS is a complex

MCT USE ONLY. STUDENT USE PROHIBITED

process. But documented transport rules are easy to recreate, and you can export transport rules to backup files by using the **Export-TransportRuleCollection** cmdlet. However, when you import transport rules to a Hub Transport server, the server replaces all of the existing transport rules for the organization.

Demonstration: Configuring Transport Rules

In this demonstration, you will see how to configure a disclaimer for all messages that are sent inside an organization. Then you will see how to configure a transport rule to block all messages that are sent with the word "Important" in the **Subject** field.

Demonstration Steps

Create a new transport rule

1. Create a new transport rule.
2. Append a disclaimer to all messages that are sent inside your organization.
3. Set the disclaimer text as: Disclaimer set on message through Transport rule.
4. Select **Ignore** as the failback action when the Transport rule cannot apply.

Configure a transport rule that blocks a message containing the word *Important* in the subject

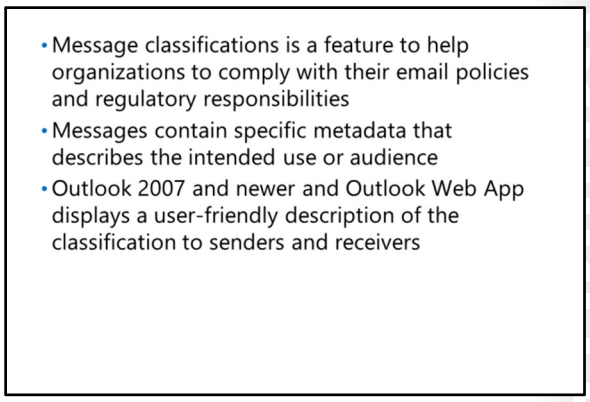
1. Create a new transport rule.
2. Apply this rule to messages where the subject contains the word *Important*.
3. Select Reject the message with the explanation as the action.

Test both new transport rules in Microsoft Outlook Web App

1. Sign in to Outlook Web App as user Ed.
2. Create an email message and send it to Ankur with the subject **Normal internal message**.
3. Create another email message and send it to Ankur with then subject **Important internal message**.
4. Check what happens after you send the messages.

Understanding Message Classification

Message classification is a feature that is available in Exchange Server 2007 or later and Microsoft Office Outlook 2007 and newer versions. This feature is intended to help organizations comply with their email policies and regulatory responsibilities by enabling users to mark messages with custom classifications. When a message is classified, specific metadata that describes the intended use or audience is added the message. Outlook 2007 and newer versions or Microsoft Outlook Web App use this metadata to display a user-friendly description of the



MCT USE ONLY. STUDENT USE PROHIBITED

classification to senders and recipients of the message. You can also configure transport rules that are triggered based on the metadata attached to a classified message.

The following three message classifications are enabled in Exchange Server 2013 by default:

- **Attachment Removed**. This classification notifies recipients when attachments have been removed from the message.
- **Originator Requested Alternate Recipient Mail**. This classification notifies recipients that the message has been redirected from delivery to the original addressed recipient.
- **Partner Mail**. This classification notifies recipients that the message was encrypted and delivered through a secure connector.

These message classifications are only used by Exchange Servers, and users cannot add these default classifications to messages.

During the installation of Exchange Server 2013, these message classifications are informational only. They are not associated with any transport rule and only provide additional information about a message to the message recipients. However, you can create transport rules based on message classifications.

When you create message classifications, you can configure the following settings:

- Display name. This property specifies the display name of the message classification that is displayed in the Permission menu in Outlook 2007 and Outlook Web App. Users of Outlook and Outlook Web App can use this property to select the appropriate message classification before sending a message. The display name is also visible to message recipients in the InfoBar of an Outlook message. The parameter name for this property is DisplayName.
- Sender description. This property provides information about the intent of a message classification to the sender. Outlook and Outlook Web App users see the description when they select the classification. The description should give users a clear understanding of the purpose of the message. The parameter name for this property is SenderDescription.
- Recipient description. This property provides information about the intent of a message classification to the recipient. The text that you enter for this field is displays to Outlook and Outlook Web App users when they receive a message with this message classification. The parameter name for this property is RecipientDescription.
- Locale. This field specifies a culture code to create a locale-specific version of the message classification. The parameter name for this property is Locale.

Users can apply the preceding message classifications to messages they send after you have enabled Outlook 2007 and newer versions to accept the default message classifications. Senders see the sender description in the InfoBar in Outlook 2007 and newer. By using the Exchange Management Shell, you can customize the sender description for each message classification and locale.

Message classifications can be used in two ways:

- A message classification can be manually added by the message sender before the message is sent.
- A message classification can be added as the result of a transport rule. For example, when the Attachment Filter agent removes an attachment from a message, the Attachment Removed message classification is attached to the message. When the sender receives the message, Outlook 2007 and newer versions and Outlook Web App display an explanation of why the attachment was removed in the recipient description in the InfoBar. You can customize the recipient description.

You must deploy the message classification configuration files and create an Outlook registry key on the end-users' computers before users of Outlook 2007 and newer versions can set and view message

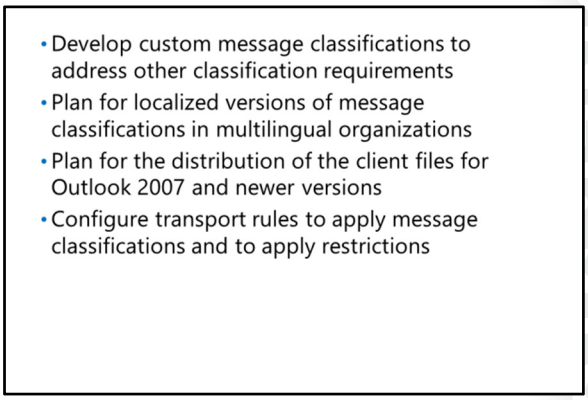
MCT USE ONLY. STUDENT USE PROHIBITED

classifications. The Outlook message classification templates are .xml files that you must generate after you create and configure the message classifications.

You manage all message classifications by using the message classification cmdlets in the Exchange Management Shell. You can bind message classifications to transport rules by using the Exchange Management Shell or the EAC

Planning Message Classification

Message classifications organize messages and provide additional information about them. The classifications also can trigger transport rules. You should consider the following when planning message classifications:



- *Develop custom message classifications*. Most organizations require custom message classifications to meet their specific needs. To do this, determine which classifications your organization requires, and define the sender and recipient descriptions that appear when the message is classified.
- *Plan for localized versions of message classifications.* Each message classification can include alternate sender and recipient descriptions associated with different locales. For multilingual organizations, create localized versions of message classification descriptions so that recipients can read the message classifications in their preferred language.
- *Configure client file distribution for Outlook 2007 and newer.* These clients do not use message classifications by default, so you must configure them to do so. To configure Outlook 2007 and newer, distribute an XML file that contains the message classifications. Redistribute this XML file each time you modify message classifications. You also need to configure registry entries. Outlook Web App supports message classifications by default.
- *Configure transport rules*. You can use transport rules to control how Exchange Server transports classified messages based on company polices. For example, you can create a transport rule that prevents messages with the Company Internal classification from being delivered outside the organization. Additionally, you can use transport rules to apply message classifications based on message content, senders, or recipients. For example, you can automatically assign the Legal classification to any message that arrives from an external lawyer.

MCT USE ONLY. STUDENT USE PROHIBITED

Understanding Message Moderation

Message moderation in Exchange Server 2013 requires that all messages to specific recipients be approved by a dedicated moderator or group of moderators. For example, this feature enables you to control messages that are sent to large distribution lists but that have unimportant content, such as "I am leaving the company." Often, these emails are of no interest and just fill up mailboxes. You also might want to control who can send messages to executive mailboxes in the organization.

- Every recipient can be moderated
- One or more moderators required
- Messages can be approved or rejected
- Messages that are not approved or rejected within a given time expire

Components of moderated transport
- Categorizer
- Mailbox Transport Service
- Information Assistant
- Arbitration Mailbox

With message moderation in Exchange Server 2013, you can assign a moderator for a recipient. The designated moderator approves or rejects messages sent to the recipient.

Components of Moderated Transport

Moderated transport consists of the following components:

- *Categorizer*. The categorizer in the Transport service on a Mailbox server initiates the approval process. If the categorizer detects a moderated recipient while processing a message, it reroutes the message to the arbitration mailbox.
- *Mailbox Transport service*. The Mailbox Transport service on a Mailbox server processes the messages that the categorizer marks for moderation. If the Mailbox Transport service encounters such a message, it delivers the original message to the arbitration mailbox and sends approval requests to the moderators. When a moderator responds with a decision, the Mailbox Transport service marks that decision on the message that is stored in the arbitration mailbox. If the Information Assistant submits an approved message again, the Mailbox Transport service removes the approval workflow wrappers so that the message that is delivered is identical to the original message that the sender submitted.
- *Information Assistant*. The Information Assistant process in the Mailbox Transport service monitors the arbitration mailbox. The Information Assistant resubmits any approved messages to the Transport service on a Mailbox server for delivery to the intended recipients, and it deletes rejected messages. The Information Assistant also sends rejection notifications to the sender. In addition, it cleans up the arbitration mailbox by deleting any stale or orphaned messages from the arbitration mailbox. For example, if a moderator simply deletes an approval request instead of making a decision, the Information Assistant removes the corresponding message waiting for approval in the arbitration mailbox.
- *Arbitration mailbox*. The arbitration mailbox stores the original message that is awaiting approval. By default, one arbitration mailbox is created for moderated transport during setup. It is used for all moderated recipients. You can add additional arbitration mailboxes for load balancing purposes. If you use multiple arbitration mailboxes, you need to specify which mailbox to use for each moderated recipient.

How Message Moderation Works

The following steps describe the process flow for a message sent to a moderated recipient:

1. A user sends a message to a moderated recipient.

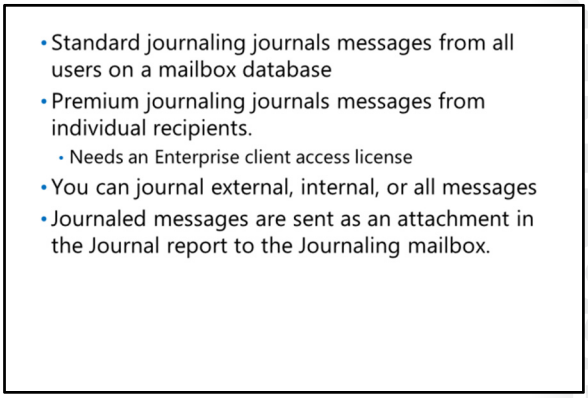
MCT USE ONLY. STUDENT USE PROHIBITED

2. In the Transport service of the Mailbox server, the categorizer intercepts the message, marks it for moderation, and then reroutes it to the Mailbox Transport service on the Mailbox server where the arbitration mailbox is stored.
3. The moderator receives an approval request from the Mailbox Transport service.
4. The moderator either accepts or rejects the message by using buttons included in the message.
5. The Mailbox Transport service marks the moderator's decision on the original message stored in the arbitration mailbox.
6. The Information Assistant in the Mailbox Transport service now reads the approval status on the message in the arbitration mailbox, and then it processes the message depending on the decision of the moderator, as follows:
    a. If the message is approved, the Information Assistant resubmits the message to the Transport service on a Mailbox server. The message is delivered to the recipient.
    b. If the message is rejected, the Information Assistant deletes the message from the arbitration mailbox, and the sender is notified that the message is rejected.
    c. If the message is not approved or rejected within five days, the Information Assistant deletes the message from the arbitration mailbox, and the sender is notified that the message expired.

Understanding Message Journaling

Journaling and archiving are often confused with one another. Both concern data retention, but the purpose behind them differs, as follows:

- *Journaling*. The process of recording all inbound and outbound email communication in an organization to meet its retention or archival strategy.
- *Archiving*. The process of managing the size of an organization's data by removing it from its primary storage location and storing it elsewhere.



Journaling does not capture existing messages that are stored in user's mailboxes. It captures new messages in transit between the sender and the recipient. Journaling also does not capture items like contacts, calendar items, and tasks.

Components of Message Journaling

Message journaling consists of the following components:

- Journaling agent. The Journaling agent processes messages on Mailbox servers.
- Journal rules. Journal rules have the following key components:
- Journal rule scope. The Journal rule scope describes which messages to journal, specifically, internal messages only, external messages only, or all messages.
- Journal recipient. A Journal recipient can be an Exchange mailbox, a distribution group, an email user, or a contact. All messages sent to or from the journaling recipient are journaled.
- Journaling mailbox. A Journal mailbox is used only for collecting journal reports.

MCT USE ONLY. STUDENT USE PROHIBITED

Journaling Options in Exchange Server 2013

Exchange Server 2013 includes the following journaling options:

- *Standard journaling*. Standard journaling is configured in the properties of the mailbox database and journals all messages that are sent to or from any mailbox that is stored on that mailbox database.
- *Premium journaling*. Premium journaling allows you to specify individual recipients or members of a distribution group to journal. For Premium journaling, an Enterprise client access license must be available.

How Message Journaling Works

The following steps describe the process flow for a journaled recipient:

1. A user creates a new message and sends it.
2. The Journaling agent on the Mailbox server processes the message, based on the following options:
    - Journal rule scope.
    - Journaling recipient.
3. A Journal report is sent to the Journaling mailbox, including the original message as an attachment.
4. The message is sent to the recipient.

Best Practices for Journaling

Here are some best practices for working with journal mailboxes:

- Configure the journal mailbox to accept messages only from the Microsoft Exchange recipient and not accept messages sent by unauthenticated users.
- Disable storage quota limits for the journaling mailbox or enable a Prohibit send and receive quota. If you disable storage quota limits, it is recommended to monitor the size of the mailbox.
- Hide the journal mailbox from the global address list (GAL).

Demonstration: Configuring Message Moderation and Journaling

In this demonstration, you will see how to configure moderation settings for the distribution group named Managers. Then you will see how to configure a journal rule to journal all messages that are sent and received in the organization.

Demonstration Steps

Configure moderation settings for a distribution group

1. Enable moderation for the distribution group Managers.
2. Designate Ed as the moderator.
3. Allow user Erwin to bypass moderation.
4. Notify senders in your organization if their messages are not approved.

Create a journal rule to journal all messages that are sent and received in the organization

1. Create a new journal rule named **ADatum Journaling**.
2. Select **All messages** as the journal scope.

MCT USE ONLY. STUDENT USE PROHIBITED

3. Select **Apply to all messages** as the journal recipient.
4. Select the mailbox **Journal** as the journal mailbox.

MCT USE ONLY. STUDENT USE PROHIBITED

Lesson 3

Designing and Implementing AD RMS Integration with Exchange Server 2013

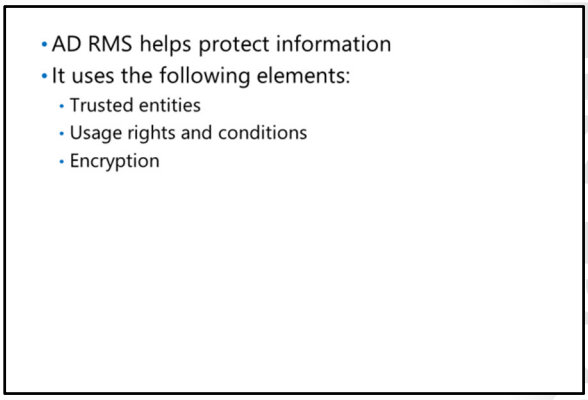
You can integrate Exchange Server 2013 with AD RMS to provide additional protection for messages. As part of planning AD RMS integration, consider how best to protect messages and how external recipients can access AD RMS to decrypt and view messages.

Lesson Objectives

After completing this lesson, you will be able to:

- Explain the features and functionality of AD RMS.
- Explain how AD RMS works.
- Describe the options for integrating AD RMS with Exchange Server 2013.
- Plan AD RMS integration.
- Configure transport protection rules.
- Configure Outlook protection rules.
- Configure transport and journal report decryption.
- Implement and manage AD RMS integration.
- Design AD RMS integration with external users.

What Is AD RMS?



AD RMS is an information protection technology that works with AD RMS–enabled applications to help safeguard digital information from unauthorized use, both online and offline, and inside and outside of a firewall. AD RMS is designed for organizations that need to protect sensitive and proprietary information, such as financial reports, product specifications, customer data, and confidential email messages. AD RMS uses persistent usage policies (also known as *usage rights and conditions*), which remain with the information no matter where it is moved. This also enables usage policies to be enforced after the information is accessed by an authorized recipient, both online and offline, and inside and outside of the organization. AD RMS has the following elements:

- *Trusted entities*. Organizations can specify the entities that are trusted participants in an AD RMS system. These can include individuals, groups of users, computers, and applications. By establishing trusted entities, AD RMS can help protect information by enabling access only by trusted participants.
- *Usage rights and conditions*. Organizations and individuals can assign usage rights and conditions that define how a specific trusted entity can use rights-protected content. Examples of usage rights are permission to read, copy, print, save, forward, and edit. Usage rights can be accompanied by

MCT USE ONLY. STUDENT USE PROHIBITED

conditions, such as when those rights expire. Organizations can exclude applications and entities from accessing the rights-protected content.

- *Encryption*. Encryption is the process by which data is locked by using electronic keys. AD RMS encrypts information, which makes access conditional on the successful validation of the trusted entities. After information is locked, only trusted entities that are granted usage rights under the specified conditions, if any, can unlock or decrypt the information in an AD RMS–enabled application or browser. The defined usage rights and conditions are then enforced by the application.

Although Exchange Server 2013 includes solutions to help protect access to data, those solutions have the following limitations:

- *Transport Layer Security (TLS)*. TLS helps protect a Simple Mail Transfer Protocol (SMTP) message only between two SMTP hosts. TLS does not protect at the message-level or information that is at rest. Messages in the sender's and recipient's mailboxes remain unprotected. TLS is a transport layer technology; it cannot control what the recipient does with the message.
- *Email encryption*. The user decides whether to encrypt a message. There are additional costs of a public key infrastructure deployment with the overhead of certificate management for users and protection of private keys. After a message is encrypted, there is no control over what the recipient can do with the information. Decrypted information can be copied, printed, or forwarded. By default, saved attachments are not protected.

Understanding AD RMS

AD RMS encompasses all of the server and client technologies that are required to support information protection by using rights management in an organization. If you use an AD RMS infrastructure, you can help protect the information in an organization by using the following client and server components to both publish and consume rights-protected content:

AD RMS uses the following components:
- AD RMS client. The client requests licenses and enforces assigned rights protection at the document level on files and messages
- AD RMS server. The server administers account certification, licensing, and publishing services that use Active Directory, and it assists clients in locating these services

- *AD RMS client*. The client requests licenses and enforces assigned rights protection at the document level on files and messages. Beginning with Windows Vista® and Windows Server® 2008, the AD RMS client is integrated into the operating system. In earlier versions such as Windows 2000 Server, Windows XP, and Windows Server 2003, the client needs to be installed separately. . The AD RMS client supports x86, x64, and Itanium.
- *AD RMS server*. The server manages account certification, licensing, and publishing services that use AD DS, and it assists clients in locating these services.

How AD RMS Clients Work

AD RMS clients request and acquire new licenses for protecting content when they publish content. These licenses are determined by the usage rights and conditions that the publisher allows for the content the license will protect. The following process describes how AD RMS clients work:

1. If a document is authored and rights protection is selected, the AD RMS client acquires a client licensor certificate. The client licensor certificate enables AD RMS to protect the content.
2. AD RMS then uses this client licensor certificate to encrypt the document.

MCT USE ONLY. STUDENT USE PROHIBITED

3. AD RMS creates and signs a publishing license (PL), and then it binds a copy of the PL to the encrypted content.
4. When the recipient wants to access the rights-protected content, they first need to use a rights-enabled application like Microsoft Office to request and acquire an end-user license for the content.
5. The AD RMS client must determine whether the recipient of the content conforms to any policies specified in the publishing license that protects the content.
6. If the user is eligible to access the content, the AD RMS client helps ensure that the user honors the conditions indicated in the end-user license, which might restrict certain actions.

How AD RMS Servers Work

AD RMS servers are implemented as a set of web service components that run on Microsoft Internet Information Services (IIS). These components work in conjunction with AD DS and the Microsoft SQL Server® database software. The following table lists the components of an AD RMS server.

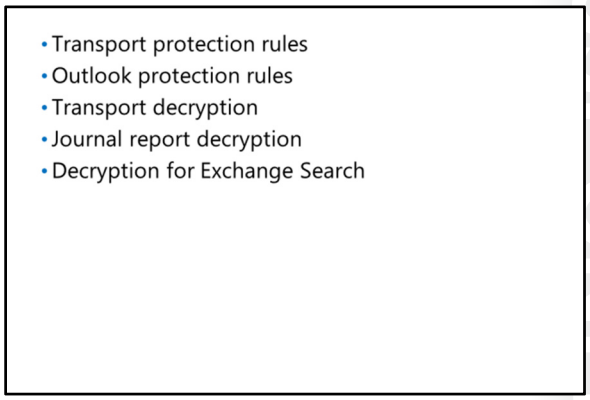
Server component	Description
Administration web service	This service hosts the Administration website, which enables you to manage AD RMS. The service runs on root certification servers and on licensing servers.
Account certification	This service creates machine certificates that identify computers in the AD RMS certificate hierarchy and it creates a rights account certificate that associates users with specific computers. This service runs on the root certification server.
Licensing	This service issues end-user licenses. The service runs on root certification servers and on licensing servers.
Publishing	This service creates the issuance licenses that define the policy that can be enumerated in an end-user license. The publishing service runs in root certification servers and on licensing servers.
Precertification	This service enables a server to request a rights account certificate on behalf of a user. The service runs on root certification servers and on licensing servers.
Service locator	This service provides the URL of the account certification, licensing, and publishing services to AD DS so that they can be discovered by AD RMS clients. The service runs on root certification servers and on licensing servers.

Pre-licensing

Exchange Server 2013 automatically attaches a pre-license that is provided by AD RMS to help protect messages. This makes it is possible to view messages and attachments that are protected by IRM features. With this license, the client does not need to make repeated trips to the AD RMS server to retrieve a use license, and users can view IRM-protected messages and attachments offline. They can also view IRM-protected messages in Outlook Web App. Pre-licensing is enabled by default if you enable IRM.

MCT USE ONLY. STUDENT USE PROHIBITED

Options for Integrating IRM Features in AD RMS with Exchange Server 2013



AD RMS is an IRM) technology in the Windows Server 2008 operating system and newer that applies persistent protection to messages and attachments in Exchange Server 2013. Your organization and its users can control the email rights that recipients have with the IRM features that are in Exchange Server 2013. IRM helps you allow or restrict recipient actions, such as forwarding a message to other recipients, printing a message or attachment, or extracting messages or attachment content by copying and pasting. Users can use IRM to help protect messages in Microsoft Office Outlook 2003 and newer. Microsoft Outlook 2010 and newer support the use of Outlook protection rules. Previous versions of Outlook do not support protection rules. Mobile devices that support the Microsoft Exchange ActiveSync® protocol version 14.1, including Windows Mobile® software powered devices, also support IRM. IRM is also supported for Outlook Web App.

IRM features in Exchange Server 2013 support Microsoft Office file formats. To use IRM with other file formats, you must deploy custom protectors.

The options for integrating AD RMS into Exchange Server 2013 are:

- *Outlook users*. To help protect messages with IRM, Outlook users can use AD RMS rights policy templates that are available to the users.
- *Outlook Web App users*. If IRM is enabled in Outlook Web App, users can protect messages they send with IRM, and they can view IRM-protected messages that they receive.
- *Windows Mobile-powered devices and Exchange ActiveSync devices*. Starting with Microsoft Exchange Server 2010 with Service Pack 1 (SP1), you can enable IRM in Exchange ActiveSync to allow users of Exchange ActiveSync devices, which includes Windows Mobile powered devices, to view, reply, forward, and create IRM-protected messages.
- *Outlook 2010 and newer*. In Outlook 2010 and newer, you can create Outlook protection rules to help protect messages automatically with IRM. These protection rules are automatically deployed to Outlook 2010 clients. Messages are IRM-protected before they leave the Outlook client. This protection is also applied to any attachments using supported file formats. When you create Outlook protection rules on an Exchange Server 2013 server, the rules are automatically distributed to Outlook 2010 by using Exchange Web Services. For Outlook 2010 to apply the rule, the AD RMS rights policy template you specify must be available on the user's computers.
- *Mailbox server*. On Exchange Server 2013 Mailbox servers, transport protection rules can be applied automatically to help protect messages with IRM. AD RMS uses XML-based policy templates to allow compatible IRM-enabled applications to apply consistent protection policies. In Windows Server 2008 and newer, the AD RMS server exposes a Web service that can be used to enumerate and acquire templates. Exchange Server 2013 ships with the Do Not Forward template. When the Do Not Forward template is applied to a message, only the recipients addressed in the message can decrypt the message. The recipients cannot forward the message to anyone else, copy content from the message, or print the message.

MCT USE ONLY. STUDENT USE PROHIBITED

Decrypting IRM-Protected Messages to Enforce Messaging Policies

If you need to enforce messaging policies for regulatory compliance, you should be able to access the encrypted message content. To meet eDiscovery requirements related to litigation, regulatory audits, or internal investigations, you must also be able to search encrypted messages. Exchange Server 2013 includes the following IRM features for this purpose:

- Transport decryption.
- Journal report decryption.
- IRM decryption for Exchange Search.

Decryption is not enabled per default. Before you can encrypt the message content, you must provide Exchange Server 2013 servers the right to decrypt content protected by your AD RMS server. This is done by adding the Federation mailbox to the super users group configured on the AD RMS cluster in your organization. Then you can use the different cmdlets to configure content decryption.

Transport Decryption

In Exchange Server 2013, IRM-protected messages are decrypted by the Decryption agent, which is a built-in agent. The Decryption agent decrypts the following types of IRM-protected messages:

- Messages encrypted manually by the user in Outlook Web App.
- Messages encrypted manually by the user in Outlook 2010.
- Messages encrypted automatically by Outlook protection rules in Exchange Server 2013 and Outlook 2010.

It is important to know that only messages that are IRM-protected by the AD RMS server in your organization are decrypted by the Decryption agent.

Transport decryption is performed on the first Exchange Server 2013 Transport service that handles a message in an Active Directory forest. If a message is transferred to a Transport service in another Active Directory forest, the message is decrypted again. After decryption, unencrypted content is available to other transport agents on that server. For example, the Transport Rules agent on a Transport service can inspect message content and apply transport rules. Any actions specified in the rule, such as applying a disclaimer or modifying the message in any other way, can be taken on the unencrypted message. Third-party transport agents, such as antivirus scanners, can scan the message for viruses and malware. After other transport agents have inspected the message and possibly made modifications to it, it is encrypted again with the same user rights that it had before being decrypted by the Decryption agent. The same message is not decrypted again by the Transport service on other Mailbox servers in the organization.

Messages decrypted by the Decryption agent do not leave the Transport service without being encrypted again. If a transient error is returned when decrypting or encrypting the message, the Transport service retries the operation twice. After the third failure, the error is treated as a permanent error. If any permanent errors occur, including when transient errors are treated as permanent errors after retries, the Transport service treats them as follows:

- If the permanent error occurs during decryption, a non-delivery report (NDR) is sent only if transport decryption is set to Mandatory, and the encrypted message is sent with the NDR.
- If the permanent error occurs during re-encryption, an NDR is always sent without the decrypted message.

Also, it is important to know that any custom or third-party agents installed on a Transport service have access to the decrypted message. You must consider the behavior of such transport agents. We recommend that you test all custom and third-party transport agents thoroughly before you deploy them in a production environment.

MCT USE ONLY. STUDENT USE PROHIBITED

Configure Transport Decryption

Transport decryption is configured by using the **Set-IRMConfiguration** cmdlet in the Exchange Management Shell. However, before you configure transport decryption, you must provide Exchange Server 2013 servers the right to decrypt content protected by your AD RMS server.

Exchange Server 2013 allows two different settings when enabling transport decryption:

- Mandatory. When transport decryption is set to Mandatory, the Decryption agent rejects the message and returns an NDR to the sender if a permanent error is returned when decrypting a message. An organization that does not want a message to be delivered if it cannot be successfully decrypted and actions such as antivirus scanning and transport rules are applied must choose this setting.
- Optional. When transport decryption is set to Optional, the Decryption agent uses a best-effort approach. Messages that can be decrypted are decrypted, but messages with a permanent error on decryption are also delivered. An organization that prioritizes message delivery over messaging policy must use this setting.

Journal Report Decryption

In Exchange Server 2013, IRM-protected messages are decrypted by the Journal Report Decryption agent, which is a built-in agent. The Decryption agent decrypts the following types of IRM-protected messages:

- Messages encrypted manually by the user in Outlook Web App.
- Messages encrypted manually by the user in Outlook 2010.
- Messages encrypted automatically by Outlook protection rules in Outlook 2010.
- Messages encrypted automatically in transit by using transport protection rules.

It is important to know that only messages that are IRM-protected by the AD RMS server in your organization are decrypted by the Journal Report Decryption agent. The agent does not decrypt an attachment if an IRM-protected file is attached to an unprotected message or if the attachment is not protected at the same time as the message, and therefore does not have the same use license.

Journal report decryption is configured by using the **Set-IRMConfiguration** cmdlet in the Exchange Management Shell. However, before you configure journal report decryption, you must assign Exchange Server 2013 servers the permissions to decrypt content that is IRM-protected by your AD RMS server.

After you enable journal report decryption, the journaling mailbox may contain journal reports with sensitive information in an unencrypted form. As a best practice, we recommend that access to the journaling mailbox be monitored closely and restricted only to authorized individuals. This is a best-practice even if you are not using IRM protection for email.

IRM Decryption for Exchange Search

Messages protected using IRM are indexed by Exchange Search and are therefore included in the search results if they match query parameters. Messages must be protected by using an AD RMS cluster in the same Active Directory forest as the Mailbox server.

When members of the Discovery Management role group perform an In-Place eDiscovery search, IRM-protected messages are returned in the search results and copied to the Discovery mailbox specified in the search. Furthermore, members of the Discovery Management role group can use Outlook Web App to access the IRM-protected messages that were copied to the Discovery mailbox as a result of the discovery search.

Please note that member of the Discovery Management role group cannot access IRM-protected messages exported from a Discovery mailbox to another mailbox or to a .pst file. IRM-protected messages in a Discovery mailbox can be accessed only by using Outlook Web App.

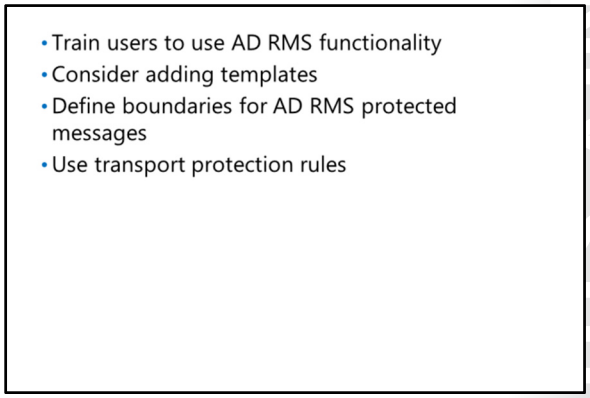
MCT USE ONLY. STUDENT USE PROHIBITED

When Exchange Search fails to index an IRM-protected message, either due to a decryption failure or because IRM is disabled, the protected message is not added to the list of failed items. If you select the option to include unsearchable items in search results, the results may not include IRM-protected messages that could not be decrypted.

You do not have to configure the IRM to allow Exchange search to index IRM-protected messages, because this is enabled by default.

Planning AD RMS Integration

The basic requirement for integrating AD RMS integration is to implement the AD RMS server. The AD RMS server generates the certificates that help protect message content and can specify restrictions.



You must perform the following tasks after configuring AD RMS:

- Train users to use the AD RMS functionality. Users have the option to apply AD RMS templates to messages. However, they are not likely to use this functionality unless you train them on how to use the templates.
- Consider adding templates. Exchange Server 2013 comes with one template, the **Do Not Forward** template. This template is useful, but you may need additional templates that prevent message modification, printing, saving, and copying. You can create additional RMS templates on the AD RMS server in your organization.
- Define the boundaries for AD RMS-protected messages. To decrypt and view protected messages, clients must be able to access the AD RMS server. Within your organization, it is easy to provide clients with access to the AD RMS server. However, if you allow AD RMS-protected messages outside of the organization, you also need to provide external users with access to your AD RMS server. If you do not coordinate external access to your AD RMS server, AD RMS-protected messages sent outside your organization will not be protected.
- Use transport protection rules to protect messages regardless of the client. Depending on the client software, users may not be able to apply AD RMS templates. To help protect messages regardless of the client software, implement transport protection rules that protect messages at the Hub Transport server level.

Demonstration: Configuring Transport Protection Rules

In this demonstration, you will see how to configure a transport protection rule that helps protect messages containing the word *Confidential* in the message Subject.

Demonstration Steps

1. Create a new transport protection rule with the name **ADatum Transport Protection Rule**.
2. Specify the condition to apply this rule if the subject or body includes the word **Confidential**.

MCT USE ONLY. STUDENT USE PROHIBITED

3. Ensure that rights protection is applied to the outgoing messages and that the message cannot be forwarded.
4. On LON-CL1, signed in as Ed, test the new transport protection rule. Create an email with the subject Confidential message, and then send it to Ankur.
5. On LON-CAS1, in Outlook Web App, verify that Ankur cannot forward the message.

Demonstration: Configuring Outlook Protection Rules

In this demonstration, you will see how to configure an Outlook protection rule that helps protect messages that are sent to the **Managers** distribution group and that have the AD RMS template **Do not forward**.

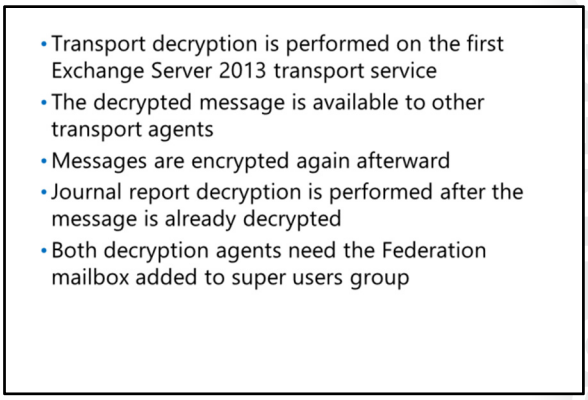
Demonstration Steps

1. Create an Outlook protection rule that helps protect messages sent to the **Managers** distribution group and that have the AD RMS template **Do not forward**.

Configuring Transport and Journal Report Decryption

You can configure transport decryption with the **Set-IRMConfiguration** cmdlet. The following two options are available:

- *Mandatory*. The decryption agent rejects the message and returns a non-delivery report NDR to the sender if a permanent error occurs.
- *Optional*. This setting uses a best-effort approach. This means that messages that can be decrypted are decrypted, but messages that cannot be decrypted are delivered.



Journal report decryption is performed by the Journal Report Decryption agent, which is a compliance-focused transport agent. Messages that are protected in transit by transport protection rules are already encrypted by the encryption agent before they get to the Journal Report Decryption agent. The Journal Report Decryption agent decrypts these messages.

You can also configure journal report decryption with the **Set-IRMConfiguration** cmdlet.

Before you can configure transport or journal report decryption, you must add the Federation mailbox, a system mailbox created during the Exchange Server 2013 setup, to the super users group that is configured on your organization's AD RMS cluster.

The AD RMS super user group is a special group that has full control over all rights-protected content managed by the cluster. Its members have full owner rights in all user licenses that are issued by the AD RMS cluster on which the super users group is configured. This means that members of this group can decrypt any rights-protected content file and remove rights-protection from it when appropriate.

The super users group is not enabled by default. When you enable the Super Users setting in the AD RMS console, you can specify an AD DS universal group as the super users group for AD RMS. The group must exist in the same forest as the AD RMS installation. Any user accounts that are members of the group that

MCT USE ONLY STUDENT USE PROHIBITED

you specify as the AD RMS super users group are automatically granted the permissions of the super users group.

You can configure a mail-enabled distribution group as a super users group in AD RMS. Members of the distribution group are granted an owner use license when they request a license from the AD RMS cluster. This allows them to decrypt all RMS-**protected** content published by that cluster. Whether you use an existing distribution group or create a distribution group and configure it as the super users group in AD RMS, we recommend that you exclusively dedicate the distribution group to this purpose and configure the appropriate settings to approve, audit, and monitor membership changes.

To add the Federation mailbox to a distribution group, perform the following steps:

1. Create a new distribution group dedicated to the use as an AD RMS super user group.
2. Add the user FederatedEmail.4c1f4d8b-8179-4148-93bf-00a95fa1e042 the new group. The mailbox is a system mailbox and therefore not visible in the EAC. You have to use the Shell to add it to the group.

To use AD RMS to set up a super user group, perform the following steps:

1. Open the Active Directory Rights Management Services console and expand the AD RMS cluster.
2. In the console tree, expand Security Policies and then click **Super Users**.
3. In the action pane, click **Enable Super Users**.
4. In the results pane, click **Change Super User Group** to open the Super Users property sheet.
5. In the Super user group box, type the email address of the distribution group you created in the previous procedure.

Implementing and Managing AD RMS Integration

Consider the following when implementing AD RMS integration:

- Provide Outlook Web App for external users
- Windows Phone 7 can read protected messages
- Develop a plan for distributing custom AD RMS templates
- Ensure that only trusted users can access the journal mailbox
- Develop a communication plan for incoming users
- Monitor the performance impact of encryption on Transport servers

- Provide Outlook Web App for external users. External users can use Outlook Web App to view protected messages in a web browser. You need to create user accounts for the external users, but you do not need to provide external access to your AD RMS server. The Client Access server hosting Outlook Web App communicates with the AD RMS server instead. By contrast, Outlook Anywhere requires the client to communicate directly with the AD RMS server.
- Clients of the Windows® Phone operating system can access protected email messages. If a Windows Phone client accesses a protected message, Exchange Server determines whether the user has access to the file, when allowed, it decrypts the file, and then it sends the file to Windows Phone. The user permissions are also transmitted to Windows Phone. You cannot create a protected message from Windows Phone.
- Develop a plan for distributing AD RMS templates. AD RMS templates must be distributed to clients so that the clients can use them. To automate template distribution to clients, you can use the Windows Vista operating system with SP1 or newer or you can use Windows Server 2008. By default,

MCT USE ONLY. STUDENT USE PROHIBITED

these tools distribute templates every 30 days. You can also copy AD RMS templates to clients as part of a Group Policy Object (GPO).

- Ensure that only trusted users can access the journal mailbox. If journal report decryption is disabled. Exchange Server stores all journaled content in an unencrypted format. This configuration means that anyone who can access the journal mailbox can read the messages. If encrypted messages contain confidential information, you should increase security on the journal mailbox.
- Develop a communication plan for users. AD RMS is a powerful tool for managing email usage, but you must teach users how to use AD RMS.
- Monitor the performance impact of encryption on Mailbox servers Transport service. Transport protection rules, transport decryption, and journal report decryption run on a Mailbox servers Transport service to encrypt or decrypt messages. Encryption and decryption are processor-intensive tasks that may cause performance issues on the Mailbox server. This is particularly true if the server processes many messages.

Planning AD RMS Integration for External Users

AD RMS integration for external users is more complex than when you simply restrict an AD RMS deployment to your own organization. Before you integrate AD RMS with external organizations, consider the following:

- Can you create external user accounts in the Active Directory forest?
- Have the external organizations deployed AD RMS?
- Do you need to enable AD RMS integration for all users in the external organizations?
- Have the external organizations deployed Active Directory Federation Services (AD FS)?

To integrate AD RMS with external organizations, do one of the following:

- Deploy an AD RMS server that is accessible from the Internet
- Configure trusted user or publishing domains
- Configure AD RMS integration with Windows Live ID
- Configure a federated trust by using AD FS

The following options are available for integrating AD RMS with external organizations:

- *Deploy an AD RMS server that is accessible from the Internet*. If your AD RMS server is accessible from the Internet, external users can communicate with the AD RMS server to obtain the necessary license certificates. This arrangement does not require the external organization to implement AD RMS, but it does require you either to create external user accounts in your Active Directory forest, or to create a separate forest with an AD RMS trust.
- *Configure trusted user or publishing domains*. You can use both trusted user and trusted publishing domains if the external organization has enabled AD RMS. With these two integration methods, users in one organization can access content that is protected by AD RMS in the other organization.
- *Configure AD RMS integration with the Windows Live® ID network of Internet services*. Configure a trust with Windows Live ID to allow protected content to be sent to any user who has a Windows Live ID. This option is suitable only for a small number of users, and it does not allow the external user to create protected content.
- *Configure a federated trust by using AD FS*. With this option, external clients contact the AD RMS server in your organization, but AD FS performs authentication. If you use this option, you do not need to create external user accounts in your Active Directory forest.

MCT USE ONLY. STUDENT USE PROHIBITED

Lab: Designing and Implementing Message Transport Security

Scenario

Several departments at A. Datum work with highly confidential information. It is critical that this information be secured and that it not leave the organization, except under approved circumstances. To meet these requirements, the security and compliance teams at A. Datum have developed new policies that define requirements for messaging security. You must create a design that meets these requirements and then implement and validate the changes on the Exchange Server 2013 deployment.

Objectives

After completing this lab, you will be able to:

- Plan a message transport implementation.
- Implement message transport security.
- Implement AD RMS and Exchange Server integration.

Lab Setup

Estimated Time: 75 minutes

Virtual machines	**20342B-LON-DC1** **20342B-LON-CAS1** **20342B-LON-MBX1** **20342B-LON-CL1**
User Name	Adatum\Administrator
Password	Pa$$w0rd

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

1. On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In Microsoft Hyper-V® Manager, click **20342B-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Sign in using the following credentials:
    a. User name: **Adatum\Administrator**
    b. Password: **Pa$$w0rd**
5. Repeat steps 2 to 4 for **20342B-LON-CAS1**, and **20342B-LON-MBX1**.
6. Start **20342B-LON-CL1**, but do not sign in to this virtual machine until directed to do so.

MCT USE ONLY. STUDENT USE PROHIBITED

Exercise 1: Planning a Message Transport Implementation

Scenario

The security and compliance teams at A. Datum have identified the following requirements for the messaging deployment:

- All email messages sent to users outside the organization must contain a disclaimer approved by the legal department.
- Messages sent from the info@adatum.com shared mailbox must contain a different disclaimer approved by the legal department.
- Copies of all messages that the Research team sends to both internal and external users must be retained in a mailbox that only the Auditing team can access.
- A message that contains characters that look like customer numbers can be sent to the Internet only if a customer service manager approves them.
- A message that contains the word *confidential* in the subject or body can be delivered only to users who have a mailbox on the A. Datum Exchange servers.
- A message that contains the word *private* in the subject must be encrypted and protected throughout the message delivery.
- A message to the AllCompany distribution group can be sent only if a member of the management team approves it.

The main tasks for this exercise are as follows:

1. Plan a message transport plan
2. Discuss your design with the class

▶ Task 1: Plan a message transport plan

Review the information in the Exercise Scenario and answer the following questions:

1. **Question:** Do you need transport rules in order to meet the requirements? If so, how many transport rules do you need and how will you implement them?
2. **Question:** Do you need journaling? If so, how will you implement it?
3. **Question:** Do you need recipient moderation? If so, how will you implement it?
4. **Question:** How can you protect messages during the message delivery? Is IRM an option? If so, which features can you use to meet the requirements?

▶ Task 2: Discuss your design with the class

1. Discuss the proposed solutions from the students all together to find the solution that best fits the requirements of A. Datum.

Results: After this exercise, you should have created a message transport plan.

MCT USE ONLY. STUDENT USE PROHIBITED

Exercise 2: Implementing Message Transport Security

Scenario

Based on a design made in the previous exercise, implement the appropriate technologies to address the security policy requirements.

The main tasks for this exercise are as follows:

1. Configure the required transport rules
2. Configure the required message moderation settings
3. Configure the required journal rules
4. Validate the message transport security configuration

▶ Task 1: Configure the required transport rules

1. On LON-CAS1, open Exchange admin center and create a transport rule named **ADatum External Disclaimer**, which applies to all recipients outside of the organization and that has the following text:

   This message contains confidential information and is intended only for the individual named. If you are not the named addressee, you should not disseminate, distribute, or copy this email. Please notify the sender immediately by email if you have received this email by mistake and delete this email from your system.

2. Select Reject as the failback action.
3. Create a transport rule named **ADatum Info Disclaimer**, which applies to all messages that are sent from the shared mailbox info@adatum.com and that has the following text:

   This message is sent on behalf of the Information Department of A. Datum and is intended for internal recipients of A. Datum only. If you are not the intended recipient, you are notified that disclosing, copying, distributing, or taking any action in reliance on the contents of this information is strictly prohibited.

4. Select Reject as the failback action.
5. Next, use the Exchange Management Shell to create a transport rule named ADatum Customer Approval, which applies to all messages that are sent to recipients where the subject or body contains customer numbers in the format \d\d\d\d(-|.)\d\d\d. These messages must be approved by the customer manager Benno before they are sent.
6. Use the Exchange Management Shell to create a transport rule named **ADatum Internal Confidential**, which applies to all messages where the subject or body contains the word Confidential and the recipients are outside the organization. These messages are rejected with the explanation *You are not allowed to send confidential messages outside the organization.*

▶ Task 2: Configure the required message moderation settings

1. Configure the distribution group **AllCompany** to be moderated by user Aiden, who is a member of the Managers group.
2. Configure the moderation to notify senders in your organization if their messages aren't approved.

▶ Task 3: Configure the required journal rules

1. Create a journal rule named **Research Journal Rule**, which journals all messages sent by the Research distribution group to internal and external users into the Journal mailbox, which only the Managers team can access.

MCT USE ONLY. STUDENT USE PROHIBITED

2. On LON-DC1, disable the Journal user account.

▶ Task 4: Validate the message transport security configuration

1. On LON-CAS1, validate the transport rule **ADatum Info Disclaimer**, which appends a different disclaimer to all messages sent from the shared mailbox info@adatum.com. Sign in to Outlook Web App as **Adatum\Info**, create an email message and send it to **Ed@adatum.com**.
2. Sign in to Outlook Web App as **Adatum\Ed** and verify that the message is delivered with the correct disclaimer.
3. Validate the transport rule **ADatum Customer Approval**, which requires that the messages that appear to contain customer numbers be approved by the customer manager Benno. Signed in as **Adatum\Ed**, create an email message that contains customer numbers in the format 1234-567 or 1234.567, and then send the message to **Adam@adatum.com**.
4. Sign in to LON-CL1 as **Adatum\Benno** and open Outlook 2013 to verify that you have a message from Ed that is waiting for approval. Reject the message.
5. Sign out of LON-CL1.
6. Validate the transport rule **ADatum Internal Confidential**, which rejects messages that have the word Confidential in the subject or body, if the recipients are outside the organization. Signed in to Outlook Web App as **Adatum\Ed**, create an email message that has the word Confidential in the subject or body, and then send it to **Troy@treyresearch.net**.
7. Validate that messages sent to the **AllCompany** distribution group are redirected to Aidan, who is the moderator of the group. Signed in as **Adatum\Ed**, create an email message and send it to the **AllCompany** distribution group.
8. Sign in to LON-CL1 as **Adatum\Aidan** and open Outlook and verify that the message from Ed is received and waiting for approval.
9. Sign out of LON-CL1.
10. Validate that all messages sent to members of the Research group are journaled into the Journal mailbox. On LON-CL1, signed in as **Adatum\Benno**, create an email message and send it to **Chloe@adatum.com**.
11. Sign out of LON-CL1.
12. Sign in on LON-CL1 as **ADatum\Aidan** with the password **Pa$$w0rd**.
13. In Outlook, add the Journal mailbox to Aidan's account.
14. Verify that the Journal mailbox is accessible, and then check for the journaled message sent from Benno to Chloe

Results: After this exercise, you should have implemented message transport security.

MCT USE ONLY. STUDENT USE PROHIBITED

Exercise 3: Implementing AD RMS and Exchange Server Integration

Scenario

A. Datum also wants to evaluate integration between Exchange Server 2013 and AD RMS. Every user in the organization must have the option to set IRM permissions on new emails in Outlook Web App that prevent the recipient of the message from forwarding it. Messages sent from members of the Manager group must be encrypted automatically with the IRM permission **Do not forward**. And finally, you need to enable journal report decryption in the organization.

The main tasks for this exercise are as follows:

1. Configure AD RMS integration
2. Configure the required transport protection rules
3. Configure journal decryption
4. Validate the AD RMS integration
5. To prepare for the next module

▶ Task 1: Configure AD RMS integration

1. Sign in to LON-CAS1 as **Adatum\administrator** with the password **Pa$$w0rd**.
2. From the Start screen, open the Exchange Management Shell.
3. Type the following command to enable IRM for internal messages, and then press Enter:

```
Set-IRMConfiguration –InternalLicensingEnabled $True
```

4. Next create a mail-enabled distribution group named **ADRMSSuperUser** that acts as the super users group in the AD RMS cluster, and then add the **FederatedEmail.4c1f4d8b-8179-4148-93bf-00a95fa1e042** system mailbox as a member. This is necessary to enable the following IRM features:
   - IRM in Outlook Web App.
   - IRM in Exchange ActiveSync.
   - Journal report decryption.
   - Transport decryption.
5. Sign in to LON-DC1, configure AD RMS to set up a super user group by using the distribution group that you just created.
6. On LON-CAS1, enable transport decryption, so that messages that cannot be decrypted are rejected and an NDR is returned to the sender.
7. On LON-CAS1, enable IRM on the Client Access servers so they can use IRM for Outlook Web App and for Exchange ActiveSync.
8. On LON-DC1, grant the Exchange Servers group and the AD RMS Service Group **Read & Execute** and **Read** permissions to the **c:\Inetpub\wwwroot_wmcs\Certification\ ServerCertification.asmx** file.

▶ Task 2: Configure the required transport protection rules

1. Create a transport protection rule that encrypts all messages that are sent from members of the Managers group and that have the IRM permission **Do not forward**.

MCT USE ONLY. STUDENT USE PROHIBITED

▶ Task 3: Configure journal decryption

1. Enable journal report decryption to allow the journaling agent to attach a decrypted copy of a rights-protected message to the journal report.

▶ Task 4: Validate the AD RMS integration

1. On LON-CL1, send an email from Aidan, who is a member of the **Managers** group, to Ed Meadows.
2. On LON-CAS1, log on to Outlook Web App as Ed. Check the received email message from Aidan and try to forward it.

▶ Task 5: To prepare for the next module

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

1. On the host computer, start the Microsoft Hyper-V® Manager.
2. In the **Virtual Machines** list, right-click **20342B-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 to 3 for **20342B-LON-CAS1**, **20342B-LON-MBX1** and **20342B-LON-CL1**.
5. In Hyper-V Manager, click **20342B-LON-DC1**, and in the Actions pane, click **Start**.
6. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
7. Sign in using the following credentials:
   a. User name: **Adatum\Administrator**
   b. Password: **Pa$$w0rd**
8. Repeat steps 5 to 7 for **20342B-LON-MBX1, 20342B-LON-MBX2,** and **20342B-LON-CAS1**.

Results: After this exercise, you should have implemented AD RMS integration in Microsoft® Exchange Server 2013.

Question: Before you enable journal decryption, what do you need to do?

Question: How can you check whether IRM is enabled in an organization for internal messages?

MCT USE ONLY. STUDENT USE PROHIBITED

Module Review and Takeaways

Best Practices

- Use transport rules to check messages for specific conditions and take defined actions on the messages before they reach the recipients.
- Use message moderation to avoid cases when messages with unimportant content are sent to large distribution lists.
- Train the users how to use the AD RMS functionality.
- Use transport protection rules to protect messages regardless of the client.

Review Question(s)

Question: What happens if a message to a moderated recipient is not answered within five days?

Tools

- Exchange Administration Center
- Exchange Management Shell

MCT USE ONLY. STUDENT USE PROHIBITED

MCT USE ONLY. STUDENT USE PROHIBITED

Module 6

Designing and Implementing Message Retention

Contents:

Module Overview **06-1**

Lesson 1: Overview of Messaging Records Management and Archiving **06-2**

Lesson 2: Designing In-Place Archiving **06-6**

Lesson 3: Designing and Implementing Message Retention **06-11**

Lab: Designing and Implementing Message Retention **06-18**

Module Review and Takeaways **06-24**

Module Overview

Microsoft® Exchange Server 2013 provides tools to address a growing number of legal, regulatory, and internal policy and compliance requirements that relate to email. Most organizations must be able to filter email delivery based on several criteria, and to manage email retention and deletion. This module shows you how to configure the Exchange Server 2013 messaging policy and compliance features.

Objectives

After completing this module, you will be able to:

- Describe message records management and archiving.
- Design in-place archiving.
- Design and implement message retention.

MCT USE ONLY. STUDENT USE PROHIBITED

Lesson 1

Overview of Messaging Records Management and Archiving

Email has become a reliable and ubiquitous communication medium for information workers in organizations of all sizes. Messaging stores and mailboxes have become repositories of valuable data. As a best practice, organizations should establish messaging policies that provide guidelines to users about how to use the messaging system responsibly. These messaging policies can also establish the kind of communication that may not be allowed.

Organizations must also create policies to manage the email lifecycle. These email lifecycle policies may require system administrators to retain messages for a length of time based on business, legal, and regulatory requirements, to preserve email records for litigation and investigation purposes, and to be prepared to search for and provide the required email records to fulfill eDiscovery requests.

This lesson provides an overview of the options available in Exchange Server 2013 that help you comply with your organization's business and legal requirements.

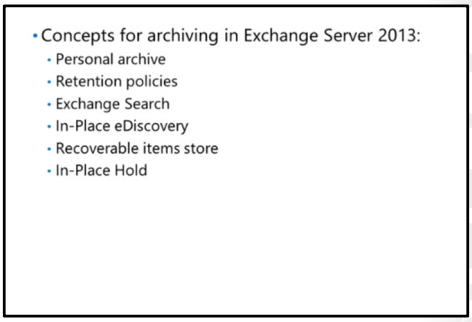
Lesson Objectives

After completing this lesson, you will be able to:

- Describe archiving.
- Describe message records management.
- Describe other options for implementing message records management and archiving.
- Describe when to use message records management and archiving.

What Is Archiving?

A compliance issue that many organizations must solve arises from the fact that much of the information users receive by email is not stored within the email system. In order to avoid mailbox size limits, many users move messages from their mailboxes to personal storage table (PST) files (also called *personal folder files*), which are typically stored on the local computer. These messages are not backed up regularly and are not available for discovery or indexing.



In the past, if a specific email message needed to be traced, it often took weeks to find it. With today's compliance legislation and legal discovery rules, many Information Technology (IT) departments now need to manage the entire organization's email archiving in bulk, so specific messages can be located in minutes. Archiving is the process of managing the size of an organization's data store by taking a backup copy of historical data, removing it from its native environment, and storing it elsewhere.

Exchange Server 2013 has archiving capabilities in addition to enhanced mailbox management features, including the capability to perform advanced multi-mailbox searches and to apply legal holds and granular retention policies for individual mailboxes. In Exchange Server 2013, archiving consists of the following concepts.

MCT USE ONLY. STUDENT USE PROHIBITED

Personal Archive

A personal archive is an additional mailbox that is associated with the user's primary mailbox. The personal archive can reside on any Mailbox server in the same Active Directory® site as the user's primary mailbox. Users can view, move, and copy messages between the primary and the archive mailbox. An archive mailbox presents a consistent view of messaging data to users, and it eliminates the need to manage PST files. Users can access the archive mailbox only in online mode.

Retention Policies

You can apply archive policies to a mailbox to automatically move messages from a user's primary mailbox to the archive mailbox after a defined retention period.

Exchange Search

With archive mailboxes, the ability to search messages quickly is more critical than ever. There are no differences for Exchange Search between the primary and archive mailbox, because the content of both mailboxes is indexed. The archive mailbox is not cached on a user's computer, so the search results for the archive are always provided by Exchange Search. If you search the entire mailbox in the Microsoft Office Outlook® 2010 messaging client or in a newer version, or in Microsoft Outlook Web App, the results always include the primary and the archive mailbox.

In-Place eDiscovery

A user's archive mailbox is also searched if a discovery manager performs an In-Place eDiscovery search. You cannot exclude an archive mailbox from a discovery search that is running from the Exchange Administration Center (EAC). You must use the Exchange Management Shell to exclude the archive from the search.

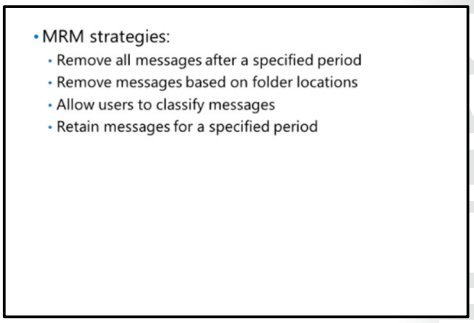
Recoverable Items Store

A Recoverable Items folder is also available in the archive mailbox, and the quotas are the same as for the folder in the primary mailbox.

In-Place Hold

If you enable a mailbox for an In-Place Hold, both the primary and the archive mailbox are placed on hold.

What Is Messaging Records Management?



Organizations handle an increasing volume of email every day. The email contains messages that are important from a business, legal, and regulatory perspective, and the organization may need to retain the messages for a certain period of time, depending on its messaging policies. However, many email messages do not have any retention value beyond a certain period. For example, a user's mailbox may contain critical messages that need to be retained, such as messages related to business strategy, transactions, product development, or customer interactions. However, messages such as newsletter subscriptions or personal messages may not have any retention value, so they do not need to be retained after a certain period.

Messaging records management (MRM) in Exchange Server 2013 helps you reduce the legal, regulatory, and business risks associated with email. MRM makes it easier to keep the messages that you need in

MCT USE ONLY. STUDENT USE PROHIBITED

order to comply with company policy, government regulations, and legal needs, and to remove content that has no legal or business value. In Exchange Server 2013, this retention is done through retention policies.

MRM Strategies

You can use retention policies to enforce basic message retention on default folders and on an entire mailbox. In combination with In-Place Hold, you can implement more effectively the MRM policies of your organization.

Remove All Messages After a Specified Period

You can implement a records management policy to remove messages after a certain period. These messages are not classified. You can implement this policy by creating a single default policy tag for the mailbox. This policy does not ensure that messages are retained for the specified period, because users can delete messages before this period ends.

Remove Messages Based on Folder Location

You can base records management policies on the location of email messages. For example, you can retain messages in the Inbox for a certain period, and you can retain messages in the Junk Mail folder for a shorter period. To implement this policy, use a combination of retention policy tags for each default folder that you want to have different settings and a default policy tag for the entire mailbox.

Allow Users to Classify Messages

You can implement MRM policies that set a baseline retention setting for messages in the mailbox and that also give users permission to set retention settings based on business or regulatory requirements.

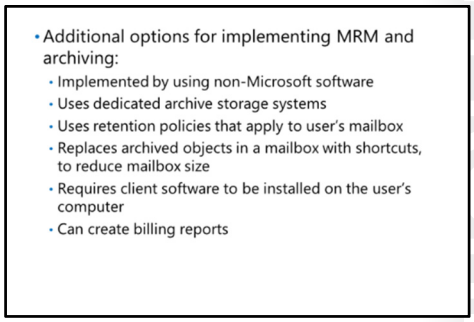
Retain Messages for a Specified Period

You can retain messages for a specified period. This means not only that messages are removed after the specified period, but also that they are retained for that period, even if the user or any process deletes them.

Note that for every mailbox you configure for MRM, you need an Exchange Server 2013 Enterprise client access license.

Additional Options for Implementing MRM and Archiving

Organizations may also use non-Microsoft products to achieve MRM and archiving according to their own rules. These products provide different features, and they are implemented in different ways. This topic describes how these products are implemented and how they work.



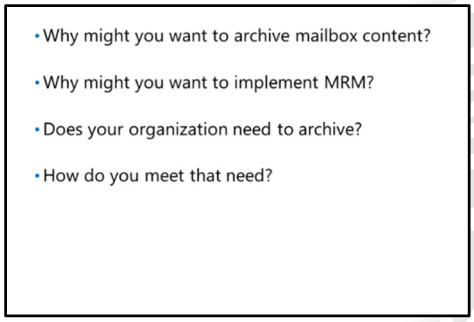
You can create retention policies for folders and configure the policies according to your organization's requirements. If a retention policy initiates the archiving process, objects that comply with the policy are removed from the user's mailbox and replaced with a smaller shortcut, which points to the object itself in the archive storage system. If the archived object is needed, the user can double-click it to retrieve the object from the archive and work on it. To set the retention settings on a user's objects in their mailbox, you can create a service account that has rights on all mailboxes that are selected for archiving. Often, after a defined period, the shortcuts are also removed from the mailbox to save

MCT USE ONLY. STUDENT USE PROHIBITED

space. After this, the user can run a program that searches the archive for the needed objects. If configured in the storage system, the user cannot delete the archived objects if these objects are moved from the user's mailbox to the archive. Also, you can get detailed reports about the space that is used by every user in the archive storage system to provide billing for the organization.

Discussion: Why Might You Want to Use Archiving and MRM?

Use the discussion questions to help examine why you would want to use archiving and MRM.



MCT USE ONLY STUDENT USE PROHIBITED

Lesson 2
Designing In-Place Archiving

With In-Place Archives, you can store all messages in the mailbox in one location, where they are accessible and manageable. In order to implement In-Place Archives successfully, you need to plan carefully.

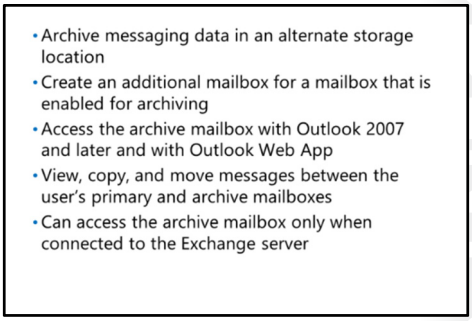
Lesson Objectives

After completing this lesson, you will be able to:

- Describe In-Place Archiving.
- Describe the benefits of In-Place Archiving.
- Describe In-Place Archiving for Microsoft Lync® Server 2013.
- Configure In-Place Archiving.
- Design storage for In-Place Archiving.

In-Place Archiving

Exchange Server 2010 introduced the ability to archive messaging data from Exchange Server into another mailbox database in the organization by using only Microsoft software. This was called a personal archive.



In Exchange Server 2013, this feature is now called In-Place Archiving. An In-Place Archive is an additional mailbox that you can enable for a user who has a mailbox. The In-Place Archive can be accessed with Outlook 2007 and newer versions, and with Microsoft Outlook Web App. With these applications, users may view, copy, and move messages between the user's primary mailbox and the In-Place Archive. It is not possible to access the archive mailbox with the Microsoft Exchange ActiveSync® technology.

How In-Place Archiving Works

The Exchange Server administrator enables the user's mailbox for In-Place Archiving. In the process, the administrator chooses where to store the archive mailbox for the user. The following locations are possible:

- The same mailbox server where the primary mailbox of the user resides.
- Another mailbox server in the same Active Directory site as the user.
- In the cloud, if the Exchange Server 2013 organization is running in hybrid mode.

The archive mailbox appears as a folder in user's primary mailbox when the user accesses their mailbox by using Outlook 2007 or newer versions, or by using Outlook Web App. Users can move their PST content, or any other messages, into the archive mailbox simply by dragging and dropping the email into an archive folder.

One of the differences between the primary and the archive mailbox is that, if Outlook is configured in cache mode, the archive mailbox is not cached on the client computer. This decreases the mailbox cache

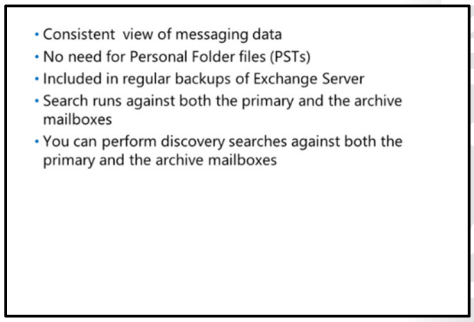
MCT USE ONLY. STUDENT USE PROHIBITED

size on the client, but it also means that the user can access the archive mailbox only when connected to the Exchange server.

You can manage the archive mailbox by using retention policies that move messages automatically from the primary to the archive mailbox. Messages are moved to the archive mailbox into a folder that has the same name as in the primary mailbox. If this folder does not exist in the archive mailbox, the Managed Folder Assistant creates it when a message is moved. This way, users can find messages easily.

Benefits of In-Place Archiving

With an In-Place Archive, the user gets a consistent view of their messaging data without needing to use PSTs. Because the content is stored in the user's In-Place Archive instead of in PSTs, the user does not need to think about saving the content. Instead, the content is saved during regular backups of Exchange Server, and it can be restored in case of a failure.



A user can search the In-Place Archive the same way they can search their mailbox, because content is indexed in both mailboxes. These search results for the In-Place Archive are provided by Exchange Search. If the user searches the entire mailbox, the search results include content from both the user's primary mailbox and the In-Place Archive.

By using In-Place eDiscovery, which is called Multi Mailbox Search in Exchange Server 2010, you can search for legally discoverable content within mailboxes and within archive mailboxes that are stored in Exchange Server 2013. In-Place eDiscovery enables you to do the following:

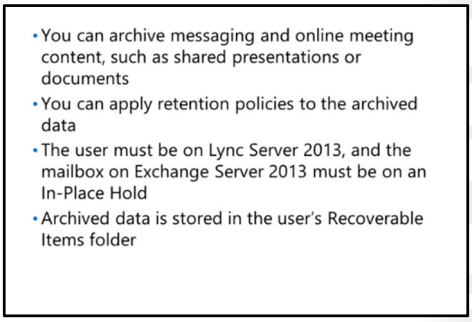
- Get the number of messages that Search will return.
- Get statistics to determine the effectiveness of keywords that are used.
- Preview search results.
- Copy messages from the search results to a discovery mailbox.

Only users who are members of the Discovery Management role group are authorized to perform In-Place eDiscovery searches.

MCT USE ONLY STUDENT USE PROHIBITED

In-Place Archiving for Lync Server 2013

In an organization using Exchange Server 2013, where the Lync Server 2013 communications software is also deployed, you can configure Lync Server to archive messaging and on-line meeting content, such as shared presentations or documents, in the mailbox of an Exchange Server 2013 user. If you do, you can apply retention policies to the data. You can also retrieve archived Lync Server content during eDiscovery searches.



Note that if you enable Exchange Server integration, Exchange Server controls purging for Exchange Server 2013 users and for their mailboxes that are on In-Place Hold. The only exception is for conferencing files, which are stored on the Lync Server file share. Conferencing files are purged from the file share only after the files are exported (that is, uploaded to Exchange Server), and only if you select the option to purge data either after the archiving data is exported or after a specified number of days.

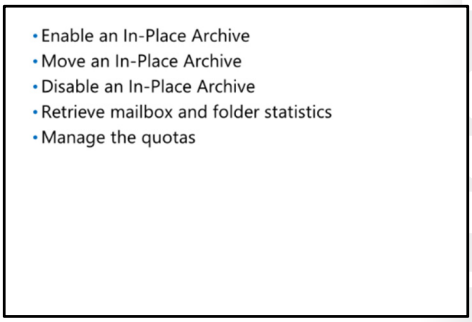
How Do I Access Archived Data?

If you choose the Exchange Server integration option, Lync Server deposits the archiving content in the Exchange Server 2013 store for all Exchange Server 2013 users. Archived data is stored in the Recoverable Items folder of the user's mailbox. Exchange Server enables federated search and discovery, as does Microsoft SharePoint® Server 2013, if it is deployed.

Configuration of In-Place Archiving

The following options are available for the In-Place Archive:

- Enable the In-Place Archive.
- Move the In-Place Archive.
- Disable the In-Place Archive.
- Retrieve mailbox and folder statistics.
- Manage the quotas.



Enable the In-Place Archive

You can create an archive mailbox when you create the primary mailbox, or you can enable an archive mailbox for an existing mailbox.

Move the In-Place Archive

You can move the archive mailbox to other mailbox databases, which might be on other servers. If the Exchange Server 2013 organization is running in hybrid mode, you can also move the archive mailbox to a mailbox database in the cloud. The location of the archive mailbox is independent of the primary mailbox.

Disable the In-Place Archive

You can disable the user's archive mailbox if it is on an Exchange server that does not support the use of archive mailboxes, meaning Exchange Server 2007 and earlier. A disabled mailbox is retained until the

MCT USE ONLY. STUDENT USE PROHIBITED

retention period for deleted mailboxes is reached. During this period, you can reconnect the archive mailbox to a mailbox user.

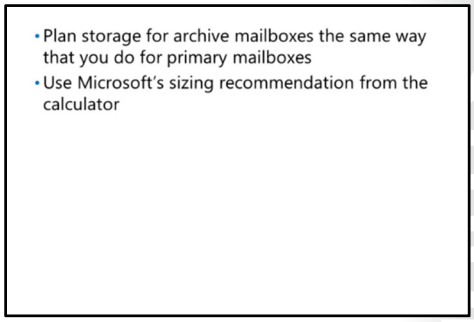
Retrieve Mailbox and Folder Statistics

You can get mailbox and folder statistics about the archive mailbox if you use the **–Archive** parameter.

Manage the Quotas

By default, an archive mailbox has an archive warning quota of 45 gigabytes (GB) and an archive quota of 50 GB. You can change this quota to meet your organization's requirements. Quota information is written to the application event log, and messages are also sent to inform the user that they are reaching their available archive mailbox space.

Designing Storage for In-Place Archiving



If you decide to use In-Place Archives in Exchange Server 2013, carefully plan your storage requirements up front. Many users store lots of messages in PST files. If you plan to move all of these messages into the In-Place Archive mailbox, be sure that you have enough storage, because the amount that you need will increase dramatically. Plan for In-Place Archives storage the same way you would when you plan for new mailboxes.

For example, if you have 1,000 users, they all get mailboxes with a maximum size of 5 GB, and they all get an archive mailbox with a maximum size of 50 GB. In this case, you have 5 terabytes (TB) of data for the primary mailboxes and 50 TB for the archive mailboxes. You can size this as if you had two classes of users, one with 5 GB quota mailboxes, and one with 50 GB quota mailboxes. Use the sizing recommendation from Microsoft in the calculator as though you were sizing other mailboxes. An option to save space is to decrease values of the default In-Place Archiving settings.

Starting with Exchange Server 2010 with Service Pack 1 (SP1), you can store the archive mailbox in a mailbox database that is separate from the user's primary mailbox. You can also store the archive mailbox in the cloud, if the organization is running in hybrid mode.

When you design storage for the In-Place Archive mailboxes, you can create the archive mailboxes on dedicated mailbox databases. This way, you can create fewer copies of the archive mailboxes than you do for the active user mailboxes.

Demonstration: Managing In-Place Archiving

Your organization is concerned about the amount of data that users store in PSTs. Some users from the IT department store several GBs of data in PSTs. Your organization has decided to provide all users in the IT department with an archive mailbox.

Demonstration Steps

1. Enable In-Place Archiving for all users who belong to the IT department.
2. Verify in the Exchange Administration Center (EAC) and in Outlook Web App that the In-Place Archiving mailbox is created for all users who belong to the IT department.

MCT USE ONLY. STUDENT USE PROHIBITED

MCT USE ONLY. STUDENT USE PROHIBITED

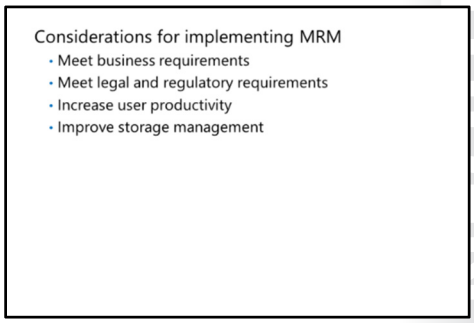
Lesson 3
Designing and Implementing Message Retention

Lesson Objectives

After completing this lesson, you will be able to:

- Describe MRM in Exchange Server 2013.
- Describe retention tags.
- Design and manage tags.
- Manage retention by using retention policies.
- Describe archive policies.
- Describe the MRM process in Exchange Server 2013.

MRM in Exchange Server 2013



MRM in Exchange Server 2013 helps you to reduce legal risks that are associated with email. With MRM in place, it is easier to keep the messages you need in order to comply with legal, business, and government requirements. As an example, say the mailbox of a user contains critical and non-critical messages. Critical messages need to be retained, which causes the user's mailbox to grow, which in turn requires additional space on the Mailbox servers. With retention policies, you can retain critical messages and remove content that has no business or legal value. When a message reaches its retention age, the retention action specified in the retention tag is taken. Messages can be moved to the Deleted Items folder, they can be deleted but remain recoverable from the Recoverable Items folder, or they can be permanently deleted. You can also use retention tags to move a message to the user's archive mailbox, if the user has one.

The strategy to make MRM and policy enforcement more reliable, effective, and easy to use is based on the following principles:

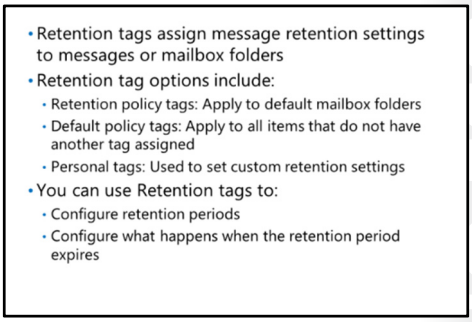
- Users can tag messages with personal tags.
- Messages that have no retention value are removed.
- Messages that have some retention value are retained.

If a user's mailbox is enabled for archiving, a default retention policy is assigned to the user's mailbox. The default retention policy contains one default policy tag and ten personal retention tags. With this policy in place, all items that reach the retention period are moved to the user's archive mailbox. This action occurs automatically every time the Managed Folder Assistant processes the mailbox. With the personal tags available, the user can select items in their mailbox and stamp them with different personal retention tags. If the user identifies items in this mailbox that are no longer needed, the user can stamp them with a personal tag that has a retention action of Delete and Allow Recovery and a retention period of one week. This stamp means that when the Managed Folder Assistant processes the mailbox, it deletes the item after the retention period is reached.

MCT USE ONLY. STUDENT USE PROHIBITED

Retention Tags in Exchange Server 2013

You can use retention tags to apply retention settings to items and folders in the user's mailbox. The applied settings specify how long a message stays in the user's mailbox and what happens when the message reaches its retention age. When a message reaches its retention age, it can be moved to the user's In-Place Archive or it can be deleted. This action depends on the retention tag settings that you choose when you create the retention tag. You can also allow users to tag items and folders in their own mailbox.



The following types of retention tags are available.

Retention tag	Description
Default policy tag	Applies to untagged mailbox items in the entire mailbox. Untagged items are mailbox items that do not have a retention tag applied.
Retention policy tag	Applies retention settings to default folders, such as Inbox, Deleted Items, and Sent Items. Items in a default folder that have an applied retention policy tag inherit the tag of the folder. Users cannot apply or change a retention policy tag that is applied to a default folder. They can apply a different tag to the items in it.
Personal tag	Is available to Outlook 2010 and newer versions, and to Outlook Web App. They are part of the user retention policy. Users can apply personal tags even if those items have a different tag applied.

When planning retention tags, consider the following:

- Messages with a personal tag applied are always processed based on the settings of the personal tag.
- You cannot include more than one retention tag for the same default folder in one retention policy.
- You cannot apply retention policy tags to the Contacts folder.

With retention policy tags, the following actions can occur when the retention age of an item is reached:

Action	Description
Move to Archive	Moves a message to the user's archive mailbox. If no archive mailbox is available, no action is taken.
Delete and Allow Recovery	Moves a message to the Recoverable Items folder. The user can recover deleted messages.
Permanently Delete	Purges a message from the mailbox. The user cannot recover deleted messages.
Mark as Past Retention Limit	Marks a message as expired. This action is available only in the Exchange Management Shell.

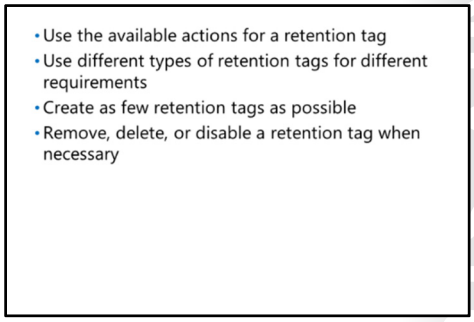
MCT USE ONLY. STUDENT USE PROHIBITED

If you combine retention tags with In-Place Hold or Single-Item Recovery, you get the following results:

- Permanently deleted items are retained in the Recoverable Items Store when an In-Place Hold is enabled for the user until the hold is disabled.
- Permanently deleted items are retained in the Recoverable Items Store when single-item recovery is enabled for the user until the deleted item retention period of the mailbox or the mailbox database is reached.

Designing and Managing Tags

Retention tags are used to apply retention settings to folders and items in the user's mailbox. The settings in the retention tags specify how long a message is retained in the mailbox and what happens when the retention period expires. The following actions are available for a retention tag:



- Move to the In-Place Archive.
- Delete the object.

Users can also use personal tags to tag their own folders and items in their mailbox for retention.

As an administrator, you can create the following types of tags:

- Default policy tags.
- Retention policy tags.
- Personal tags.

Each type of tag has its own retention settings that you can apply to a user's mailbox by using a retention policy.

As a best practice, before you define the tags, you should collect all of your organization's compliance requirements. This way, you can create only the retention tags that you really need, which reduces the work required to manage all of the available retention tags in your organization.

For example, assume your organization's compliance requirements state that all email messages older than 60 days must be moved to an archive mailbox. All objects in the Deleted Items folder must be deleted permanently after 30 days. Users cannot have the option to tag items themselves.

In this case, create one default policy tag that moves all items into the archive mailbox after 60 days. Additionally, create one retention policy tag that applies to the Deleted Items folder and that permanently deletes all objects in that folder after 30 days. Then, create one retention policy that links these two tags, and apply it to all of the users. You have now created the tags that enforce your organization's compliance requirements.

Remove and Disable Retention Tags

If you remove a retention tag from a retention policy, consider the following:

- If the tag is a personal tag, it is no longer available to the user and therefore cannot be applied to items in the mailbox.
- Items stamped with a removed tag continue to be processed by the Managed Folder Assistant with the settings and actions specified in the tag.

MCT USE ONLY. STUDENT USE PROHIBITED

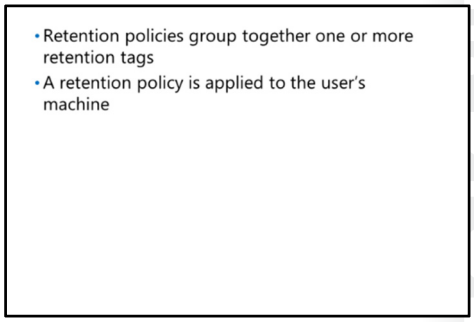
If you delete a retention tag, consider the following:

- When a tag is deleted, the definition is removed from Active Directory Domain Services (AD DS).
- The Managed Folder Assistant must now process all items in the mailbox and restamp the messages that have the deleted tag applied. This can consume significantly more resources on the Mailbox servers where the mailboxes are located.

You can also disable a retention tag as a first step before you remove it from a retention policy. If you disable a retention tag, an item that has this tag applied is ignored during the Managed Folder Assistant process. A retention period for a disabled retention tag is displayed as Never to the user.

Managing Retention by Using Policies

To apply retention tags to a user's mailbox, you need to create a retention policy, link the retention tags to that policy, and then apply the policy to the user's mailbox. You can link or unlink retention tags from a retention policy at any time. The changes take effect automatically for all mailboxes that the policy applies to.



The following retention tags can be included in a retention policy:

- One or more retention policy tags for supported default folders.
- One default policy tag with the Move to Archive action.
- One default policy tag with the Delete and Allow Recovery or Permanently Delete actions.
- One default policy tag for voicemail messages with the Delete and Allow Recovery or Permanently Delete action.
- Any number of personal tags.

Note that you cannot link more than one retention policy tag for a particular default folder, such as Deleted Items, to the same retention policy.

Although it is not necessary, you should make sure that that your retention policies have retention tags linked to them. Mailboxes that have retention policies with no retention tags linked to them may cause mailbox items to never expire.

Default MRM Policy

Exchange Server 2013 has a default retention policy, which is called the Default MRM Policy. When a mailbox that does not already have a retention policy is enabled for In-Place Archiving, this default policy is applied automatically. With this policy in place, all objects in the primary mailbox that are more than two years old are moved to the user's archive mailbox. If a retention policy is later applied to that mailbox, tags from the Default MRM Policy are no longer available to the mailbox. Items that have tags applied from the Default MRM Policy are still processed.

Managed Folder Assistant

Mailboxes that have retention policies applied are processed through the Managed Folder Assistant, which runs in every Mailbox server. The Management Folder Assistant is throttle-based, which means that the assistant is always running and does not need to be scheduled.

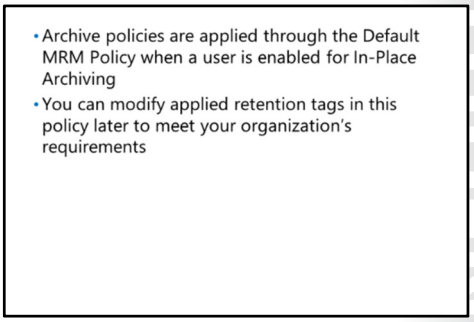
MCT USE ONLY. STUDENT USE PROHIBITED

Note that the Managed Folder Assistant takes no action on messages that are not subject to retention. You can make a message not subject to retention by using a personal tag and setting the retention period to Never.

When you move items and remove retention tags, keep the following in mind:

- When an item is moved from one folder to another, the item inherits any tags that already exist on the destination folder. If the destination folder does not have any tags applied to it, the default policy tag is applied. If the item has a tag explicitly applied, that tag always takes precedence.
- A tag is no longer available to the user if the retention tag is removed from the retention policy.
- Existing items that are stamped with the removed tag continue to be processed.
- Deletion of a tag removes the tag from AD DS. Items with this tag applied are restamped through the Managed Folder Assistant and no longer have this tag applied. The tag is also deleted from all retention policies.

Using Archive Policies



Starting with Exchange Server 2010 with SP1, you can use archive policies to automatically move mailbox items to either personal, also called *on-premises*, or cloud-hosted archives. The Move to Archive retention action is done through retention tags that are linked to the Archive policy.

When you set up Exchange Server, a retention policy called Default MRM Policy is created. This policy has a default policy tag linked to it that moves items to the archive mailbox after two years. This policy also includes several personal tags that users can apply to folders or mailbox items to automatically move or delete messages. If a mailbox does not have a retention policy assigned when it is archive-enabled, Exchange Server automatically applies the Default MRM Policy to it through the Managed Folder Assistant. You can also create your own archive policies and apply them to the mailbox users.

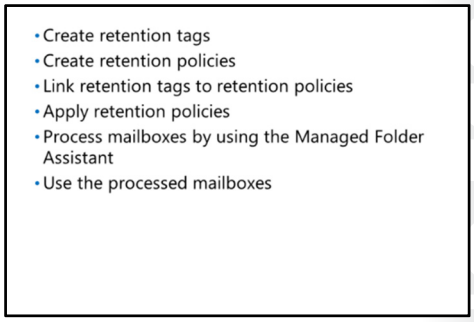
You can modify the retention tags that are linked to the default policy to meet your organization's requirements. For example, you can modify the archive default policy tag to move items to the archive after five years instead of two. You can also create personal tags and add them to the retention policy, or you can allow users to add personal tags to their mailboxes from Exchange Control Panel.

MCT USE ONLY. STUDENT USE PROHIBITED

Messaging Records Management Process in Exchange Server 2013

Setting up an MRM process consists of the following steps:

1. Create retention tags.
2. Create retention policies.
3. Link retention tags to retention policies.
4. Apply retention policies.
5. Process mailboxes by using the Managed Folder Assistant.
6. Use the processed mailboxes.



Create Retention Tags

Create retention tags to apply different retention settings to the items in the user's mailbox.

Create Retention Policies

Create retention policies to group retention tags together. A retention policy is applied to the user's mailbox.

Link Retention Tags to Retention Policies

You can link existing retention tags to retention policies. This linking makes it easy to apply the retention tags to the users' mailboxes in your organization. Note that a retention policy can have one default policy tag, one retention policy tag for each supported default folder, and any number of personal tags.

Apply Retention Policies

Apply retention polices to the users. You can apply different policies to different users.

Process Mailboxes by Using the Managed Folder Assistant

The Managed Folder Assistant on the Mailbox server processes all mailboxes, applies the retention settings to the mailbox items, and takes specific retention actions.

Use the Processed Mailboxes

The default policy tag and retention policy tags are applied to the users' mailboxes, and the personal retention tags are now available for the users in Outlook and Outlook Web App.

Demonstration: Using MRM in Outlook Web App

This demonstration shows you how to configure a retention tag and how to configure the content settings for this tag. It then shows you how to merge the retention tag with a retention policy. Finally, it shows you how to assign the retention policy to a user's mailbox and how the tag appears in both Outlook and Outlook Web App.

Demonstration Steps

1. Log on to LON-CAS1 as **adatum\administrator** with the password **Pa$$w0rd**.
2. Open the Exchange Management Shell, and then create a new default retention policy tag, named **DefaultTag**, that applies to all folders.
3. The retention policy content applies to all messages that do not have another retention tag assigned, and it permanently deletes all messages after 365 days.

MCT USE ONLY. STUDENT USE PROHIBITED

4. Create a new retention policy tag for the folder Inbox, and configure a content setting to move all messages to the Deleted Items folder after 30 days.
5. Create a personal tag named **BusinessCritical** that sets a retention period of three years and that moves the messages to the user's archive mailbox after the retention period expires.
6. Create a retention policy named **AllTagsPolicy**, and then add all of the newly created retention tags to it.
7. Assign the retention policy to the user Kim.
8. Open Outlook Web App to check that the policy is applied to the mailbox.

MCT USE ONLY. STUDENT USE PROHIBITED

Lab: Designing and Implementing Message Retention

Scenario

Email is a critical business tool at A. Datum. Employees constantly use email to communicate with other employees and with external contacts. Business groups and the security team at A. Datum are concerned that some of this email is stored on client computer hard drives in PST files, where the data is not backed up and where it cannot be accessed by using server-side email search tools. Many users started using PST files with the previous deployment of Exchange Server, when the mailbox size limits were much smaller.

To ensure that this data can be moved back in to Exchange Server 2013 mailbox databases, A. Datum has decided to implement archive mailboxes for users who have large PST files. The storage team at A. Datum is concerned about the impact of adding all of this data to the Exchange mailbox databases, so you need to implement policies that enable automatic purging of messages that are not important to the business.

Objectives

After completing this lab, you will be able to:

- Design message retention and archiving.
- Implement message retention and archiving.

Lab Setup

Estimated Time: 60 minutes

Virtual machines	20342B-LON-DC1 20342B-LON-CAS1 20342B-LON-MBX1 20342B-LON-MBX2
User name	Adatum\Administrator
Password	Pa$$w0rd

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

1. On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In Microsoft Hyper-V® Manager, click **20342B-LON-DC1**, and then in the **Actions** pane click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Log on by using the following credentials:
   - User name: **Adatum\Administrator**
   - Password: **Pa$$w0rd**
5. Repeat steps 2 through 4 for 20342B-LON-CAS1, 20342B-LON-MBX1, and 20342B-LON-MBX2.

Exercise 1: Designing Message Retention and Archiving

Scenario

As part of planning Mailbox servers, A. Datum has decided to implement In-Place Archiving to replace PST files. However, the storage team at A. Datum is concerned about the amount of storage that this will

MCT USE ONLY. STUDENT USE PROHIBITED

require. To reduce the amount of additional storage needed to enable archiving, A. Datum has identified the following requirements related to message retention:

- All user mailboxes that currently have a PST file need to be configured with an In-Place Archive mailbox.
- You have identified that these users' files are all located in the Managers and Sales groups.
- A. Datum is deploying three copies of every mailbox database that is used for regular mailboxes. Mailbox databases where the archive mailboxes are stored need less redundancy.
- The users of the Sales and Manager groups have PSTs with an average size of 10 GB, so every user gets an archive mailbox with 20 GB as the archive quota for future growth. In order to meet backup and reseed needs, the maximum size of a database for the archive mailboxes should not exceed 200 GB.
- For Sales users who have archive mailboxes, all messages older than one year old should be archived.
- Messages in Managers mailboxes should be archived only after three years.
- Any Deleted Item messages should be purged after 30 days.
- Any messages in the default folders in Outlook should be moved to the Deleted Items folder after two years. Managers should be exempt from this requirement.
- All users must be able to control whether messages are deleted. They should be given the option of deleting messages in six months, one year, two years, or never.
- All users must be able to control whether messages are archived. They should be given the option of archiving messages in six months or never.

The main tasks for this exercise are as follows:

1. Design the Mailbox database configuration required for this deployment
2. Design the retention tags required for this deployment
3. Design the retention policies required for this deployment

▶ Task 1: Design the Mailbox database configuration required for this deployment

1. A. Datum deploys three copies of every mailbox database that contains regular mailboxes. It has decided that the databases for archive mailboxes can be less redundant than the databases for the regular mailboxes.
2. To make the databases for the archive mailboxes redundant, they need a minimum of two copies.
3. The Managers group has 42 users, and the Sales group has 77 users.
4. To get a rough estimate of the space you need, multiply the number of users by the size of the planned archive mailboxes.
5. Decide how many databases to create for each group.
6. Create an additional copy of the database on the second mailbox server.

   With this rough calculation, you get a fast estimate of how much additional space you need to support the additional archive mailboxes.

▶ Task 2: Design the retention tags required for this deployment

1. A. Datum has identified that it needs to implement different MRM settings for two different groups of users. These groups are:
   - Sales

MCT USE ONLY. STUDENT USE PROHIBITED

- Managers

2. To do this, A. Datum decides to create new retention tags and to use already available tags.
3. A. Datum needs to pay attention to which type of retention tag fulfills its requirements. For moving items to the archive, the company needs to create default policy tags, because only this type of tag allows it to move items to the archive.
4. The company needs retention policy tags to set the needed retention settings on the default folders in the user's mailbox (including, for example, Inbox and Sent Items).
5. The company needs personal tags so it can give users the option to tag messages themselves.
6. After these retention tags are created, they can be linked together with the already available retention tags in the appropriate retention policy.

▶ Task 3: Design the retention policies required for this deployment

1. Datum has identified that it needs to implement different MRM settings for different group of users. These two groups are:
   - Sales
   - Managers
2. To do this, it decides to create two retention polices, one for each group of users.
3. After these retention policies are created, they can be linked with the appropriate retention tags, and then the retention policy can be set on the user's mailbox.

Results: After completing this exercise, you will have successfully:

- Designed the required mailbox database configuration.
- Designed the required retention tags.
- Designed the required retention policies.

Exercise 2: Implementing Message Retention and Archiving

Scenario

In this exercise, use the design that you created in Exercise 1 to implement retention tags and policies, as well as In-Place Archive functionality.

The main tasks for this exercise are as follows:

1. Enable in-place archiving on mailboxes
2. Create the required retention tags
3. Create and apply the retention policies
4. Verify the configuration
5. To prepare for the next module

▶ Task 1: Enable in-place archiving on mailboxes

1. On LON-CAS1, open the Windows® Internet Explorer® browser.
2. Type **https://lon-cas1.adatum.com/ecp**, and then press Enter.

MCT USE ONLY. STUDENT USE PROHIBITED

3. Sign in as **adatum\administrator** with the password **Pa$$w0rd**.
4. Enable In-Place Archiving for all users who are members of the **Sales** and **Managers** department. Use the **Research** database for the archives.
5. Sign out of Exchange admin center.
6. In Internet Explorer, type **https://LON-CAS1.adatum.com/owa**, and then press Enter.
7. Type **adatum\Dan** and the password **Pa$$w0rd** to log on to Dan's mailbox with Outlook Web App.
8. Verify in the EAC and in Outlook Web App that an In-Place Archiving mailbox is created for all users who have large PSTs files.
9. Send a test email to Dan and Bill that will test the retention policies.
10. Close Internet Explorer.

▶ Task 2: Create the required retention tags

Based on your design result from Exercise 1, on LON-CAS1, sign in to the Exchange admin center as adatum\administrator and create the following retention tags for your organization:

- Default policy tag:
  - Name: Sales User 1 year move to archive
  - Retention Action: Move to Archive
  - Retention Period: 365 days
- Default policy tag:
  - Name: Default 2 year move to Deleted Items
  - Retention Action: Delete and Allow Recovery
  - Retention Period: 730 days
- Default policy tag:
  - Name: Manager 3 year move to archive
  - Retention Action: Move to archive
  - Retention Period: 1095 days
- Retention policy tag on deleted items folder:
  - Name: Purge Deleted Items 30 days
  - Retention Action: Permanently Delete
  - Retention Period: 30 days
- Personal tag:
  - Name: 2 Year Delete
  - Retention Action: Delete and Allow Recovery
  - Retention Period: 730 days
- Personal tag:
  - Name: Never archive
  - Retention Action: Move to Archive
  - Retention Period: Never

MCT USE ONLY. STUDENT USE PROHIBITED

▶ Task 3: Create and apply the retention policies

1. Based on your design result from Exercise 1, create the following retention policies for your organization:
   a. Retention Policy for Sales users
   b. Name: Sales MRM Policy
   c. Retention Tags included:
   d. 6 Month Delete
   e. 1 Year Delete
   f. 2 Year Delete
   g. Never Delete
   h. Sales user 1 year move to archive
   i. Default 2 year move to Deleted Items
   j. Purge Deleted Items 30 days
   k. Personal 1 year move to archive
   l. Never archive
   m. Retention Policy for Managers
   n. Name: Manager MRM Policy
   o. Retention Tags included:
   p. 6 Month Delete
   q. 1 Year Delete
   r. 2 Year Delete
   s. Never Delete
   t. Manager 3 year move to archive
   u. Purge Deleted Items 30 days
   v. Personal 1 year move to archive
   w. Never archive
2. Apply the retention policies based on their names to all Sales and all Manager mailboxes. Close Internet Explorer when complete.
3. Open the Exchange Management Shell as an administrator.
4. Type the following command to apply the retention polices to the mailboxes immediately:

```
Get-Mailbox | Where {$_.RetentionPolicy –ne $NULL} | Start-ManagedFolderAssistant
```

▶ Task 4: Verify the configuration

1. On LON-CAS1, log on to Outlook Web App as user Dan, a member of the Sales department, and check that the correct retention policy with the linked retention tags is applied to the mailbox.
2. Sign out from Outlook Web App.
3. Log on to Outlook Web App as user Bill, a member of the Managers department, and check that the correct retention policy with the linked retention tags is applied to the mailbox.

MCT USE ONLY. STUDENT USE PROHIBITED

4. Sign out from Outlook Web App.

▶ Task 5: To prepare for the next module

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

1. On the host computer, start Hyper-V® Manager.
2. In the **Virtual Machines** list, right-click **20342B-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
   a. Repeat steps 2 to 3 for 20342B-LON-CAS1, 20342B-LON-MBX1, and 20342B-LON-MBX2.
4. In Hyper-V Manager, click **20342B-LON-DC1**, and in the Actions pane, click **Start**.
5. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
6. Sign in using the following credentials:
   - User name: **Adatum\Administrator**
   - Password: **Pa$$w0rd**
7. Repeat steps 4 to 6 for **20342B-LON-MBX1, 20342B-LON-MBX2, 20342B-LON-CAS1,** and **20342B-LON-CAS2**.
8. Repeat steps 4 and 5 for **20342B-LON-CL1**. Do not log on until directed to do so.

Results: After completing this exercise, you will have successfully:

- Enabled In-Place Archiving on mailboxes.
- Created the required retention tags.
- Created and applied the retention policies.
- Verified the configuration.

Question: Which retention tags can users use to stamp items in their mailboxes themselves?

Question: Where can you store the In-Place Archive mailbox, if you enable this feature on a user's mailbox?

MCT USE ONLY. STUDENT USE PROHIBITED

Module Review and Takeaways

Review Question(s)

Question: What happens when a user's mailbox is enabled for In-Place Archiving?

Question: What happens if the quota for a user's In-Place Archive mailbox is reached?

MCT USE ONLY. STUDENT USE PROHIBITED

Module 7

Designing and Implementing Messaging Compliance

Contents:

Module Overview **07-1**

Lesson 1: Designing and Implementing Data Loss Prevention **07-2**

Lesson 2: Designing and Implementing In-Place Hold **07-10**

Lesson 3: Designing and Implementing In-Place eDiscovery **07-14**

Lab: Designing and Implementing Messaging Compliance **07-19**

Module Review and Takeaways **07-25**

Module Overview

Microsoft® Exchange Server 2013 provides many features that can prevent data loss from email. Today, email is used intensively as a way to communicate both business and personal matters. The intensity of this usage presents a significant security risk that business critical data might leave the organization. Also, some organizations need to monitor email traffic and content in their Exchange organization. Exchange Server provides several features to help you minimize data loss and monitor email traffic and content, including data loss prevention policies, In-Place Hold, and eDiscovery.

Objectives

After completing this module, you will be able to:

- Design and implement data loss prevention.
- Design and implement In-Place Hold.
- Design and implement In-Place eDiscovery.

MCT USE ONLY. STUDENT USE PROHIBITED

Lesson 1
Designing and Implementing Data Loss Prevention

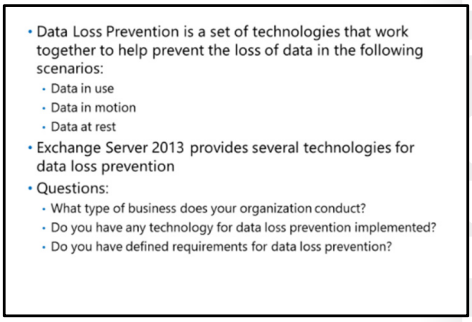
Preventing data loss is a key task for every administrator. Exchange Server 2013 provides a specialized feature, called data loss prevention (DLP) policies, that helps you to prevent the loss of data in email. This lesson explains how to design and implement DLP in Exchange Server 2013.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe DLP.
- Describe how DLP works.
- Describe DLP policies.
- Apply DLP policies.
- Create a DLP policy from a template.
- Configure policy tips.
- Configure and test custom DLP policies and policy tips.
- Plan a DLP implementation.
- Implement Document Fingerprinting.

Overview of Data Loss Prevention



In today's business, email is a critical communication resource. Email is a primary means of communication in many organizations, and users typically send a great deal of business information by email. This information may include confidential information, such as customer data or business intelligence. In some cases, business critical information can leak out of a company in unprotected email. Sometimes, a malicious user can leak business critical information intentionally, but users can also inadvertently send this information in an email. In that case, sensitive data, such as personal or company information, credit card details, or social security numbers is disclosed to unauthorized users.

In many countries, governments have implemented legislation that restricts the storage and movement of certain information. Additionally, many organizations have corporate security policies that limit how to share information within the organization. Because email is such a critical business tool, it is important that you configure your organization's messaging system so that it complies with government legislation and with corporate policies. This configuration is usually achieved by designing and implementing a DLP strategy that aims to minimize data loss in email traffic and in other areas.

A DLP strategy defines software and hardware solutions that monitor data in the following scenarios:

- Data in use. The strategy monitors data that is in use, such as data that is being copied to a USB drive or that is being printed.

MCT USE ONLY. STUDENT USE PROHIBITED

- Data in motion. The strategy monitors email traffic, web traffic, instant messages, and other communications transmitted over corporate networks.
- Data at rest. The strategy monitors data stored in file shares or on users' hard drives.

DLP technologies work together to minimize the possibility of users intentionally or inadvertently transmitting business critical data beyond the domain of the organization. Various types of policies and rules in Exchange Server 2013 enable you to manage email messages that are in transit or at rest, and to help ensure that your organization meets compliance requirements.

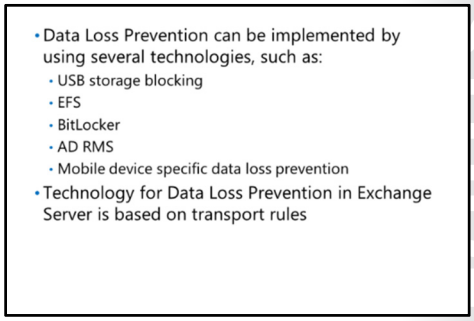
Question: What type of business does your organization conduct?

Question: Does your organization have any technology for DLP in place?

Question: Does your organization have defined requirements for DLP?

How Data Loss Prevention Works

To help prevent data loss, you should implement various technologies on several layers of your network and system infrastructure. For example, you can disable the use of USB removable storage devices to help prevent users from saving company data on USB storage. Also, you can encrypt critical data by using Encrypting File System (EFS), so that only authorized users can access the data. To help prevent data loss if a computer or hard drive is lost, you can also use the Windows® BitLocker® Drive Encryption.



Active Directory® Rights Management Services (AD RMS) is a technology that can also help prevent data from leaking. You can apply AD RMS manually by using Office 2013 applications or automatically by using Exchange Server transport rules.

Increased use of smartphones that synchronize data with an organization's internal infrastructure can also pose a security risk. You can address this risk with a mobile-specific DLP strategy.

However, when it comes to securing email traffic, it is difficult to prevent users from sending email messages outside of the organization. Email messages, in general, can be directed to any email address on the Internet, and that poses a data leakage risk. Because of that risk, you must identify potentially compromising data in email messages before it leaves the organization and prevent this data from being sent to email addresses outside of the organization or to unauthorized users inside the organization.

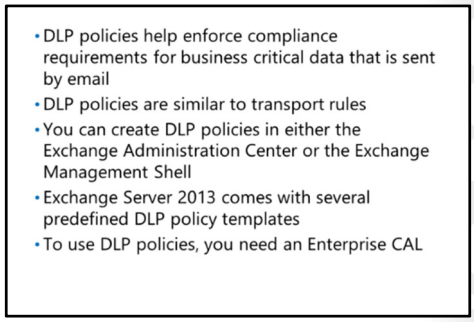
Exchange Server 2013 includes several methods to identify and control email messages that might compromise your organization's data, and all of these methods are based on transport rules. Transport rules can inspect a message while it is between the sender and the recipient, and, based on the message content and previously created rules, these rules can then reject the message, discard it, or forward it for moderation or approval before it leaves the organization.

The most important part of this approach is the identification of critical data in email messages. It is generally inefficient to look for particular words in messages. Message inspection must be based on patterns rather than a fixed set of words. Older versions of Exchange Server provide a limited ability to define these patterns. Exchange Server 2013 includes new technologies and features to define these patterns, so that Exchange Server 2013 can more easily detect security or business critical information in email messages.

MCT USE ONLY. STUDENT USE PROHIBITED

What Are DLP Policies?

To prevent data leaking through email, Microsoft has implemented DLP policies in Exchange Server 2013. The primary purpose of these policies is to enforce compliance requirements for business critical data and to manage how that data is used in email, without hindering the productivity of workers. For example, you can configure a policy to prevent sending credit card numbers, Social Security numbers, and IP addresses in email messages.



Note: DLP is a premium feature that requires an Enterprise client access license (CAL).

DLP policies are a set of conditions that contain transport rules, actions, and exceptions. When you apply DLP policies, the policies filter email traffic to help prevent business critical information that is in email from leaving the company. DLP policies are very similar to transport rules—in fact, they are transport rules with an extended set of options. The difference between transport rules and DLP policies is an approach to classifying sensitive information that can be incorporated into mail flow processing. This approach includes performing deep content analysis through keyword matches, dictionary matches, and regular expression evaluation to detect content that violates the organization's DLP policies.

You can create DLP policies in the Exchange admin center and in the Exchange Management Shell. You can create DLP policies for testing, in order to observe the effects of DLP policies, or you can enforce these polices on all email traffic in the organization.

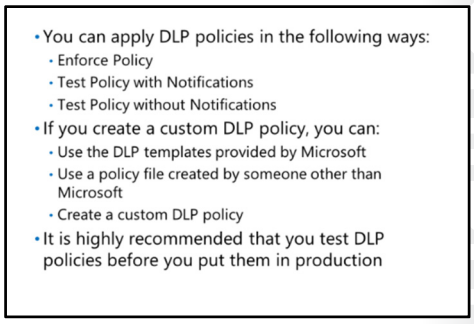
One of the benefits of DLP policies is that you can inform email senders that they may be violating one of the policies even before they send a message. You can inform users by using DLP Policy Tips, which are similar to MailTips, but they are preconfigured for use with DLP policies.

Exchange Server 2013 includes numerous DLP policy templates. You can also define custom policies and transport rules as an alternative to the DLP policy templates.

In order to implement DLP policy features, you must have at least one mailbox active in Exchange Server 2013.

Applying DLP Policies

When you implement DLP policies, you can use any of the following methods:



- Enforce Policy. The DLP policy is enabled, and all rules and actions specified in it are applied to messages in transport.
- Test Policy with Notifications. The DLP policy is enabled, but the actions defined in it are not executed. They are logged in Message Tracking Logs. Users are notified by Policy Tips.
- Test Policy without Notifications. Similar to Test Policy with Notifications, but no Policy Tips are

MCT USE ONLY. STUDENT USE PROHIBITED

displayed to users.

In addition to deciding how to apply policies, you should also decide how you want to create them, based on business requirements. You have the following choices:

- Use the DLP policy templates provided by Microsoft. This option provides the fastest way to start using DLP policies, and you do not need to build a complete set of rules yourself. However, if you use this method, make sure that that the template meets your compliance requirements. Some of the available DLP policy templates are the following:
- U.S. financial data. Helps to detect the presence of data that is commonly associated with financial information in the United States. This includes credit card numbers, account numbers, and debit card data.
- Germany financial data. Helps to detect the presence of data that is commonly associated with financial information in Germany. This includes credit card numbers, account numbers, and debit card data.
- U.S. Health Insurance Portability and Accountability Act (HIPAA). Helps to detect the presence of data that is commonly associated with health information that is subject to HIPAA.
- U.S.A. PATRIOT Act. Helps to detect the presence of data that is commonly subject to the U.S.A. PATRIOT Act.
- U.K. Access to Medical Reports Act. Helps to detect the presence of data that is commonly associated with health information in the United Kingdom.
- Israel Protection of Privacy. Helps to detect the presence of data that is commonly associated with private information in Israel.
- Saudi Arabia Anti-Cyber Crime Law. Helps to detect the presence of data that is commonly associated with the cyber-crime law in Saudi Arabia.
- Use a policy file provided by a company other than Microsoft. You can import policies that are created by independent software vendors, so you can extend the functionality of DLP policies to better meet your compliance requirements. You can import these policies from the policy file.
- Create a custom policy. If any of the predefined policies do not meet your requirements, you have the option of creating a custom policy in order to start checking and acting on your organization's unique message data. To implement a custom DLP policy, you need to know the requirements and constraints of the environment in which the DLP policy will be enforced.

If you create DLP policies, you can include rules that check for sensitive information. These information types should be used in your policies. You can customize the conditions within a policy in order to meet your specific policy requirements, such as how many times something is found before an action is taken.

 **Note:** As a best practice, you should test the DLP policies before you run them in the production environment. During the tests, you should configure sample user mailboxes and send test messages that invoke your test policies in order to confirm the results.

Demonstration: Creating a Data Loss Prevention Policy from a Template

Demonstration Steps

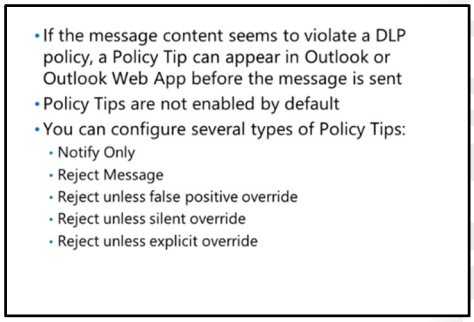
1. On LON-CAS1, sign in to the Exchange admin center as **Adatum\Administrator**.
2. Go to **compliance management,** and then go to **data loss prevention**.
3. Start a wizard to create a new DLP policy from a template.

MCT USE ONLY. STUDENT USE PROHIBITED

4. Choose the **U.S. Financial Data** template.
5. Choose **Test DLP Policy with Policy Tips**.
6. Configure Policy Tips with the text **This message contains information that you are not allowed to send**.
7. Click **mail flow** in the feature pane, and then in the **rules** tab notice that several transport rules are created as a result of creating the DLP policy.

Configuring Policy Tips

Sometimes, users inadvertently send inappropriate information in email. You can help prevent your users from inappropriately sending sensitive information, by implementing Policy Tip notification messages.



Policy Tips are similar to Mail Tips. In the Microsoft Office Outlook® 2010 messaging client or later, and in Microsoft Outlook Web App, if the content of a message seems to violate a DLP policy, a Policy Tip appears. By using Policy Tips you can help prevent users from trying to send a message that might contain data that should not be sent by email. If a user who is composing a message learns of the organization's expectations and standards through the Policy Tip, they are less likely to violate the standards. For example, if an organization has a DLP policy that forbids sending financial data outside the company and a user includes a credit card number in the message body and addresses the message to someone outside the organization, a Policy Tip can be displayed to warn the user that the message violates company rules.

Policy Tips are not enabled on each DLP policy by default. You need to manually configure a DLP policy to enable Policy Tips. To display Policy Tips to the users, the rules within the DLP policy must include the action **Notify the sender with a Policy Tip**. You can add this action in the rules editor from the Exchange admin center.

Policy Tips are implemented by using a transport rule agent that enforces DLP policies. This agent does not differentiate between email message attachments, body text, and subject lines while evaluating messages and the conditions within the DLP policies. When you configure Policy Tips, you can choose from the following options:

- Notify Only. This is just an informative message that does not prevent the user from sending a message. The user sees the following text: **This message may contain sensitive content. All recipients must be authorized to receive this content.**
- Reject Message. The message is not delivered if the message body contains data that violates the DLP policy. The user can report the message as a false positive, so that the administrator can examine it, but the user cannot send the message. The user sees the following text: **This message may contain sensitive content. Your organization won't allow this message to be sent until that content is removed.**
- Reject unless false positive override: This is similar to the Reject Message option, except the user can override the limitation and send the message if they think that the message contains no sensitive content. Before the user overrides the limitation, the following text appears: **This message may contain sensitive content. Your organization won't allow this message to be sent until that**

MCT USE ONLY. STUDENT USE PROHIBITED

content is removed. If the user overrides the limitation, the following text appears: **Your feedback will be submitted to your administrator when the message is sent.**

- Reject unless silent override. The message is not sent if sensitive content is detected, but the user can override this rule. Before the override, the user sees the same message as in the previous case. But if the user chooses to override, the following text appears: **You have overridden your organization's policy for sensitive content in this message. Your action will be audited by your organization.**
- Reject unless explicit override. This option is similar to the previous one, except the user must provide a justification for overriding the policy if they choose to override the limitation.

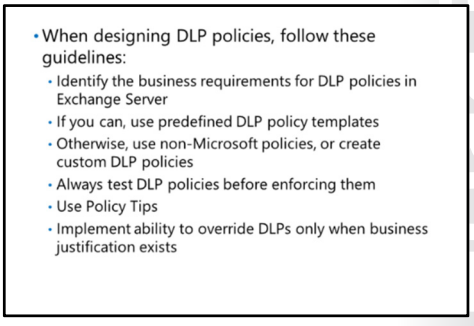
Demonstration: Configuring and Testing Custom DLP Policies and Policy Tips

Demonstration Steps

1. In the Exchange admin center, go to **compliance management**, and then go to **data loss prevention**.
2. Click **new custom DLP Policy**.
3. Configure the policy as follows:
   - Policy is **Enforced**
   - Name of policy: **IP address block**
   - Include rule: **Block messages with sensitive information**
   - Sensitive information type: **IP address**
   - Action: **Generate incident report and send it to Administrator**
   - Action: **notify the sender with a Policy Tip with text "your message is blocked."**
4. Activate and save the policy.

Planning to Implement Data Loss Prevention Policies

DLP policies can be very useful in minimizing the possibility of data leaking through email. But if DLP policies are implemented improperly, they can block regular email traffic by generating false positive results. Therefore, you need to carefully plan the design and implementation of DLP policies. Detailed plans, evaluations, and documentation of DLP policies help you to precisely define conditions for DLP policies, and therefore minimize data loss and false positives.



When you plan and design DLP policies, follow these guidelines:

- Identify a business scenario for DLP policies in Exchange Server. DLP policies in Exchange Server 2013 cannot prevent data loss in general. They focus only on email traffic. Therefore, you also need to plan DLP on other layers of the network. You also need to identify your organization's compliance and business needs that require DLP policies on Exchange Server.

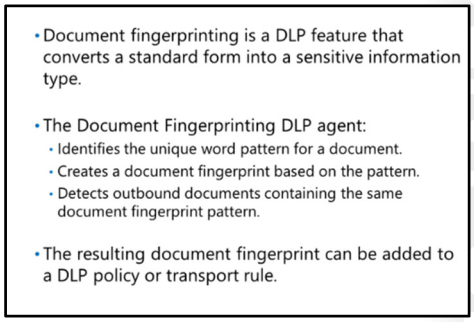
MCT USE ONLY. STUDENT USE PROHIBITED

- If you can, use the predefined DLP policy templates. Exchange Server 2013 comes with many predefined DLP policy templates. After you identify your DLP policy requirements, review the templates in Exchange Server 2013 and see if they meet your organization's needs.
- If you need to, use other DLP policies. If the predefined DLP policy templates in Exchange Server 2013 do not meet your requirements, either use non-Microsoft policies or create custom policies.
- Always test DLP policies before enforcing them. You can test and monitor the functionality of each DLP policy. As a best practice, test DLP policies before you put them into production in order to prevent false positives and to minimize data leaks.
- Use Policy Tips. We recommend that you use Policy Tips with DLP policies. With Policy Tips, you can warn users that content in a message might violate the organization's rules. By increasing user awareness about data leaks and data loss prevention, you can complement the functionality of DLP policies in general.
- Implement an override for DLP policies only if you have a business justification to do so. If you allow users to override DLP policies, you are actually allowing them to send potentially sensitive data out of the organization. Be sure to precisely define any scenarios in which users can override DLP policies.

In addition to using DLP policies, you should also prevent data leaks that can occur in other ways, such as by copying data to USB drives or collecting data on mobile phones.

Implementing Document Fingerprinting

Information workers handle many kinds of sensitive information during a typical day. To protect this information, organizations can convert standard forms that they use into a sensitive information type. This is known as document fingerprinting, which is a DLP feature that you can use to define transport rules and DLP policies. For example, you can create a document fingerprint based on a blank patent template and then create a DLP policy that detects and blocks all outgoing patent templates with sensitive content filled in.



How Document Fingerprint works

In the same way that a person's fingerprints have unique patterns, documents have unique word patterns. When you upload a file, the Document Fingerprinting DLP agent identifies the unique word pattern in the document, creates a document fingerprint based on that pattern, and uses that document fingerprint to detect outbound documents containing the same pattern. This is why uploading a form or template creates the most effective type of document fingerprint. Everyone who fills out a form uses the same original set of words and then adds his or her own words to the document. As long as the outbound document is not password protected and contains all the text from the original form, the DLP agent can determine if the document matches the document fingerprint.

For example, consider a patent template that contains the blank fields "Patent title," "Inventors," and "Description" and descriptions for each of those fields. These fields make up the word pattern for that document. When you upload the original patent template, it is in one of the supported file types and in plain text. The Document Fingerprinting DLP agent uses an algorithm to convert this word pattern into a document fingerprint, which is a small Unicode XML file containing a unique hash value representing the original text, and the fingerprint is saved as a data classification in Active Directory.

MCT USE ONLY. STUDENT USE PROHIBITED

 **Note:** As a security measure, the original document itself is not stored on the service; only the hash value is stored, and the original document cannot be reconstructed from the hash value.

The patent fingerprint then becomes a sensitive information type that can be associated with a DLP policy. After you associate the fingerprint with a DLP policy, the DLP agent detects any outbound emails containing documents that match the patent fingerprint and deals with them according to your organization's policy.

For example, you might want to set up a DLP policy that prevents regular employees from sending outgoing messages containing patents. The DLP agent will use the patent fingerprint to detect patents and block those emails. Alternatively, you might want to let your legal department send patents to other organizations because it has a business need for doing so. You can allow specific departments to send sensitive information by creating exceptions for those departments in your DLP policy, or you can allow them to override a policy tip with a business justification.

Document Fingerprinting supports the same file types that are supported In transport rules. The Document Fingerprinting DLP agent does not detect sensitive information in password protected files, files that contain only images, and documents that do not contain all the text from the original form used to create the document fingerprint.

Using the EAC to Create a Document Fingerprint

To use document fingerprinting, simply upload a blank form, such as an intellectual property document, government form, or other standard form used in your organization. Then perform the following steps to add the resulting document fingerprint to a DLP policy or transport rule:

1. In the Exchange Administration Center EAC, go to **Compliance Management > Data Loss Prevention**.
2. Click Manage document fingerprints.
3. On the document fingerprints page, click the New Add icon to create a new document fingerprint.
4. Give the document fingerprint a Name and Description. (The name you choose will appear in the sensitive information types list.)
5. To upload a form, click the Plus sign (**+**).
6. Choose a form, and click **Open**. (Make sure that the file you upload contains text, is not password protected, and is in one of the File types that are supported in the transport rules (otherwise, you will receive get an error when you try creating the fingerprint.) Repeat for any additional files you want to add to the document list for this document fingerprint. You can also add or remove files from this document fingerprint later if you want.
7. Click **Save**.

The document fingerprint is now part of your sensitive information types, and you can add it to a DLP policy or add it to a transport rule.

Using PowerShell to Create Classification Rule Packages

DLP uses classification rule packages to detect sensitive content in messages. To create a classification rule package based on a document fingerprint, use the New-Fingerprint and **New-DataClassification** cmdlets. Because the results of **New-Fingerprint** aren't stored outside the data classification rule, you always run **New-Fingerprint** and **New-DataClassification** or **Set-DataClassification** in the same PowerShell session.

MCT USE ONLY. STUDENT USE PROHIBITED

Lesson 2
Designing and Implementing In-Place Hold

You might have legal or business requirements to preserve email data and to keep it in an unaltered form. You need to be able to preserve this data for a limited or unlimited amount of time. To provide data storage to meet compliance or operational requirements, Exchange Server 2013 includes a feature called In-Place Hold.

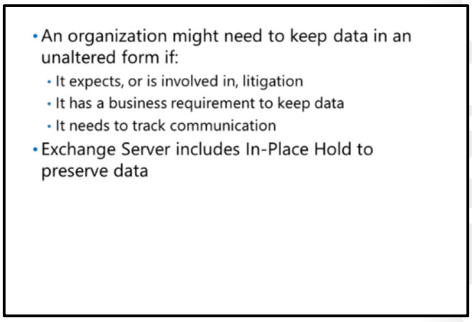
Lesson Objectives

After completing this lesson, you will be able to:

- Describe In-Place Hold scenarios.
- Explain how In-Place Hold works.
- Decide among options for implementing In-Place Hold.
- Configure In-Place Hold.
- Plan In-Place Holds.

Overview of Data Preservation Scenarios

Organizations that expect or are already involved in litigation might need to retain electronically stored information that is relevant to the litigation. Because organizations might need to keep data before a legal issue arises or is clearly defined, the scope of data preservation can be broad. For example, organizations might need to preserve all email that relates to a specific project or all email that is sent by or to specific employees who are involved in specific issues.

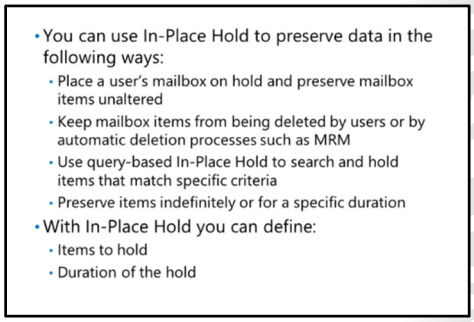


To achieve this, organizations can take several approaches. For example, they can ask users to preserve email messages by not deleting messages that relate to a specific issue. However, users might delete email either knowingly or inadvertently. Organizations can also suspend automatic deletion mechanisms, such as messaging records management (MRM). However, this suspension can cause large volumes of email to be stored in the users' mailboxes, which can reduce user productivity. Also, this approach doesn't prevent users from manually deleting email. Some organizations use archiving solutions to preserve data, but this can significantly increase cost. But if an organization fails to preserve data in email, the organization can be exposed to financial and legal risks.

For all of these reasons, you need a reliable and cost-effective way to preserve data from one or more mailboxes, while preventing users from changing or deleting data related to specific issues. Exchange Server 2013 provides several options. One of these options is In-Place Hold, which you can enable on particular mailboxes. You can use In-Place Hold, together with eDiscovery, to help prevent changes to a user's mailbox for a specific period of time.

MCT USE ONLY. STUDENT USE PROHIBITED

What is In-Place Hold?



In-Place Hold in Exchange Server 2013 is a successor to the litigation hold feature in Exchange Server 2010. Exchange Server 2013 provides the same functionality as litigation holds in Exchange Server 2010, with the additional functionality of In-Place Hold.

To start, identify the business cases that you want to address with In-Place Hold. In Exchange Server 2013, you can use In-Place Hold to accomplish the following goals:

- Place a user's mailbox on hold and preserve the mailbox items unaltered.
- Help prevent mailbox items from being deleted by users or by automatic deletion processes such as MRM.
- Search for and hold items matching specified criteria by using In-Place Hold.
- Preserve items indefinitely or for a specific duration.

You can also use In-Place Hold to place multiple holds on a user's mailbox for different cases or investigations. Also, you can implement In-Place Hold without telling the user, because you do not need to disable the MRM system. You can search all items that are on hold by using eDiscovery search, which is explained in the next lesson.

Compared to litigation hold, In-Place Hold has several advantages. With a litigation hold, you cannot select types of items to hold or the duration of the hold. In-Place Hold uses a different model and is much more precise. With In-Place Hold, you can specify following options:

- Items to hold. You can specify the types of items to place on hold. For example, you can define keywords and then hold only items that contain a keyword. You can also hold only messages from a specific sender. In other words, you can define a query based hold using several parameters. If you create a query-based In-Place Hold, you help to preserve all mailbox items that match the query, including existing items, items that are created after the In-Place Hold is created, and messages that are received at a later date.
- Duration of hold. You can define how long to apply an In-Place Hold. If the In-Place Hold duration is indefinite, the functionality is similar to that of a litigation hold.

Note: To put a mailbox on an In-Place Hold, you must have permission. By default, no one has this permission, not even an Administrator. To grant permission, assign the Legal Hold role to a user or a group.

How In-Place Hold Works

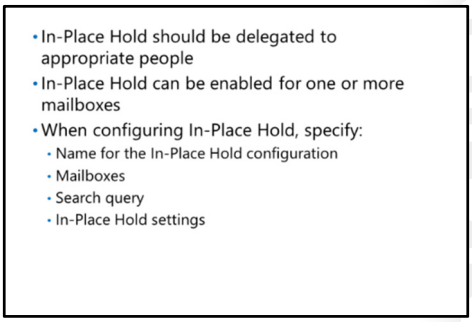
In-Place Hold is based on the mailbox dumpster folder that is also called Recoverable Items. This folder is located in the user's mailbox in the Non-IPM subtree, and it is not viewable through the user interface. The Recoverable Items folder is indexed, it can be searched, and you can prevent deletions from it by implementing In-Place Hold. With In-Place Hold in Exchange Server 2013, if a user deletes an item, the item is no longer marked only with a **ptagDeletedOnFlag** flag. Instead, the item goes to the Deletions subfolder within the Recoverable Items folder. From the Deletions folder, the user can retrieve items that were deleted. However, the user cannot permanently delete items from this folder. If the user deletes an item from Recoverable Items, the item goes to the Purges subfolder of Recoverable Items. The user

MCT USE ONLY STUDENT USE PROHIBITED

cannot access this item, but an administrator can, which helps prevent users from hiding or destroying items. Items in the Recoverable Items folder are not included in the user's mailbox quota. The Recoverable Items folder has its own quota, and two parameters apply to this quota: **RecoverableItemsWarningQuota** and **RecoverableItemsQuota**. Quotas for these values are set at the mailbox database level. By default, these parameters have no limit.

You can use In-Place Hold to place multiple holds on a user's mailbox. If you do, the search parameters of all In-Place Holds are applied together, by using a logical OR operator.

Options for Implementing In-Place Holds



In older versions of Exchange Server, you can configure hold options only by using the Exchange Management Shell. But in Exchange Server 2013, you can put a mailbox on hold by using the Exchange admin center. Administrators are not normally the people who put a mailbox on hold. Usually, the people who put a mailbox on hold are in the human resources or legal departments.

No one has the ability to put mailboxes on hold by default. To give someone this ability, put them in the Discovery Management security group in Active Directory Domain Services (AD DS). People in this group can also search through mailboxes. If you want to delegate only the ability to activate In-Place Hold and not the ability to search and perform queries, configure a new role group for this purpose. Also, if you want to separate users who can perform mailbox searches from users who can activate In-Place holds, you need to use two different groups.

After permissions are delegated, configure an In-Place Hold by using either the Exchange admin center or the Exchange Management Shell. In the Exchange admin center, you use the same interface to put a mailbox on In-Place Hold as you use to search for items on hold.

When configuring an In-Place Hold, configure following options:

- Name for the In-Place Hold configuration. This should be something descriptive.
- Mailboxes. You can choose one or more mailboxes in the Exchange Server organization to put on an In-Place Hold.
- Search query. You can define a query for the In-Place Hold. The result of this query is items that will be preserved. You can base the query on the following values:
- Keywords in mailbox items.
- Start or end date.
- **From** field.
- **To** or **CC** field.
- Item type. You can also choose to preserve all of the user's mailbox content, which is, in practice, the same as a litigation hold.
- In-Place Hold settings. Specify that you are placing the mailbox on hold instead of searching it. You can also choose how long to preserve items.

MCT USE ONLY. STUDENT USE PROHIBITED

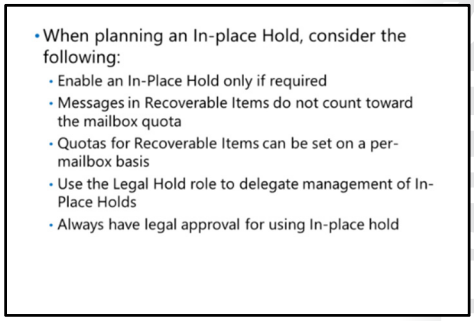
Demonstration: Configuring In-Place Hold

Demonstration Steps

1. Delegate the ability to configure an In-Place Hold to April by making her a member of the Discovery Management security group.
2. Sign in to the Exchange admin center as **Adatum\April**.
3. Configure an In-Place Hold for Amr Zaki that uses the following parameters:
   - Search criteria: Preserve all emails from administrator@adatum.com
   - Content to keep on hold: Only emails
   - Hold duration: 365 days

Considerations for Planning an In-Place Hold

Consider the following when planning an In-Place Hold:



- Enable In-Place Hold only if required. If you enable In-Place Hold for a large number of mailboxes, the mailbox database size can grow quickly, because messages cannot be deleted. Also, be sure to remove an In-Place Hold when it is no longer required. To remove an In-Place Hold by using the Exchange admin center, clear the option **Place content matching the search query in selected mailboxes on hold on search query**. To remove an In-Place Hold by using Exchange Management Shell, use the **Set-MailboxSearch** cmdlet.
- Messages that are in the Recoverable Items folder and that are removed from Deleted Items do not count toward the mailbox quota. You do not need to plan for user quotas if a litigation hold is enabled for a mailbox.
- You can set quotas for Recoverable Items on a per-mailbox basis. The **RecoverableItemsWarningQuota** is set to 20 gigabytes (GB) by default, and an event is generated in the Application log of the Mailbox server if the quota is reached. The **RecoverableItemsQuota** is set to 30 GB by default, and users cannot delete items from their Deleted Items folder if the quota is reached.

Note: The **RecoverableItemsQuota** default configuration is derived from a setting in the mailbox database that holds the mailbox.

- Use the Legal Hold role to delegate management of In-Place Holds. The manager who is responsible for designating which users are subject to an In-Place Hold may not want to share that information with Exchange Server administrators. You can delegate the ability to enable an In-Place Hold by using the Legal Hold role.
- Always get approval from your company's legal department before you implement an In-Place Hold, to make sure that you are not violating any compliance standards.

MCT USE ONLY. STUDENT USE PROHIBITED

Lesson 3
Designing and Implementing In-Place eDiscovery

Searching data in users' mailboxes has security, legal, and privacy implications, but sometimes a search is necessary. In-Place eDiscovery is a feature of Exchange Server 2013 that delegated and authorized people can use to search users' mailboxes. However, the nature of eDiscovery requires that you plan and implement it with care. This lesson explains how eDiscovery works and how you should plan for and implement it.

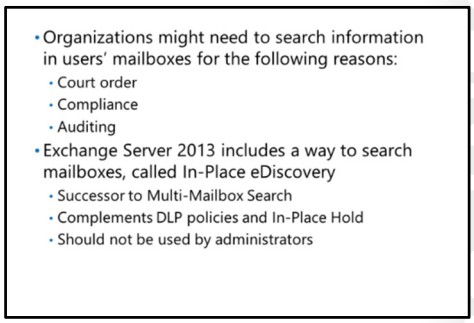
Lesson Objectives

After completing this lesson, you will be able to:

- Describe In-Place eDiscovery scenarios.
- Explain how In-Place eDiscovery works.
- Configure In-Place eDiscovery.
- Integrate eDiscovery with Microsoft SharePoint® Server 2013.
- Explain what to consider when planning In-Place eDiscovery.

Overview of In-Place eDiscovery Scenarios

Many organizations need to be able to search mailboxes for specific content when they are performing compliance audits. As part of a data loss prevention strategy, you might need a way to identify data in users' mailboxes that might violate the organization's compliance policy.



Exchange Server 2013 provides a way to search through users' mailboxes. This feature is called In-Place eDiscovery, and it is a successor to Multi-Mailbox Search in Exchange Server 2010. Authorized personnel can use In-Place eDiscovery to search one or more mailboxes in the Exchange organization and to see mailbox items resulting from the search query.

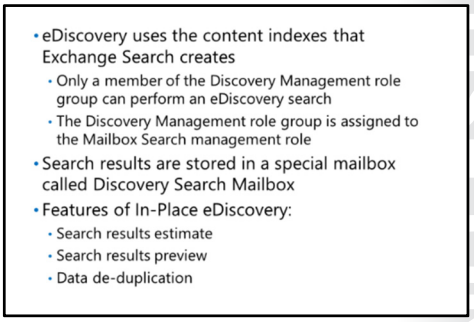
A search is usually performed only if there is a legal or business need for it. For example, a court might order a search for specific information in a mailbox. Or, an organization that suspects confidential information is being sent by email might use specific criteria to audit email traffic.

In-Place eDiscovery complements DLP policies and In-Place Hold. In-Place eDiscovery is reactive, and DLP policies and In-Place Hold are proactive. With In-Place eDiscovery, you search for emails that are already sent or that are placed on hold, but you cannot prevent information from being sent in the first place.

As with In-Place hold, eDiscovery is not a procedure that Exchange administrators should perform. Also, eDiscovery is not delegated to anyone by default.

MCT USE ONLY. STUDENT USE PROHIBITED

How In-Place eDiscovery Works



In Exchange Server 2013, the mailbox search functionality is available through the eDiscovery feature in the Exchange admin center. You can use eDiscovery to search multiple mailboxes for mailbox items, such as email messages, attachments, Calendar items, Tasks, and Contacts, across both primary and archive mailboxes. You can filter an eDiscovery search by sender, receiver, send or receive date, CC or BCC, and regular expressions.

eDiscovery uses the content indexes that the Exchange Search service creates. Having a single content-indexing engine ensures that additional resources are not used to crawl and index mailbox databases during the mailbox search.

Discovery Management Role

A user who belongs to the Discovery Management role group can perform an eDiscovery search. The Discovery Management role group is a universal security group that is created in AD DS when you install Exchange Server 2013. The Discovery Management role group is assigned to the Mailbox Search management role, which has permission to search all mailboxes in the organization. The Discovery Management role group is also assigned the Legal Hold management role.

Note: Exchange Server 2013 uses role-based access control (RBAC) to define what actions users can perform in the Exchange Server organization. RBAC uses management roles and management role groups to manage these permissions.

All search results are stored in a special mailbox called Discovery Search Mailbox. It is not possible to store results in any other mailbox. The Discovery Search Mailbox is created when you install Exchange Server 2013, and it cannot be used for standard purposes such as sending and receiving email, because delivery restrictions are applied to it. The user account associated with the Discovery Search Mailbox is disabled, so no one can log on to this mailbox without explicit permissions to do so. The Discovery Management group has full access rights to the Discovery Search Mailbox.

Because the Discovery Search Mailbox should be able to store a large amount of data, it is assigned a storage quota of 50 GB when it is created. If you have multiple teams or individuals that perform discovery searches and you do not want them to see results from other searches, create additional Discovery Search Mailboxes. You can create these extra mailboxes by using the Exchange Management Shell.

When you perform a search, a new folder is created in the Discovery Search Mailbox that has the same name as the search. Within that folder, a subfolder is created for each mailbox that was searched. Messages that the search returns are copied to that folder.

The eDiscovery search functionality in Exchange Server 2013 includes the following features:

- Search results estimate. In Exchange Server 2013, discovery managers can determine the number of items that an eDiscovery search will return before the items are copied to the selected discovery mailbox. Discovery managers are users who belong to the Discovery Management role group. Discovery managers can view the number of hits the specified keywords will return, and then they can modify the search query—if appropriate—before returned messages are copied to the discovery mailbox.

MCT USE ONLY. STUDENT USE PROHIBITED

- Search results preview. Before results are copied to Discovery Search Mailbox, the discovery manager can preview the results in the Exchange admin center.
- Data de-duplication. eDiscovery search includes an optional data de-duplication feature. When selected, eDiscovery search copies only a single instance of a message returned from multiple folders within the same mailbox or from different mailboxes. Do not select de-duplication if you want to see each instance of a message and its location.

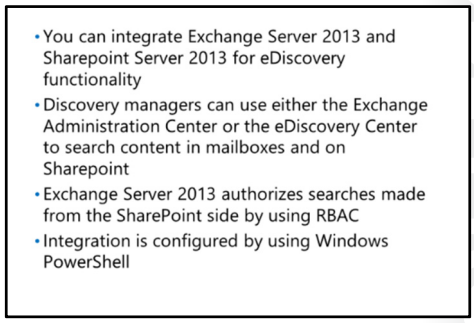
Demonstration: Configuring In-Place eDiscovery

Demonstration Steps

1. Log on to LON-CL1 as **Adatum\Allie**.
2. Start Outlook 2013 and send an email message to Aidan Delaney with the text: **This is my password: Pa$$w0rd.**
3. On LON-DC1, log on as **Administrator**. Open Active Directory Users and Computers, and then add April Reagan to the Discovery Management security group.
4. On LON-CAS1, sign in as **Administrator** to the Exchange admin center, and ensure that April is a member of the Discovery Management role group.
5. Log off from the Exchange admin center.
6. Log on to the Exchange admin center as **Adatum\April**.
7. Start a wizard to create a new eDiscovery search.
8. Specify to search through all mailboxes based on the keyword **password**.
9. Review the results.

Integrating eDiscovery with SharePoint Server 2013

A new feature of mailbox searches in Exchange Server 2013 is integration with SharePoint Server 2013. In addition to searching through users' mailboxes, you can also integrate Exchange Server 2013 with SharePoint Server 2013.



If you integrate these two products, users with the Discovery Management role can use the eDiscovery Center in SharePoint Server 2013 to do the following:

- Search and preserve content from a single location. A discovery manager can use this integration to search, and—if needed—preserve content in both SharePoint and Exchange. This content includes content in the Microsoft Lync® communication software, such as instant messaging conversations and shared meeting documents that are archived in Exchange Server mailboxes.
- Case management. The eDiscovery Center uses a case management approach to eDiscovery, so you can create cases and search and preserve content across different content repositories for each case.

MCT USE ONLY. STUDENT USE PROHIBITED

- Export search results. You can use the eDiscovery Center to export results from a search that you define. You can export mailbox content that is included in search results to a .PST file.

As with Exchange Server 2013, SharePoint Server 2013 uses Microsoft Search to index content and to query it. The discovery manager can therefore use either the Exchange admin center or the eDiscovery Center to search for the content, because the same results are returned. Searches performed from the SharePoint side are authorized by the Exchange Server RBAC mechanism. Because of this, the person who performs the search on SharePoint must have a mailbox on Exchange Server.

Configuring eDiscovery Integration

Before you can use the eDiscovery Center in SharePoint to search Exchange mailboxes, you must establish trust between the two applications. In Exchange Server 2013 and SharePoint Server 2013, you establish this trust by using Open Authorization (OAuth) authentication.

First, create Exchange Server 2013 as a trusted security token issuer in SharePoint Server 2013, by running the following cmdlet in the Windows PowerShell® command-line interface:

```
New-SPTrustedSecurityTokenIssuer -Name Exchange -MetadataEndPoint https://<Exchange
Server Name or FQDN>/autodiscover/metadata/json/1
```

Next, grant the Exchange service principal full control permissions to the SharePoint site subscription, by running the following commands:

```
$exchange=Get-SPTrustedSecurityTokenIssuer
$app=Get-SPAppPrincipal -Site http://<SharePoint ServerName> -NameIdentifier
$exchange.NameId
$site=Get-SPSite http://<SPServerName>
Set-SPAppPrincipalPermission -AppPrincipal $app -Site $site.RootWeb -Scope
sitesubscription -Right fullcontrol -EnableApplyOnlyPolicy
```

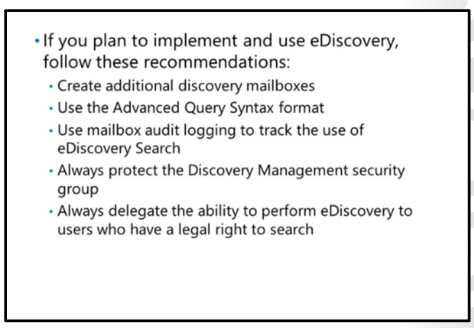
Next, configure the SharePoint partner application on the Exchange Server 2013 side, by running the following command on Exchange Server 2013:

```
cd c:\'Program Files'\Microsoft\'Exchange Server'\V15\Scripts
.\Configure-EnterprisePartnerApplication.ps1 -AuthMetadataUrl <path to SharePoint
AuthMetadataUrl> -ApplicationType SharePoint
```

Finally, add users who need to use SharePoint Server 2013 to perform eDiscovery searches to the Discovery Management role group in Exchange Server 2013.

Considerations for Planning eDiscovery

Similar to other compliance technologies discussed in this module, eDiscovery requires precise planning. Because misuse of eDiscovery can jeopardize users' privacy, it is very important that you plan eDiscovery so that it is used only in appropriate situations and is protected from unauthorized users.



When you plan to implement eDiscovery, consider the following guidelines:

- Create additional discovery mailboxes for distinct users or groups that perform discovery searches. This approach helps ensure that

MCT USE ONLY. STUDENT USE PROHIBITED

access to the results of the searches is limited to people who are authorized to perform the searches. For example, a team performing searches for legal purposes may have access to different mailboxes than help desk staff who are recovering deleted messages from mailboxes.

- Use the Advanced Query Syntax format to generate search queries that are more specific than the options provided in the basic user interface. If users perform many discovery searches, provide them with information about Advanced Query Syntax so they can search more efficiently.
- Use mailbox audit logging to track the use of eDiscovery. Mailbox audit logging is not enabled by default, and you must enable it on each mailbox. Mailbox audit logging can generate a significant amount of data, and you should enable it only if you need to.
- Always protect the Discovery Management security group in AD DS. You can protect membership in this group by using the Restricted Groups feature in Group Policy.
- Always delegate the ability to perform eDiscovery searches to users who have a legal right to search.

MCT USE ONLY. STUDENT USE PROHIBITED

Lab: Designing and Implementing Messaging Compliance

Scenario

The security and auditing teams at A. Datum have developed new policies that define requirements for messaging compliance. The teams have identified several categories of messages that must be retained within the messaging system, even if users delete the messages from their mailboxes. The teams have also identified specific rules for enabling auditors to search users' mailboxes and integrated SharePoint sites for messages. You need to plan and implement the messaging compliance configuration to meet the business requirements.

Objectives

After completing this lab, you will be able to:

- Design messaging compliance.
- Implement data loss prevention.
- Implement In-Place eDiscovery.

Lab Setup

Estimated Time: 60 minutes

Virtual machines	**20342B-LON-DC1** **20342B-LON-CAS1** **20342B-LON-MBX1** **20342B-LON-MBX2** **20342B-LON-CL1**
User name	Adatum\Administrator
Password	Pa$$w0rd

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

1. On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In Microsoft Hyper-V® Manager, click **20342B-LON-DC1**, and then in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Log on by using the following credentials:
    - User name: **Adatum\Administrator**
    - Password: **Pa$$w0rd**
5. Repeat steps 2, 3, and 4 for 20342B-LON-MBX1, 20342B-LON-MBX2 and 20342B-LON-CAS1.
6. For **20342B-LON-CL1**, repeat steps 1-3. Do not log on until directed to do so.

MCT USE ONLY. STUDENT USE PROHIBITED

Exercise 1: Designing Messaging Compliance

Scenario

As a part of its messaging security policy, A. Datum has defined following rules that must be implemented:

- No messages with financial information can leave the organization. Specifically, credit card numbers should be blocked. If anyone attempts to send this information in an email message, the administrator should be notified.
- Users can never send information about IP addresses in email.
- Email messages about ProjectX must be preserved in the mailboxes of users Amr Zaki, Brad Sutton, and Ed Meadows for at least two years. Email messages about this project are marked with the word *ProjectX* in the message subject and body.
- Members of the Auditing department must be able to search the contents of all mailboxes.
- Only members of the Auditing department can put mailboxes on a legal hold.

You must define and implement a DLP strategy that meets these requirements.

The main tasks for this exercise are as follows:

1. Design DLP rules that meet the security requirements
2. Design hold policies that meet the security requirements
3. Design the eDiscovery configuration

▶ Task 1: Design DLP rules that meet the security requirements

1. Read the exercise scenario requirements for DLP.
2. Propose a policy to address the requirements for DLP.

▶ Task 2: Design hold policies that meet the security requirements

1. Read the exercise scenario requirements for data preservation.
2. Propose a policy to address the requirements for a data hold.

▶ Task 3: Design the eDiscovery configuration

1. Read the exercise scenario requirements for a mailbox search.
2. Propose a policy to meet the requirements for a mailbox search.

Results: After completing this lab, you will have designed and implemented a DLP strategy.

Exercise 2: Implementing Data Loss Prevention

Scenario

Based on the design that you created, you implement the DLP policies to meet the requirements of the lab scenario and the Exercise 1 scenario.

The main tasks for this exercise are as follows:

1. Configure a DLP policy by using a template
2. Configure a DLP by using a custom rule
3. Validate the DLP deployment

MCT USE ONLY. STUDENT USE PROHIBITED

▶ Task 1: Configure a DLP policy by using a template

1. On LON-CAS1, open Internet Explorer, and then type **https://lon-cas1.adatum.com/ecp** to open the Exchange admin center.
2. Sign in as **adatum\administrator**, and then navigate to **compliance management**.
3. Create a new policy from a template. Name the policy **Prevent financial data flow**.
4. Click the **U.S. Financial Data** template.
5. Enforce the policy, and configure its **U.S. Financial: Scan email sent outside – high count** rule to reject messages that contain credit card numbers. The sender will receive a message that says, "This message contains financial information and can't be sent outside the organization," and the report should be sent to Administrator.
6. Activate and save the DLP policy.
7. Configure a general Policy Tip message to say, "This message contains information that you are not allowed to send."

▶ Task 2: Configure a DLP by using a custom rule

1. In the Exchange admin center, create a new custom DLP policy.
2. Name the policy **Prevent IP addresses**.
3. Configure the policy to prevent IP addresses from being sent in email.
4. Configure the policy to block messages that contain sensitive information unless the sender overrides the block with a business justification.
5. Configure the policy to send a report to the administrator if an email message violates the policy.
6. Configure the policy to reply to the sender with the following text: **You are not allowed to send an IP address in email.**
7. Activate the policy.

▶ Task 3: Validate the DLP deployment

1. Sign in to LON-CL1 as **Adatum\Aidan** with the password **Pa$$w0rd**.
2. From the Desktop, open File Explorer and browse to **C:\Files**. Open the file **Northwind Customer Data.** Examine the content of the file. Close the Microsoft Excel® spreadsheet software.
3. Start Outlook 2013. Configure an email account automatically if necessary.
4. Send an email message to **Ben@contoso.com** with the subject **Northwind data** and attach the file **C:\Files\Northwind Customer Data.xlsx**. Type **Find attached data** in the message body, and then send the message. Examine the content of the policy tip.
5. Sign in to Outlook Web App as **Adatum\Administrator** with the password of **Pa$$w0rd**.
6. Ensure that the administrator has received a report about the message that Aidan tried to send and then close Outlook Web App.
7. On LON-CL1, from Outlook, send another email message to **Ben@contoso.com**, with the subject **My IP** and the following content: **This is my IP address: 172.16.0.100**. Wait for a few moments before you send email for policy tip to appear. After you read the content of the policy tip, try to send the message.
8. Ensure that you receive a message that the message cannot be sent.
9. Leave **Outlook 2013** open, and stay logged on as Aidan on LON-CL1.

MCT USE ONLY. STUDENT USE PROHIBITED

Results: After completing this exercise, you will have implemented DLP.

Exercise 3: Implementing In-Place eDiscovery

Scenario

Based on a design provided in Exercise 1, implement In-Place eDiscovery to meet the security and auditing requirements. Specifically, you need to identify whether users are exchanging confidential information about a new project with the Contoso Corporation. Give auditor April Reagan permission to perform the searches.

The main tasks for this exercise are as follows:

1. Send emails between users
2. Configure permissions required for In-Place eDiscovery
3. Perform an In-Place eDiscovery search
4. Configure In-Place Hold
5. To prepare for the next module

▶ Task 1: Send emails between users

1. Logged on to LON-CL1 as **Adatum\Aidan**, open Outlook 2013.
2. Send a message to Bill Malone with the following content: **It seems like the company won the project for delivering tools to Contoso. We must make sure that we take advantage of this information before authorized personnel do. Let me know what you think.**
3. Open Outlook Web App and sign in as **Adatum\Bill** with the password of **Pa$$w0rd**.
4. Read the message from Aidan, and then reply with the following content: **We must meet with Contoso people as soon as possible. Can you keep this confidential?**
5. Close Outlook Web App.

▶ Task 2: Configure permissions required for In-Place eDiscovery

1. On LON-DC1, from Server Manager, open **Active Directory Users and Computers**.
2. Add **April Reagan** to the **Discovery Management** role group.
3. Open the Exchange admin center as **Adatum\Administrator**, and then ensure that **April** is a member of the **Discovery Management** role group.

▶ Task 3: Perform an In-Place eDiscovery search

1. On LON-CL1, open Internet Explorer, and then type **https://lon-cas1.adatum.com/ecp** to open the **Exchange admin center**, and then sign in as **Adatum\April** with the password of **Pa$$w0rd**.
2. Start a wizard to perform a new eDiscovery search.
3. Name the search **Contoso search**.
4. Choose to search mailboxes of Aidan Delaney and Bill Malone.
5. Define a filter based on keywords, and then type the following:

(Contoso NEAR(3) tools) AND (confidential OR authorized)

6. Choose to search only emails.

MCT USE ONLY. STUDENT USE PROHIBITED

7. Finish the wizard, and then wait until the search finishes.
8. Preview the search results.

▶ Task 4: Configure In-Place Hold

1. Ensure that you are signed in as April to the Exchange admin center.
2. Create a new search.
3. Specify the following mailboxes to search: Amr Zaki, Brad Sutton, and Ed Meadows.
4. Base the search on the keyword **ProjectX**.
5. Place items that the search finds on hold for 720 days.

   Note: After you configure mailboxes for In-Place Hold, you can search for deleted or modified items in these mailboxes by using same procedure for eDiscovery search.

▶ Task 5: To prepare for the next module

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

1. On the host computer, start Hyper-V® Manager.
2. In the **Virtual Machines** list, right-click **20342B-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 to 3 for 20342B-LON-CAS1, 20342B-LON-MBX1, 20342B-LON-MBX2, and 20342B-LON-CL1.
5. In Hyper-V Manager, click **20342B-LON-DC1**, and in the Actions pane, click **Start**.
6. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
7. Sign in using the following credentials:
   a. User name: **Adatum\Administrator**
   b. Password: **Pa$$w0rd**
8. Repeat steps 5 to 7 for **20342B-LON-MBX1, and 20342B-LON-CAS1**.

Results: After completing this exercise, you will have configured eDiscovery.

Exercise 4: Comparing Messaging Policy and Compliance Options

Scenario

This is a discussion-based activity. The preceding three exercises cover various ways to implement different types of messaging compliance. In this instructor-led activity, you review some of the compliance requirements from previous labs, and then you discuss the different ways that you can meet the compliance requirements with Exchange Server 2013.

The main tasks for this exercise are as follows:

1. Discuss alternative solutions

▶ Task 1: Discuss alternative solutions

1. Review the solution that you provided in Exercises 1, 2, and 3. Discuss with the instructor and the other students any alternative solutions that produce a result that is similar or the same.

MCT USE ONLY. STUDENT USE PROHIBITED

2. Review the usage scenarios for DLP policies, In-Place Holds, and In-Place eDiscovery.
3. Which compliance situations does each of these technologies address?
4. Which technologies have overlapping features?
5. How should you decide which technology to use?

Results: After completing this exercise, students will have discussed alternative solutions for messaging policy and compliance options.

Question: When should you use custom DLP policies instead of policies based on templates?

Question: How can you notify users that they are about to violate DLP policy before they actually send an email message?

MCT USE ONLY. STUDENT USE PROHIBITED

Module Review and Takeaways

Best Practice

- Always define business and legal requirements before implementing DLP strategies.
- Review the rules that a DLP policy template contains before you apply a DLP policy template.
- Always test DLP policies with Policy Tips before you enforce them.
- Avoid applying an In-Place Hold on all mailbox items.

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
DLP policy does not block email messages that contain sensitive information	
No compliance management in the feature pane	

Review Question(s)

Question: What is data loss prevention?

Question: What is the main purpose of In-Place Hold?

Question: How do you give someone permission to perform an eDiscovery search?

MCT USE ONLY. STUDENT USE PROHIBITED

MCT USE ONLY. STUDENT USE PROHIBITED

Module 8

Designing and Implementing Administrative Security and Auditing

Contents:

Module Overview **08-1**

Lesson 1: Designing and Implementing RBAC **08-2**

Lesson 2: Designing and Implementing Split Permissions **08-14**

Lesson 3: Planning and Implementing Audit Logging **08-19**

Lab: Designing and Implementing Administrative Security and Auditing **08-25**

Module Review and Takeaways **08-31**

Module Overview

Administrative security and auditing are key elements of any Microsoft® Exchange Server implementation, and they help prevent unauthorized access or system modification. You want to make sure that only authorized users who are well trained and who understand configuration problems can change your Exchange Server configuration. If an unauthorized or untrained user changes a configuration setting that causes the deletion of message databases, your users will be very unhappy.

For that reason, Exchange Server offers role-based access control (RBAC) and includes various options for auditing administrative access. This module explains how to design and implement auditing and administrative security.

Objectives

After completing this module, you will be able to:

- Design and implement RBAC.
- Design and implement split permissions.
- Plan and implement audit logging.

MCT USE ONLY. STUDENT USE PROHIBITED

Lesson 1
Designing and Implementing RBAC

Exchange Server 2013 uses the RBAC permissions model to restrict which administrative tasks each user may perform on Exchange servers. With RBAC, you control the resources that administrators can configure and the features that users can access. This lesson describes how to design and implement RBAC permissions in Exchange Server 2013.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe RBAC.
- Describe management role groups.
- Identify built-in management role groups.
- Discuss guidelines for implementing management role groups.
- Manage RBAC permissions.
- Design custom role groups.
- Configure custom role groups.
- Troubleshoot RBAC.
- Describe management role assignment policies.
- Describe built-in user roles for role assignment policies.
- Describe how to work with management role assignment policies.
- Configure management role assignment policies.

What Is RBAC?

RBAC is a permissions model that was first introduced in Exchange Server 2010. With RBAC, you do not need to modify and manage access control lists (ACLs) on Exchange Server or Active Directory® Domain Services (AD DS) objects. In Exchange Server 2013, RBAC controls the administrative tasks that users can perform and the extent to which they can administer their own mailboxes and distribution groups.

- RBAC defines all Exchange Server 2013 permissions, and all Exchange Server management tools use it
- RBAC defines which cmdlets the user can run:
  - Who: Can modify objects
  - What: Objects and attributes can be modified
  - Where: Scope or context of objects that can be modified
- RBAC options include:
  - Management role groups
  - Management role assignment policies
  - Direct policy assignment (avoid using)

When you configure RBAC permissions, you can define precisely which Exchange Management Shell cmdlets a user can run and which objects and attributes the user can modify.

Both Exchange Server administration tools, namely Exchange Administration Center and Exchange Management Shell, use RBAC to determine user permissions. Therefore, permissions are consistent regardless of which tool you use.

If RBAC allows the cmdlet to run, the cmdlet actually runs in the security context of the Exchange Trusted Subsystem and not the context of the user. The Exchange Trusted Subsystem is a highly privileged

MCT USE ONLY. STUDENT USE PROHIBITED

universal security group that has read/write access to every Exchange Server–related object in the Exchange organization. The Exchange Trusted Subsystem is also a member of the Administrators local security group and the Exchange Windows Permissions universal security group, which enables Exchange Server 2013 to create and manage AD DS objects.

RBAC Options

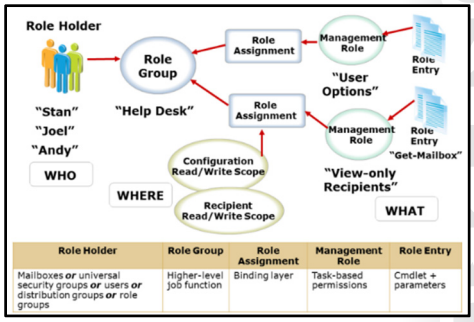
RBAC assigns permissions to users in two primary ways, depending on whether the user is an administrator or an end user:

- Management role groups. RBAC uses management role groups to assign permissions to administrators. These administrators may require permissions to manage the Exchange organization or some part of it. Some administrators may require limited permissions to manage specific Exchange Server features, such as compliance or specific recipients. To use management role groups, add users to the appropriate built-in management role group or to a custom management role group. RBAC assigns each role group one or more management roles that define the precise permissions that RBAC grants to the group.
- Management role assignment policies. RBAC uses management role assignment policies to assign management roles to end users. Role assignment policies consist of roles that control what users can do with their mailboxes or distribution groups. These roles do not allow users to manage any features that the users are not associated with directly.

Note: You can also use direct role assignment to assign permissions. Direct role assignment is an advanced method for assigning management roles directly to a user or to a universal security group, without the need to use a role group or role assignment policy. Direct role assignments are useful if you need to provide a granular set of permissions to a specific user only. However, we recommend that you avoid using direct role assignment, because it is significantly more complicated to configure and manage.

What Are Management Role Groups?

A *management role group* is a universal security group that simplifies the process for assigning management roles to a group of users. All members of a role group are assigned the same set of roles. In Exchange Server 2013, groups, such as organization management and recipient management, are assigned administrator and specialist roles that define major administrative tasks. By using role groups, you can more easily assign a broader set of permissions to a group of administrators or specialist users.



You use management role groups to assign administrator permissions to groups of users. To understand how management role groups work, you need to understand their components.

Components of Management Role Groups

Management role groups use several underlying components to define how RBAC assigns permissions:

- *Role holder*. A user or security group that you can add to a management role group. For example, when a user becomes a member of a management role group, RBAC grants the user all of the

MCT USE ONLY. STUDENT USE PROHIBITED

permissions that the management roles provide. You can add user accounts to the group in AD DS, or you can use the **Add-RoleGroupMember** cmdlet.

- *Management role group*. A universal security group that contains users or groups that are role-group members. Management role groups are assigned to management roles. The combination of all of the roles assigned to a role group defines everything that users who are added to a role group can manage in the Exchange organization.
- *Management role*. A container for a group of management role entries. These entries define the tasks that users can perform if RBAC assigns them the management role by using management role assignments.
- *Management role entry*. A cmdlet, including its parameters, which you add to a management role. By adding a cmdlet to a role as a management role entry, you are granting rights to manage or view the objects associated with that cmdlet.
- *Management role assignment*. Assigns a management role to a role group. After you create a management role, you must assign it to a role group so that the role holders can use it. If you assign a management role to a role group, role holders can use the cmdlets that the management role defines.
- *Management role scope*. The scope of influence or impact that the role holder has after RBAC assigns a management role. When you assign a management role, you use management scopes to target which objects that role controls. Scopes can include servers, databases, organizational units (OUs), and recipient objects, and more.

Examples of Management Role Groups

Management role groups define who can perform specific tasks and the scope within which administrators can perform those tasks. For example, you can use RBAC to assign permissions as the following table shows.

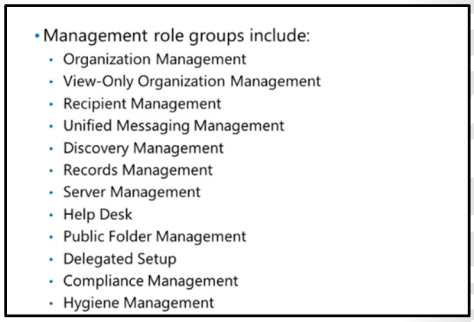
Role holder	Management role group	Management role	Management role entries	Management role scope
Stan	Organization Management	Organization Management	All Exchange cmdlets	Organization
Joel	Help Desk	HelpDesk	Cmdlets related to mailbox and user account management	Organization
Andy	Sales Admins	SalesAdminRole	Cmdlets related to Recipient management only	Sales department OU in AD DS or Active Directory

MCT USE ONLY. STUDENT USE PROHIBITED

Built-In Management Role Groups

Exchange Server 2013 includes several built-in role groups that you can use to provide varying levels of administrative permissions to user groups. You can add users to, or remove users from, any built-in role group. You also can add role assignments to or remove role assignments from most role groups.

The following table describes each of the built-in role groups.



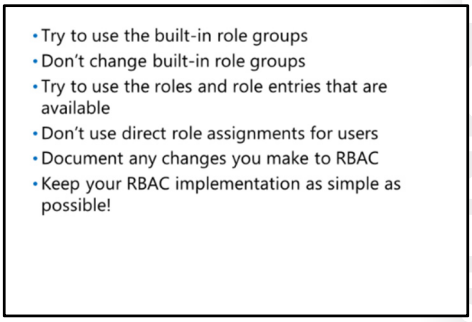
Role group	Description
Organization Management	Role holders can access the entire Exchange organization and can perform almost any task against any Exchange Server object.
View-Only Organization Management	Role holders can view the properties of any object in the organization.
Recipient Management	Role holders have access to create or modify Exchange recipients within the Exchange organization.
UM Management	Role holders can manage the Unified Messaging features within the organization, such as Unified Messaging server configuration, properties on mailboxes, prompts, and auto-attendant configuration.
Discovery Management	Role holders can perform searches of mailboxes in the Exchange organization for data that meets specific criteria.
Records Management	Role holders can export audit logs, and they can configure compliance features, such as retention policy tags, message classifications, and transport rules.
Server Management	Role holders have access to the Exchange Server configuration settings, such as database copy locations, certificates, and transport queues. They cannot administer the recipient configuration.
Help Desk	Role holders can perform limited recipient management.
Public Folder Management	Role holders can manage public folders and databases on Exchange servers.
Delegated Setup	Role holders can deploy previously provisioned Exchange servers.
Compliance Management	Role holders can configure and manage compliance settings. This role group is new in Exchange Server 2013.
Hygiene Management	Role holders can manage Exchange Server anti-spam and antimalware features, and they can grant permissions for antivirus products to integrate with Exchange Server. This role group is new in Exchange Server 2013.

MCT USE ONLY. STUDENT USE PROHIBITED

Note: All of these role groups are located in the Microsoft Exchange Server Security Groups OU in AD DS.

Guidelines for Implementing Management Role Groups

The built-in role groups in Exchange Server can meet the permission delegation needs of most organizations. However, there are situations that might require you to change or modify role groups, such as special security requirements from your auditing department. Before you start modifying the role groups, consider the following guidelines for implementing management role groups:



- Try to use the built-in role groups first. Always use the built-in role groups first to apply permissions, and then decide if you need to modify the role group to assign or remove permissions.
- Don't change built-in role groups. If you want to modify a role group, such as if you want to apply a scope to a built-in role group, copy it to a new role group, and then configure a scope. This approach helps ensure the original built-in group is still configured with the default settings.
- Try to use the roles and role entries that are available. You should create your own roles or role entries only when necessary, in order to reduce complexity and to keep an overview of the permissions you configured.
- Don't use direct role assignments for users. We recommend that you create a role group and add the user there, because a direct role assignment is hard to find.
- Always thoroughly document any changes you make to RBAC. If you decide to change roles, role assignments, scopes, or role entries, write down all changes so that all administrators understand what you changed in RBAC. It might be hard for somebody new to understand your RBAC if you don't explain to them how it was implemented.

Note: As a best practice, keep your RBAC implementation as simple as possible. The more you change it, the harder it becomes to manage.

Demonstration: Managing Permissions by Using the Built-In Role Groups

In this demonstration, you will review how to manage RBAC permissions in Exchange Server 2013 by using the built-in role groups. You will see how to add users to the built-in role groups and how RBAC assigns the resulting permissions to the user accounts.

Demonstration Steps

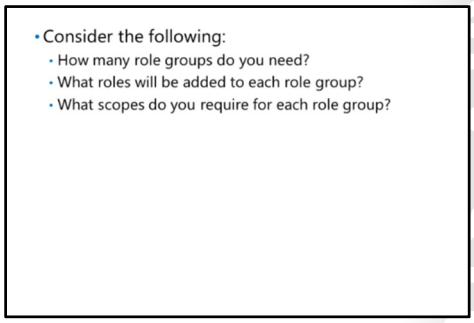
1. On LON-DC1, open Active Directory Users and Computers.
2. Add **Tony** to the **Recipient Management** group located in Microsoft Exchange Security Groups OU.
3. Switch to the LON-CAS1 virtual machine. In the Exchange Administration Center, sign in as **Adatum\Tony** with the password **Pa$$w0rd**.

MCT USE ONLY. STUDENT USE PROHIBITED

4. In the feature pane, view **Servers**. Notice that Tony has Read access to the Exchange organization configuration. He has this access because the Recipient Management group has been granted implicit Read permission.
5. In the feature pane, click **permissions**. Notice that no tabs for administrator or user groups are available for Tony.
6. In **recipients feature**, verify that you can modify the user properties of **Adam Barr**.

Designing Custom Role Groups

In addition to the built-in role groups, you also can create custom role groups to delegate specific permissions within the Exchange Server organization. Use this option if your ability to limit permissions is beyond the scope of the built-in role groups.



RBAC enables complete flexibility in how you assign permissions in an Exchange Server 2013 environment. In order to design custom role groups, consider the following questions:

- How many role groups do you need? Always start with a single one, and then add another one if that one is not enough.
- What roles will be added to each role group? Decide what roles you want to add to the custom role group, try to use the built-in roles, and create custom roles only if the built-in roles are not suitable.
- What scopes do you require for each role group? Define a scope, such as a database scope, only if you require it.

For example, you can use RBAC to assign permissions to a group of administrators in a branch office who only need to manage recipient tasks for branch-office users and mailboxes on branch office Mailbox servers. To implement this scenario, do the following:

1. Create a new role group, and then add the branch office administrators to the role group. You can use the **New-RoleGroup** cmdlet to create the group. When you create the group, you must specify the management roles. You also can specify the management scope for the role.
2. Assign management roles to the branch office administrators. To delegate permissions to a custom role group, you can use one or more of the default built-in management roles, or you can create a custom management role that is based on one of the built-in management roles. Exchange Server 2013 includes approximately 70 built-in management roles that provide fine-grained levels of permissions. To view a complete list of all of the management roles, use the **Get-ManagementRole** cmdlet. To view detailed information about a management role, type **Get-ManagementRole rolename | FL**, and then press Enter.

Note: You also can configure a new management role rather than use one of the existing management roles. To do this, use the **New-ManagementRole** cmdlet to create a custom management role based on an existing one. You can then add and remove management role entries as needed. By default, the new management role inherits all of the permissions assigned to the parent role. You can remove permissions from the role as necessary, by using the **Remove-ManagementRoleEntry** cmdlet. However, it can be complicated to create a new

MCT USE ONLY. STUDENT USE PROHIBITED

management role or to remove unnecessary management role entries, so we recommend that you use an existing role whenever possible.

3. Identify the management scope for the management role. For example, in the branch office scenario, you can create a role assignment with an OU scope that is specific to the branch office OU.
4. Create the management role group by using the information that you collect. Use the **New-RoleGroup** cmdlet to create the link between the role group, the management roles, and the management scope. For example, consider the following cmdlet:

```
New-RoleGroup – Name BranchOfficeAdmins –roles “Mail Recipients”, “Distribution
Groups”, “Move Mailboxes”, “Reset Password”, “Mail Recipient Creation” –Members
BranchOfficeAdmins -RecipientOrganizationalUnitScope Adatum.com/BranchOffice
```

These tasks perform the following:

- Create a new role group named BranchOfficeAdmins.
- Assign the Mail Recipients, Distribution Groups, Move Mailboxes, Reset Password, and Mail Recipient Creation management roles to the BranchOfficeAdmins role group.
- Configure a management role scope that is limited to the BranchOffice OU in the Adatum.com domain.

Demonstration: Configuring Custom Role Groups

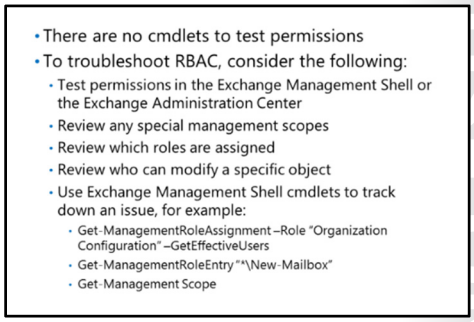
In this demonstration, you will see how to create a custom role group, add members and roles to the role group, and test that permissions are applied correctly.

Demonstration Steps

1. On LON-CAS1, in the Exchange Admin Center, create a new role group and configure it as follows:
    - Name: **MarketingAdmins**
    - Write scope: Click **Organizational Unit**, and then type **adatum.com/Marketing**
    - Roles: **Add Mail Recipients** and **Mail Recipient Creation**
    - Members: **Andreas Schou**
2. On LON-MBX1, in **Active Directory Users and Computers**, in **Adatum\ Microsoft Exchange Security Groups OU**, verify that the **MarketingAdmins** group exists and that **Andreas** is a member of it.
3. Open Windows® Internet Explorer®, and then connect to **https://LON-CAS1.adatum.com/ECP**. Sign in as **Adatum\Andreas** with the password **Pa$$w0rd**.
4. In **recipients**, double-click **Anil Elson**. Note that all fields are grey, indicating that you do not have permission to change this user because the user is not in the Marketing OU.
5. Modify one of the properties of **Alan Steiner**.
6. Create a new mailbox in the default **Users** container. Verify the error message that the Users OU is not in your write scope. Change the OU to **Marketing**. Verify that the mailbox is successfully created.

MCT USE ONLY. STUDENT USE PROHIBITED

Troubleshooting RBAC



Troubleshooting permissions in RBAC is not a trivial task. First, make sure that you understand how RBAC is working and that it consists of the following components: role groups, roles, role entries, scopes, and role assignments. The previous topics in this module explain this. Also, keep in mind that no test cmdlet is available to test effective permissions. The only way to test them is to look into the system that is configured in RBAC and try it. To troubleshoot RBAC, consider the following:

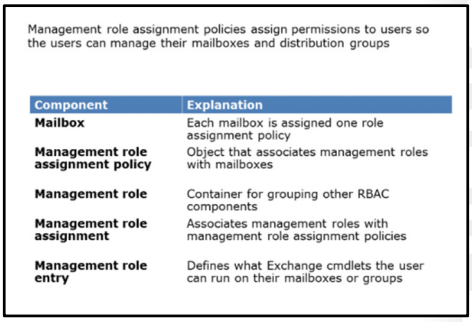
- Test permissions in the Exchange Management Shell or the Exchange Administration Center.
  To do this, sign in to the Exchange Management Shell or the Exchange Administration Center with the user account that you want to troubleshoot. Then use, for example, the **Get-Command** cmdlet in the Exchange Management Shell to retrieve the list of cmdlets that the user can run. If the cmdlet is not in the list, the user does not have permission to run it. Next, run the cmdlets that are available with the **–WhatIf** parameter to test whether there are permission errors when you run them.
- Review any special management scopes that are configured. Before you dig deeper in the system, review any management scopes that are defined in your Exchange organization. You can list all of the configured scopes by using the **Get-ManagementScope** cmdlet.
- Review what roles are assigned. To learn why an administrator might not have a specific permission, find out what roles are assigned to them. To do this, run the **Get-ManagementRoleAssignment -RoleAssignee <user>** cmdlet. Then, you can use the **Get-ManagementRoleEntry** cmdlet to investigate each role and role entry.
- Review who can modify a specific object. To troubleshoot permissions on an object, you can also find out who has permission to modify it. For example, if you are investigating who can modify Tony's mailbox, you can run the **Get-ManagementRoleAssignment -WritableRecipient Tony –GetEffectiveUsers** cmdlet.
- Use Exchange Management Shell cmdlets to track down the issue. You can use the following cmdlets:
- **Get-ManagementRoleAssignment**. Shows all assignments and allows you to filter, for example, roles or role groups. For example, run the **Get-ManagementRoleAssignment –Role "Organization Configuration" –GetEffectiveUsers** cmdlet to find which users have permission to configure the Exchange organization.
- **Get-ManagementRole**. Displays all roles configured in the system. For example, to display all roles that include the cmdlet **New-InboxRule**, run the **Get-ManagementRole –cmdlet "New-InboxRule"** cmdlet.
- **Get-ManagementRoleEntry**. Shows all role entries that are assigned. The syntax is **Get-ManagementRoleEntry <Role>\<Cmdlet>**. For example, to identify all roles that include the cmdlet **New-Mailbox**, run the **Get-ManagementRoleEntry "*\New-Mailbox"** cmdlet.
- **Get-Management Scope**. Displays all custom management scopes.
- **Get-RoleGroup**. Shows the management role groups.
- **Get-RoleGroupMember**. Lists all members of a management role group. For example, to list all members of the management role group Organization Management, run the **Get-RoleGroupMember "Organization Management"** cmdlet.

MCT USE ONLY. STUDENT USE PROHIBITED

Note: If you are working in a multi-domain environment, make sure that the Exchange Management Shell is configured so you can view the entire forest. Otherwise, you cannot see, for example, role groups. To configure this, run the **Set-ADServerSettings –ViewEntireForest $true** cmdlet.

What are Management Role Assignment Policies?

Management role assignment policies associate end-user management roles with users. You do not configure administrative permissions with management role assignment policies. Instead, you use management role assignment policies to configure what changes users can make to their mailbox settings and to distribution groups that they own. Every user with an Exchange Server 2013 mailbox receives a role assignment policy by default. You can do the following:


Management role assignment policies assign permissions to users so the users can manage their mailboxes and distribution groups

Component	Explanation
Mailbox	Each mailbox is assigned one role assignment policy
Management role assignment policy	Object that associates management roles with mailboxes
Management role	Container for grouping other RBAC components
Management role assignment	Associates management roles with management role assignment policies
Management role entry	Defines what Exchange cmdlets the user can run on their mailboxes or groups

- Decide which role assignment policy to assign by default.
- Choose what to include in the default role assignment policy.
- Override the default policy for specific mailboxes.

In Exchange Server 2013, you can use the Exchange Administration Center to view and modify the default management role assignment policy and to configure additional management role assignment policies with different permissions. If you create a custom management role assignment policy, you must assign it to the applicable mailboxes.

Components

Role assignment policies consist of the following components, which define what users can do with their mailboxes:

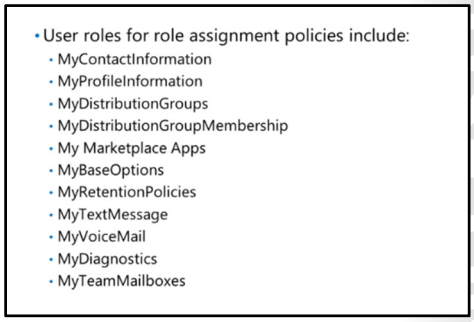
- Mailbox. Assigned a single role assignment policy. When a mailbox is assigned a role assignment policy, the policy is applied to the mailbox. This assignment grants the mailbox all of the permissions that the management roles provide.
- Management role assignment policy. An object in Exchange Server 2013. Users become associated with a role assignment policy when you create their mailboxes or when you change the role assignment policy on their mailboxes. The combination of all of the roles in a role assignment policy defines everything that the associated users can manage on their mailboxes or distribution groups.
- Management role assignment. A link between management roles and role assignment policies. Assigning a management role to a role assignment policy grants users the ability to use the cmdlets in the management role. When you create a role assignment, you cannot specify a scope. The scope that the assignment applies is based on the management role, and it is either Self or MyGAL.
- Management role. A container for a group of management role entries. Roles define the specific tasks that users can do with their mailboxes or distribution groups.
- Management role entry. A cmdlet, script, or special permission that enables users to perform a specific task. Each role entry consists of a single cmdlet and the parameters that the management role can access.

MCT USE ONLY. STUDENT USE PROHIBITED

Built-in User Roles for Role Assignment Policies

Similar to built-in role groups, there are also built-in user roles for role assignment policies. These user roles can be enabled or disabled for each role assignment policy to configure user self-management.

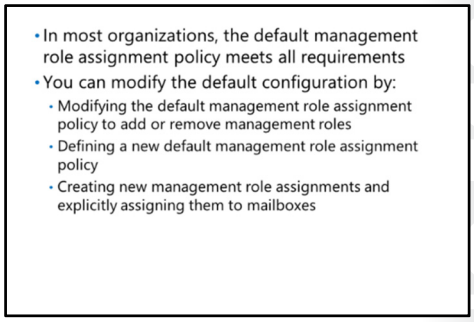
The following table describes each of the built-in user roles.



User role	Description of what end users can do	Enabled by default
MyContactInformation	Modify the user's own contact information. Can be divided into address, mobile, and personal information.	Yes
MyProfileInformation	Change their name and display name.	No
MyDistributionGroups	Create, modify, and view distribution groups, and manage members of groups they own.	No
MyDistributionGroupMembership	View and modify their membership in distribution groups, such as by joining or leaving a group.	Yes
My Custom Apps	View and modify their custom applications.	No
My Marketplace Apps	View and modify their marketplace applications.	Yes
MyBaseOptions	View and modify basic configuration, and access Microsoft Outlook® Web App settings.	Yes
MyRetentionPolicies	View and modify retention tags and retention tag settings.	No
MyTextMessage	Manage options for text messaging.	Yes
MyVoiceMail	Manage options for voice mail.	Yes
MyDiagnostics	Perform basic diagnostics on their mailbox.	No
MyTeamMailboxes	Create site mailboxes, and connect them to Microsoft SharePoint® services sites.	Yes

MCT USE ONLY. STUDENT USE PROHIBITED

Working with Management Role Assignment Policies



Exchange Server 2013 includes a default role assignment policy that grants end users with the most commonly used permissions. For most organizations, you do not need to modify this policy configuration. However, you can change the management role assignment policy if your organization has different requirements regarding how users can interact with their mailboxes or groups.

To view the default management role assignment policy configuration, use the **Get-ManagementRoleAssignment –RoleAssignee "Default Role Assignment Policy"** cmdlet. This cmdlet lists all of the management roles that are assigned to the default role assignment policy. To view the details of each management role, use the **Get-ManagementRoleEntry "<rolename>*"** cmdlet. For example, the **Get-ManagementRoleEntry "MyBaseOptions*"** cmdlet displays all management role entries that are associated with the MyBaseOptions management role.

Working with Assignment Policies

You can modify the default role-assignment configuration in several ways:

- Change the default permissions on the default role assignment policy by adding or removing management roles. For example, if you want to enable all users to create and manage groups, you can identify the management role that grants them the necessary permissions, and then add that role to the default role assignment policy.
- Define a new role assignment, and then configure that role assignment to be the default for all mailboxes. Use the **Set-RoleAssignmentPolicy** cmdlet to replace the built-in default role assignment policy with your own. When you do this, RBAC assigns the role assignment policy that you specify to new mailboxes by default.

Note: When you define a new role assignment policy, RBAC does not assign it automatically. You need to use the **Set-Mailbox** cmdlet to update previously created mailboxes to the new default role assignment policy.

- Configure additional role assignment policies, and then assign them to a mailbox manually by using the **RoleAssignmentPolicy** parameter on the **New-Mailbox**, **Set-Mailbox**, or **Enable-Mailbox** cmdlets. When you assign an explicit role assignment policy, the new policy takes effect immediately and replaces the previously assigned explicit role assignment policy. If you have many different user groups with specific needs, you can create role assignment policies for each group.

Demonstration: Configuring Management Role Assignment Policies

In this demonstration, you will see how to:

1. Use the Outlook Web App settings with default user permissions.
2. Change the default role assignment policy.
3. Verify the changed user permission in the Outlook Web App settings.

MCT USE ONLY. STUDENT USE PROHIBITED

Demonstration Steps

1. On LON-CAS1, sign in to the Exchange Admin Center as **Kelly**, and then modify the street and city in **contact location**.
2. On LON-MBX1, in the Exchange Admin Center, in **user roles**, modify **Default Role Assignment Policy** as follows:
   - Clear **MyContactInformation**. The user no longer has permission to modify their contact information in Outlook Web App.
   - Select **MyDistributionGroups**. With this permission, the user can create groups and manage them.
3. Open Kelly's mailbox and verify that you cannot modify the street and city any longer. Create a new distribution group to verify that the permissions you configured are working correctly.

MCT USE ONLY. STUDENT USE PROHIBITED

Lesson 2

Designing and Implementing Split Permissions

Generally speaking, Exchange Server administrators do not automatically manage the users or group accounts in Active Directory. Especially in large Exchange Server organizations, there is a clear differentiation between objects that relate to Exchange Server, such as the mailbox, and objects based on AD DS, such as user or group objects. For example, the AD DS administrator creates the user object according to the company's standards, and then the Exchange Server administrator creates a mailbox for that user. By default, Exchange administrators can also create AD DS objects.

Exchange Server 2013 provides a split permissions feature that you can use to separate the AD DS administrator and Exchange administrator roles. The AD DS administrator can only create or delete the objects. Exchange administrators don't have permission to create or delete objects in AD DS, but they can create mailboxes for these objects. This lesson describes what you need to consider when you plan and implement split permissions for your organization.

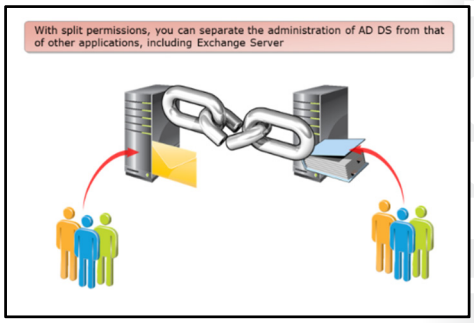
Lesson Objectives

After completing this lesson, you will be able to:

- Describe shared permissions versus split permissions.
- Describe how to configure RBAC split permissions.
- Describe how to configure Active Directory split permissions.
- Discuss guidelines for configuring split permissions.

Shared Permissions vs. Split Permissions

In smaller organizations, the same group of people is often responsible for administering AD DS and Exchange Server deployments. In larger organizations, it is usual to separate the management of AD DS from other applications, including Exchange Server.



The default configuration, in which Exchange administrators can also create and manage AD DS objects such as user objects, is called *shared permissions*. If you separate AD DS administration and Exchange administration, which is called *split permissions*.

Exchange Server 2010 introduced split permissions, which provides a degree of administrative separation between these two facets of the messaging infrastructure. Exchange Server 2013 has two types of split permissions:

- RBAC split permissions. If you implement RBAC split permissions, you prevent Exchange administrators from running cmdlets that create security principals in AD DS. Administrators can still use the AD DS management tools to create security principals. If an Exchange administrator has AD DS permissions to create security principals, they can do so by using the AD DS tools. They can then configure the Exchange attributes by using the Exchange management tools. Also, if you configure RBAC split permissions, you do not modify the underlying RBAC principle that Exchange servers, through the Exchange Trusted Subsystem group, have permissions to create security

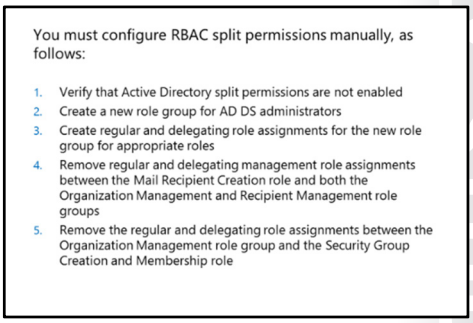
MCT USE ONLY. STUDENT USE PROHIBITED

principals in AD DS. RBAC split permissions does not remove permissions from the Exchange Trusted Subsystem account. RBAC split permissions only removes permission for Exchange administrators to run cmdlets.

- Active Directory split permissions. Contrary to RBAC split permissions, if you implement Active Directory split permissions, the Exchange servers no longer have permission to create AD DS security principals. Permissions to create security principals in the AD DS domain partition are removed completely from any Exchange Server administrator, service, or server. No option is provided in RBAC to create security principals. Administrators can create security principals in AD DS only by using AD DS management tools.

How to configure RBAC Split Permissions

By default, administrators who are assigned to either the Mail Recipient Creation role or the Security Group Creation and Membership role can create security principals in AD DS. In Exchange Server 2013, the Organization Management role group is assigned both of these role assignments, while the Recipient Management role group is assigned the Mail Recipient Creation Role role assignment.



When you configure RBAC split permissions, you remove these management role assignments from the default management role groups. This removal means that the members of the management role groups no longer have permission to run the cmdlets that are used to create security principals, as this blocks them from creating these objects by using any of the Exchange Server 2013 management tools. If you enable RBAC split permissions, Exchange Server administrators cannot use the following cmdlets:

- **New-Mailbox**
- **New-MailContact**
- **New-MailUser**
- **New-RemoteMailbox**
- **Remove-Mailbox**
- **Remove-MailContact**
- **Remove-MailUser**
- **Remove-RemoteMailbox**

If you configure RBAC split permissions, you do not prevent administrators from using the AD DS management tools to create security principals. If an Exchange Server administrator has AD DS permissions to create security principals, they can do so by using the AD DS tools. They can then configure the Exchange Server attributes by using the Exchange Server management tools.

Additionally, configuring RBAC split permissions does not modify the underlying RBAC principle that Exchange servers through the Exchange Trusted Subsystem group have permissions to create security principals in AD DS. RBAC split permissions doesn't remove permissions from the Exchange Trusted Subsystem account— it only removes permission to run cmdlets from Exchange Server administrators.

To configure RBAC split permissions, you must do the following:

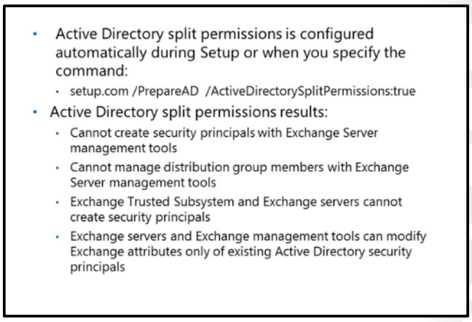
MCT USE ONLY. STUDENT USE PROHIBITED

1. Disable Active Directory split permissions if it is enabled. To do this, run Exchange Server Setup with setup.com with the **/PrepareAD** and **/ActiveDirectorySplitPermissions** parameters set to **false**. If Active Directory split permissions are not enabled and your organization is using the shared permissions model, you can skip this step.
2. Create a new role group that will contain the administrators that can create security principals in AD DS. This is an optional step, but it enables you to configure a special group of Exchange Server administrators that can still use the Exchange Server management tools to create security principals.
3. Create regular and delegating role assignments between the Mail Recipient Creation role and the new role group. This step is optional, and it applies only if you created the special role group mentioned in the previous step.
4. Create regular and delegating role assignments between the new role group and the Security Group Creation and Membership role. This step is optional.
5. Remove the regular and delegating management role assignments between the Mail Recipient Creation role and both the Organization Management and Recipient Management role groups.
6. Remove the regular and delegating role assignments between Organization Management role group and the Security Group Creation and Membership role.

After you configure RBAC split permissions, only members of the new role group that you create can create security principals, such as mailboxes. The new role group can only create the objects. It cannot configure the Exchange Server attributes on the new objects. An AD DS administrator who is a member of the new group needs to create the object, and then an Exchange Server administrator needs to configure the Exchange Server attributes on the object. If you want the new role group also to be able to manage the Exchange Server attributes on the new object, you need to assign the Mail Recipients role to the new role group.

How to Configure Active Directory Split Permissions

Active Directory split permissions differ from RBAC split permissions in that, when you implement Active Directory split permissions, the Exchange servers no longer have permission to create AD DS security principals. This result occurs because these permissions, which are normally granted to the Exchange Windows Permissions group, are removed. Because the Exchange Trusted Subsystem group, which contains all of the servers running Exchange Server 2007, Exchange Server 2010, or Exchange Server 2013, is the only member of the Exchange Windows Permissions group, these permissions are removed from the Exchange servers.



Enabling Active Directory split permissions means that:

- You can no longer create or delete mailboxes, mail-enabled users, distribution groups, or other security principals from the Exchange Server management tools.
- You cannot add distribution group members to or remove distribution group members from the Exchange Server management tools.
- The Exchange Trusted Subsystem and Exchange servers no longer have permissions to create security principals.

MCT USE ONLY. STUDENT USE PROHIBITED

- Exchange servers and the Exchange Server management tools can modify only the Exchange Server attributes of existing security principals in AD DS.

You can enable Active Directory split permissions only when you run the Exchange Server 2013 setup program. When you run the GUI version of setup during the initial deployment of Exchange Server 2013, you can choose to enable split permissions. You can also use the command line setup program with the **/PrepareAD** option and the **/ActiveDirectorySplitPermissions** parameters set to **true** when you first install Exchange Server 2013, or you can run this command after installing Exchange Server to change an existing deployment to use Active Directory split permissions.

When you run setup to implement Active Directory split permissions, the setup program makes the following changes to the AD DS and Exchange Server deployments:

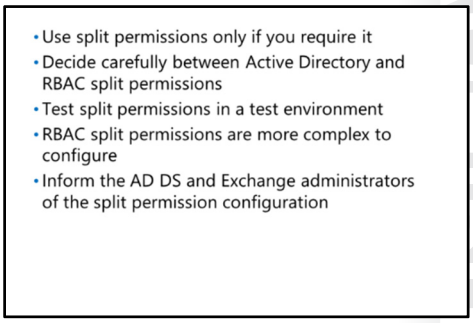
- It creates a new OU called Microsoft Exchange Protected Groups.
- It creates the Exchange Windows Permissions security group in the Microsoft Exchange Protected Groups OU.
- It does not add the Exchange Trusted Subsystem security group to the Exchange Windows Permissions security group.
- It does not create non-delegating management role assignments to management roles with the following management role types:
- MailRecipientCreation.
- SecurityGroupCreationandMembership.
- It does not add access control entries (ACEs) that would have been assigned to the Exchange Windows Permissions security group to the Active Directory domain object.

To disable Active Directory split permissions, you can rerun Exchange setup with the **/PrepareAD** and the **/ActiveDirectorySplitPermissions** parameters, setting the **ActiveDirectorySplitPermissions** parameter to **false**. Additionally, you need to run the following cmdlets to gain back full permissions:

```
New-ManagementRoleAssignment "Mail Recipient Creation_Organization Management" -Role
"Mail Recipient Creation" -SecurityGroup "Organization Management"
New-ManagementRoleAssignment "Security Group Creation and Membership_Org Management"
-Role "Security Group Creation and Membership" -SecurityGroup "Organization
Management"
New-ManagementRoleAssignment "Mail Recipient Creation_Recipient Management" -Role
"Mail Recipient Creation" -SecurityGroup "Recipient Management"
```

Guidelines for Configuring Split Permissions

Configuring split permissions significantly affects your Exchange organization, so you should carefully consider the following guidelines first:



- Use split permissions only if you require it. Don't enable it just because it is there. Make sure that you understand split permissions in Exchange Server and use it only if your security department requires it. In a split permission environment, you face more complex troubleshooting issues if a permission does not work in RBAC, for example.
- Decide carefully between Active Directory and

MCT USE ONLY. STUDENT USE PROHIBITED

RBAC split permissions. Both options in Exchange Server are very similar, but there are tradeoffs. Decide beforehand what split permissions option you want to use.

- Test split permissions in a test environment. Before you configure split permissions, make sure that you test it in your test environment to understand precisely what happens to your Exchange organization after you enable it.
- RBAC split permissions are more complex to configure. Active Directory split permissions are quite simple to configure—just run Setup with the appropriate command line parameters. RBAC split permissions are more complex to configure because you need to run a series of cmdlets. Therefore, if you decide to use RBAC split permissions, make sure that you configure them correctly and plan which cmdlets you need to run.
- Inform the AD DS and Exchange Server administrators about the split permissions configuration. Make sure that they all know about the configuration so they consider it if they are troubleshooting RBAC issues.

MCT USE ONLY. STUDENT USE PROHIBITED

Lesson 3
Planning and Implementing Audit Logging

If you work in an Exchange organization that has only one Exchange Server administrator, you know exactly what was configured and who configured it. In very large organizations, the fact that many administrators can change the Exchange configuration can cause many problems. To prevent these problems, Exchange Server 2013 includes a logging functionality that can provide you with information about the administrative tasks performed on the Exchange servers, in addition to logging information about any operation that is performed on the mailboxes.

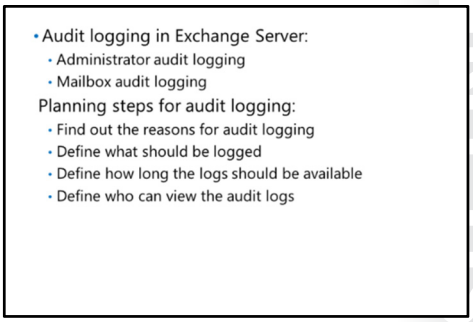
Lesson Objectives

After completing this lesson, you will be able to:

- Plan for audit logging.
- Describe administrator audit logging.
- Describe mailbox audit logging.
- Configure mailbox audit logging.
- Provide options for viewing audit Sign information.

Planning for Audit Logging

Planning for audit logging is a crucial component of any organization's security plan. It is required so that any configuration that is done to the system is logged appropriately, including information about what was changed and who changed it. If there is a problem with the system, you can use the logs to find out what was changed.



Exchange Server 2013 provides the following audit logging features:

- Administrator audit logging. Includes any changes to the Exchange Server configuration.
- Mailbox audit logging. Includes any changes to a user's mailbox.

When you plan for audit logging, consider the following key areas:

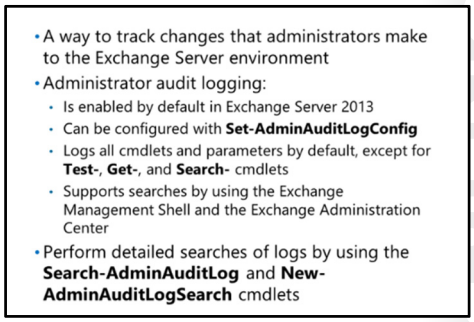
- Find out the reasons for audit logging. The first thing that you should do to start planning for audit logging is to identify why you want to implement it. One reason might be that you work in a large environment where the administrators do not necessarily talk to each other often, so you want a central place to log any change to Exchange Server. Another reason might be that you often get complaints about deletions from user mailboxes, so you need to investigate who makes changes to what mailbox. Deletions are also a concern in a regulated environment, so any attempt to delete data should be flagged.
- Define what should be logged. After you understand the reasons, plan exactly what should be logged. Find the best balance between logging everything, which consumes storage, and logging nothing. A good approach is to write down what areas you need, such as mailbox delegate access, and then configure logging for these areas.

MCT USE ONLY. STUDENT USE PROHIBITED

- Define how long the logs should be available. By default, the logs are available in Exchange Server for 90 days, so consider whether that is sufficient.
- Define who can view audit logs. Make sure that a clearly defined set of people can access the audit logs.

What Is Administrator Audit Logging?

In Exchange Server 2013, administrator audit logging captures data about changes that users and administrators make to the Exchange configuration. By default, administrator audit logging captures information about all changes made to the Exchange Server deployment.



Exchange Server 2013 administrator audit logging logs all Exchange Management Shell cmdlets that make changes to the Exchange Server environment. Because all tasks performed in the Exchange Administration Center are translated to Exchange Management Shell cmdlets, all changes are logged no matter which tool is used to perform the task.

Each time a cmdlet execution is logged, Exchange Server creates an audit log entry. Exchange Server 2013 stores audit logs in a hidden, dedicated system mailbox that you can access only by using the Exchange Administration Center Auditing Reports page, or by using the **Search-AdminAuditLog** or **New-AdminAuditLogSearch** cmdlets. The logs are not accessible from Outlook or Outlook Web App. In addition, no one can delete audit log entries. You cannot modify this dedicated mailbox.

Audit logging shows what actions were taken to modify objects in an Exchange Server organization, rather than what objects were viewed. A cmdlet is audited if it is on the cmdlet auditing list and one or more parameters on that cmdlet are on the parameter-auditing list. By default, the **Test-**, **Get-**, and **Search-** cmdlets are not logged because they are usually not critical to security and they cannot directly change anything on Exchange Server objects. All other cmdlets are logged.

Note: Administrator audit logging logs changes only in Exchange Server. If administrators use tools that directly write to AD DS, such as the Active Directory Users and Computers console, these changes are not logged in the administrator audit logs.

You can configure administrator audit logging in the Exchange Management Shell by using the **Set-AdminAuditLogConfig** cmdlet. You can use several parameters in this cmdlet to configure audit logging. Some of the most important parameters for this cmdlet are the following:

- **AdminAuditLogEnabled**. When set to **False**, logging is not enabled. By default, logging is enabled in Exchange Server 2013.
- **TestCmdletLoggingEnabled**. Enables **Test-** cmdlet logging.
- **AdminAuditLogCmdlets**. Specifies which cmdlets are logged when administrator audit logging is enabled. By default, all cmdlets are logged, as indicated by the wildcard (*) character.
- **AdminAuditLogParameters**. Specifies whether cmdlet parameters are logged. By default, this parameter is set to log all cmdlet parameters, as indicated by the wildcard (*) character.

MCT USE ONLY. STUDENT USE PROHIBITED

- **AdminAuditLogAgeLimit**. Specifies how long each log entry should be kept before it is deleted. The default age limit is 90 days.

If you want to see how administrator audit logging is currently configured, run the **Get-AdminAuditLogConfig** cmdlet.

In the Exchange Administration Center, you can view only administrator audit logging reports. If you want to search the logs by specifying your own search parameters, you must use the Exchange Management Shell.

For example, suppose you want to search **Set-Mailbox** usage between 01/30/2013 and 01/31/2013 and send the search results to Andreas@adatum.com. Run the following cmdlet:

```
New-AdminAuditLogSearch -Cmdlets Set-Mailbox -StartDate 01/30/2013 -EndDate
01/31/2013 -StatusMailRecipients Andreas@adatum.com -Name "Mailbox changes report"
```

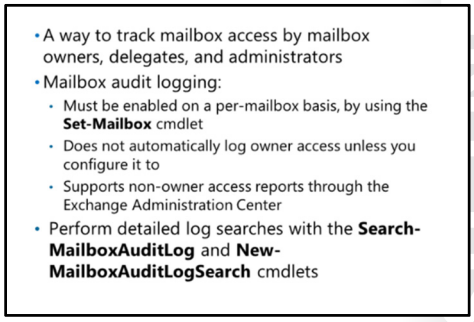
After you run the **New-AdminAuditLogSearch** cmdlet, Exchange Server may take up to 15 minutes to deliver the report to the specified recipient. The search output is limited to 10 MB, so if you have many changes in the system, consider limiting the search.

You can use the same parameters with the **Search-AdminAuditLog** cmdlet, except for the **StatusMailRecipients** parameter, which specifies to send a report by email. The **Search-AdminAuditLog** cmdlet provides a report inside the Exchange Management Shell window.

What is Mailbox Audit Logging?

Mailbox audit logging is a way to log mailbox access by mailbox owners, administrators, and delegates, including administrators who have full mailbox access permissions. Mailboxes are considered accessed by an administrator only in the following scenarios:

- For discovery searches.
- When Mailbox exports are specified through the **New-MailboxExportRequest** cmdlet.
- For mailbox access by the Exchange Server Messaging Application Programming Interface (MAPI) Editor.



When you enable audit logging for a mailbox, you can specify which user actions should be logged. You can also specify whether to log actions by the mailbox owner, delegate, and administrator. Audit log entries also include important information, such as the client IP address, host name, and process or client that was used to access the mailbox. For items that are moved, the entry includes the name of the destination folder.

Mailbox audit logs are generated for each mailbox that has mailbox audit logging enabled. Log entries are stored in the Audits subfolder of the Recoverable Items folder of the audited mailbox. If you move a mailbox to another Mailbox server, the mailbox audit logs for that mailbox also move, because they are located in the mailbox. By default, mailbox audit log entries are retained in the mailbox for 90 days.

Planning for Mailbox Audit Logging

Unlike administrator audit logging, mailbox audit logging is not enabled by default, so you must activate it manually. In addition, mailbox audit logging is activated on a per-mailbox basis and not as a general

MCT USE ONLY. STUDENT USE PROHIBITED

option. If you enable mailbox audit logging for a mailbox, access to the mailbox and certain administrator and delegate actions are logged by default.

To log actions by the mailbox owner, specify which owner actions you want to audit. However, for mailboxes such as the Discovery Search Mailbox, which may contain more sensitive information, consider enabling mailbox AUDIT logging for mailbox owner actions such as message deletion. We recommend that you enable auditing only of the specific owner actions that are necessary to meet business or security requirements.

To enable mailbox auditing on a specific mailbox, use the Exchange Management Shell. The following example enables mailbox auditing on Tony Smith's mailbox:

```
Set-Mailbox -Identity "Tony Smith" -AuditEnabled $true
```

To disable mailbox auditing, change the **$true** parameter to **$false**.

To search the mailbox audit log, you can use either the Exchange Administration Center or the Exchange Management Shell. In the Exchange Administration Center, you can generate reports of who accessed a mailbox other than the owners, which is the most common report for this type of auditing. However, in this report, you can set only a date range as the filter. If you want to specify all available options, use the Exchange Management Shell to perform your search.

The following example searches for users who accessed Tony's mailbox during 2013, limiting results to 2,000:

```
Search-MailboxAuditLog -Identity Tony -LogonTypes Admin,Delegate -StartDate 1/1/2013
-EndDate 12/31/2013 -ResultSize 2000
```

The results return to the Exchange Management Shell window.

The following example searches Terri's and Jan's mailboxes and sends the results to a specific mailbox:

```
New-MailboxAuditLogSearch –Name "Admin and Delegate Access" -Mailboxes "Terri
Chudzik"," Jan Dryml " -LogonTypes Admin,Delegate -StartDate 1/1/2013 -EndDate
12/31/2013 -StatusMailRecipients "auditors@adatum.com"
```

This cmdlet locates access attempts by administrators and delegates during 2013. Results are sent to the email alias auditors@adatum.com.

Demonstration: Configuring Mailbox Audit Logging

In this demonstration, you will learn how to configure mailbox audit logging and how to search audit logs from both the Exchange Administration Center and the Exchange Management Shell.

Demonstration Steps

1. On LON-MBX1, in **Exchange Management Shell**, run the following cmdlet:

```
Add-ADPermission "Anil Elson" –ExtendedRights “Send As”
-user “Allie Bellew”
```

2. On LON-MBX1, in **Exchange Management Shell**, run the following cmdlet:

```
Set-Mailbox -Identity "Anil Elson" -AuditDelegate SendAs,SendOnBehalf -AuditEnabled
$true
```

3. Sign in to Outlook Web App as **Adatum\Allie** with the password **Pa$$w0rd**, and then create an email message that has the following properties:

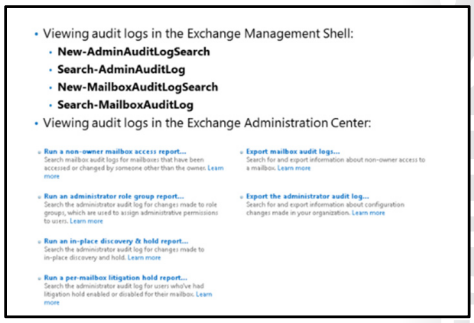
MCT USE ONLY. STUDENT USE PROHIBITED

- From: **Anil@adatum.com**
- To: **Administrator**
- Subject: **Testing Send As logging**

4. Switch to the LON-CAS1 virtual machine, open the Exchange Admin Center, and then sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.
5. In **compliance management**, on the **auditing** tab, run a non-owner mailbox access report. To run the report, in **Search for access by**, click **All non-owners**.
6. In the search results, click **Anil Elson**, and then view the report that shows that Allie Bellew accessed Anil's mailbox.

Options for Viewing Audit Log Information

The audit logs are hidden from the normal users and can be displayed only by using either Exchange Management Shell or Exchange Administration Center. Mailbox audit logs are stored in the Audits subfolder of the Recoverable Items folder of the audited mailbox, and administrator audit logs are stored in the AdminAuditLogs folder in the Recoverable Items folder in the system mailbox.



In Exchange Management Shell, you can use the following cmdlets to search administrator and mailbox audit logs.

Cmdlet	Description
New-AdminAuditLogSearch	Searches the administrator audit log, and then sends the results to one or more mailboxes that you specify.
Search-AdminAuditLog	Searches the administrator audit log.
New-MailboxAuditLogSearch	Searches mailbox audit logs, and sends the search results via email to specified recipients.
Search-MailboxAuditLog	Searches mailbox audit log entries matching the specified search terms.

MCT USE ONLY. STUDENT USE PROHIBITED

The Exchange Administration Center provides you with the following pre-defined reports.

Report	Description
Non-owner mailbox access report	Searches for all non-owner mailbox access on one or all mailboxes. Shows the results in the Exchange Administration Center.
Export mailbox audit logs	Searches for and exports all non-owner access of a mailbox, and sends the report via email.
Administrator role group report	Searches for all changes to management role groups.
Export administrator audit log	Searches for and exports information about any changes made to the Exchange configuration, and sends the report via email.
In-place discovery & hold report	Searches the administrator audit log for changes made to in-place discovery and hold. Shows the results in the Exchange Administration Center.
Per-mailbox litigation hold report	Searches the administrator audit log for one or all users who have had litigation holds enabled or disabled on their mailbox.

Note: An exported mailbox or administrator audit log report in the Exchange Administration Center consists of an .XML file that contains the report for further analysis attached to a message. The .XML generated is limited to 10 MB, so if the output is more, you should change the scope of the search.

MCT USE ONLY. STUDENT USE PROHIBITED

Lab: Designing and Implementing Administrative Security and Auditing

Scenario

A. Datum wants to increase security on their AD DS and Exchange Server infrastructure, and they are also working on separating administrative roles. Your manager assigned you to design a solution and then implement it.

Objectives

After completing this lab, you will be able to:

- Design a solution to implement RBAC and audit logging.
- Implement RBAC by using built-in or custom management role groups.
- Implement mailbox auditing.
- Use administrative auditing for troubleshooting.
- Implement and test Active Directory split permissions.

Estimated Time: 60 minutes

Lab Setup

Virtual machines	**20342B-LON-DC1** **20342B-LON-CAS1** **20342B-LON-MBX1**
User Name	Adatum\Administrator
Password	Pa$$w0rd

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

1. On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In the Microsoft Hyper-V® Manager, click **20342B-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Sign in using the following credentials:
   a. User name: **Adatum\Administrator**
   b. Password: **Pa$$w0rd**
5. Repeat steps 2 to 4 for **20342B-LON-CAS1** and **20342B-LON-MBX1**.

MCT USE ONLY. STUDENT USE PROHIBITED

Exercise 1: Designing a Solution

Scenario

Your manager asks you to design a security model that fulfills following requirements:

- The Exchange Organization Administrators should be able to administer AD DS objects only by using Exchange Server management tools.
- A group of delegated Exchange Server mailbox administrators should exist who have permission to manage only Exchange recipient objects.
- A group of delegated Exchange administrators should exist who can manage Exchange servers, but not mailboxes.
- A group of site administrators must be able to install provisioned servers, but the administrators should not be able to configure Exchange servers.
- The Mailbox Managers group should be able only to enable or disable mailboxes.
- All access to mailboxes from non-owners must be audited.
- The audit department must be able to access all administrative and mailbox audit logs, but it should not be able to change the configuration of audit logs in Exchange Server.
- All administrative tasks must be audited.

The main tasks for this exercise are as follows:

1. Read and analyze scenario requirements
2. Design a solution
3. Discuss your solution with the class

▶ Task 1: Read and analyze scenario requirements

1. Read the exercise scenario and analyze the requirements from a security perspective. Identify the permissions or audit requirements needed to satisfy the requirements.

▶ Task 2: Design a solution

Propose a solution for administrative security and auditing. Use the following questions as a guideline:

1. What built-in management role groups can you use to address all requirements in the exercise scenario?
2. What additional custom role groups do you need to create to satisfy all requirements?
3. What mailbox or administrative audit logging is required to satisfy all requirements?

▶ Task 3: Discuss your solution with the class

1. Present your proposed solution to the class. Discuss alternative solutions with other students and with the instructor.

Exercise 2: Implementing Role-Based Access Control

Scenario

As a part of a security enhancement, you need to implement the custom role groups that are identified in your security plan. Brad Sutton is a member of the Exchange Mailbox Administrators, so he requires permission to create and manage recipients. Erwin Zischka is a member of the Mailbox Managers group and needs permission only to enable and disable mailboxes.

MCT USE ONLY. STUDENT USE PROHIBITED

The main tasks for this exercise are as follows:

1. Use the Recipient Management role group to delegate permissions to create and manage recipients

2. Create a custom role group to allow only enabling and disabling of mailboxes

3. Verify that the administrators have permission to perform their tasks

▶ Task 1: Use the Recipient Management role group to delegate permissions to create and manage recipients

1. On LON-MBX1, sign in to the Exchange Admin Center as **Adatum\Administrator** with the password **Pa$$w0rd**.
2. In **permissions feature**, on the **admin roles** tab, add **Brad Sutton** to the **Recipient Management** role group.

▶ Task 2: Create a custom role group to allow only enabling and disabling of mailboxes

1. In the Exchange Admin Center, in admin roles, create a new role group that has the following settings:
   - Name: **Mailbox Managers**
   - Description: **Enable and disable mailboxes**
   - Roles: **Mail Recipients**
   - Members: **Erwin Zischka**
2. Close Internet Explorer.

▶ Task 3: Verify that the administrators have permission to perform their tasks

1. Switch to LON-CAS1, open Internet Explorer, and then connect to **https://LON-CAS1.adatum.com/ecp**. Sign in as **Adatum\Brad** with the password **Pa$$w0rd**.
2. In **recipients feature**, in **mailboxes**, create a user mailbox with a new user that has the following properties:
   - Alias: **Test**
   - First name: **Test**
   - Last name: **Test**
   - User logon name: **Test**
   - New password: **Pa$$w0rd**
   - Confirm password: **Pa$$w0rd**

   This confirms that Brad can create new mailboxes.
3. Verify that Brad cannot modify any server settings.
4. Close Internet Explorer, re-open Internet Explorer, and then connect to **https://LON-CAS1.adatum.com/ecp**. Sign in as **Adatum\Erwin** with the password **Pa$$w0rd**.
5. Disable the mailbox of **Amr Zaki**.
6. Verify that, on the feature pane, servers are not available to Erwin, because of his restricted permissions.
7. Close Internet Explorer.

MCT USE ONLY. STUDENT USE PROHIBITED

Exercise 3: Implementing Mailbox Audit Logging

Scenario

As a part of a security enhancement initiative, you have the following requirements for auditing:

- All access to mailboxes from non-owners must be audited.
- Make sure that audit logging works as expected.

The main tasks for this exercise are as follows:

1. Enable mailbox audit logging of non-owners on all mailboxes
2. Delete items from a mailbox by using a different account
3. Verify that the activity is logged

▶ Task 1: Enable mailbox audit logging of non-owners on all mailboxes

1. On the LON-MBX1 virtual machine, open the Exchange Management Shell, and then run the following cmdlet:

```
Add-MailboxPermission Tony
-AccessRights fullaccess
–user Administrator
```

2. In the Exchange Management Shell, at the **PS** prompt, type the following:

```
Get-Mailbox | Set-Mailbox -AuditEnabled $true
```

3. Minimize the Exchange Management Shell.

▶ Task 2: Delete items from a mailbox by using a different account

1. Open Internet Explorer, type **https://LON-CAS1.adatum.com/owa**, and then sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.
2. In Outlook Web App, click **Open another mailbox**, and then open **Tony@adatum.com**.
3. Delete one message from Tony's Inbox.
4. Empty the **Deleted Items** folder, and then close Internet Explorer.

▶ Task 3: Verify that the activity is logged

1. Open Internet Explorer, type **https://LON-CAS1.adatum.com/ecp**, and then sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.
2. In the **Run a non-owner mailbox access report**, search for access by **All non-owners**.
3. Click **Tony Smith**, and then notice in the report that the Administrator performed a soft-delete operation in the mailbox.
4. Close Internet Explorer.

Exercise 4: Using Administrative Audit Logging to Troubleshoot

Scenario

You get a call from a department leader that one of their user accounts, for April Reagan, has been disabled. From speaking to the other administrators, you know that none of them has recently modified April's mailbox. You need to track down what happened to April's mailbox and who modified it.

The main tasks for this exercise are as follows:

MCT USE ONLY. STUDENT USE PROHIBITED

1. Change the Exchange configuration

2. Troubleshoot what happened by using the administrative audit log

▶ Task 1: Change the Exchange configuration

1. On the LON-CAS1 virtual machine, open the Exchange Management Shell, and then run the following command:

2. E:\LabFiles\Mod08\Mod08Ex4.bat

▶ Task 2: Troubleshoot what happened by using the administrative audit log

1. On the LON-MBX1 virtual machine, maximize the Exchange Management Shell, and then run the following cmdlet:

```
Search-AdminAuditLog -ObjectIds “April Reagan”
Get-Mailbox “Ed Meadows”
```

2. You find out that the account for Ed Meadows does not exist anymore. The mailbox must have been removed.

3. At the **PS** prompt, run the following:

```
Search-AdminAuditLog –StartDate yesterday –EndDate tomorrow | Sort-object rundate
```

   In this command, the dates are written in the mm/dd/yyyy format. So May 7, 2013 would be written as 05/07/2013.

4. Can you identify how Ed's mailbox was deleted and who enabled the permission change that made it possible for Ed to disable April’s mailbox?

Exercise 5: Implementing and Testing Active Directory Split Permissions

Scenario

A. Datum strictly separates the administration of AD DS and Exchange Server. The AD DS administrators complain that Exchange Server administrators accidentally created user accounts that do not follow A. Datum’s naming conventions. As the previous exercise shows, someone can create a user object and then can assign themselves permissions to hide tasks that they perform.

For that reason, IT management decides to physically separate AD DS objects created from Exchange Server. To satisfy this requirement, you need to configure a separation between the AD DS administrators and the Exchange Server administrators.

The main tasks for this exercise are as follows:

1. Configure Active Directory split permissions for Exchange Server

2. Verify that the Exchange Server administrators cannot change objects directly in AD DS

3. To prepare for the next module

▶ Task 1: Configure Active Directory split permissions for Exchange Server

1. On the LON-MBX1 Virtual Machine Connection, attach the Exchange ISO file **D:\Program Files\Microsoft Learning\20342\Drives\ExchangeServer2013CU1.iso**.

2. Open the Windows PowerShell® command-line interface, change the path to **D:** and then run **Setup /PrepareAD /ActiveDirectorySplitPermissions:true /IAcceptExchangeServerLicenseTerms**.

3. Wait until the process finishes, and then close Windows PowerShell.

MCT USE ONLY. STUDENT USE PROHIBITED

▶ Task 2: Verify that the Exchange Server administrators cannot change objects directly in AD DS

1. On LON-MBX1, open Internet Explorer, and then connect to **https://LON-CAS1.adatum.com/ecp**. Sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.
2. In **recipients**, on the **mailboxes** tab, try to create a mailbox with **New user**. Note that all of the fields for creating a user, such as **First name**, **Last name**, and **User logon name**, are grayed out. Therefore, even though this administrator is a Domain Admin, you cannot create a user object in Exchange Server anymore through their account.
3. Create a mailbox for the existing user account **Ales Ruzicka**.
4. On the **groups** tab, try to add Ales Ruzicka to the **IT** group. When you save the group, an error appears that says, "You don't have sufficient permissions." This error appears because you cannot manage groups any longer from Exchange Server.

▶ Task 3: To prepare for the next module

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

1. On the host computer, start Microsoft Hyper-V® Manager.
2. In the **Virtual Machines** list, right-click **20342B-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for 20342B-LON-CAS1 and 20342B-LON-MBX1.
5. In Hyper-V Manager, click **20342B-LON-DC1**, and in the Actions pane, click **Start**.
6. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
7. Sign in using the following credentials:
   a. User name: **Adatum\Administrator**
   b. Password: **Pa$$w0rd**
8. Repeat steps 5 to 7 for **20342B-LON-MBX1, and 20342B-LON-CAS1**.

Results: In this exercise, you configured Active Directory split permissions.

Question: During the lab, what was the reason to create a custom role group for mailbox managers?

Question: What is the difference between Active Directory split permissions and RBAC split permissions? When should you implement each?

MCT USE ONLY. STUDENT USE PROHIBITED

Module Review and Takeaways

Best Practice

- Use built-in management role groups wherever possible.
- To prevent uncontrolled growth of mailbox size, enable mailbox audit logging only if you need to.
- Never use split permissions unless you absolutely need to.

Review Question(s)

Question: How will you configure role assignment policies in your own organization?

MCT USE ONLY. STUDENT USE PROHIBITED

MCT USE ONLY. STUDENT USE PROHIBITED

Module 9

Managing Exchange Server 2013 with Exchange Management Shell

Contents:

Module Overview 9-1

Lesson 1: Overview of Windows PowerShell 3.0 9-2

Lesson 2: Managing Exchange Server Recipients by Using the Exchange Management Shell 9-7

Lesson 3: Using Windows PowerShell to Manage Exchange Server 9-17

Lab: Managing Microsoft Exchange Server 2013 by Using Exchange Management Shell 9-23

Module Review and Takeaways 9-27

Module Overview

The Windows PowerShell® command-line interface is a core feature of the Windows Server® operating system. Windows PowerShell enables command-line management and configuration of the operating system and of Microsoft® Exchange Server 2013. It is a standardized, task-based command-line shell and scripting language that offers you flexibility and choice in how you manage computers running Windows Server. The Exchange Management Shell enables you to access Exchange management features from within Windows PowerShell. Understanding the basics of Windows PowerShell is important when learning to use the Exchange Management Shell effectively.

Windows PowerShell 3.0 offers more functionality and features than older versions. In addition to managing Exchange Server 2013, you can now use Windows PowerShell to manage all of the Windows Server roles and features. This capability enables you to automate configuration tasks quickly by using a single tool, instead of needing to use multiple tools, such as batch scripts, Microsoft Visual Basic® Script Edition scripts (VBScripts), and manual configuration steps.

In this module, you will learn key Windows PowerShell concepts and how to use Exchange Management Shell.

Objectives

After completing this module, you will be able to:

- Describe the Windows PowerShell Integrated Scripting Environment (ISE).
- Use Windows PowerShell remotely.
- Manage Exchange Server configuration and recipients by using the Exchange Management Shell.

MCT USE ONLY. STUDENT USE PROHIBITED

Lesson 1
Overview of Windows PowerShell 3.0

As an Exchange Server administrator, you can use Windows PowerShell to configure Exchange Server features and recipients, as well as other software like Windows Server and Microsoft System Center 2012. Although you can use the Exchange Administration Console, a graphical user interface (GUI), for administration, if you use Windows PowerShell you can create automation scripts to administer and access configuration options that are not available in the GUI. Some tasks that you can perform in Windows PowerShell, such as listing the contents of a directory, may already be familiar to you. To use Windows PowerShell effectively, you must have a basic understanding of how this command line environment works and how to use it.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe what is new in Windows PowerShell 3.0.
- Describe the features of the Windows PowerShell ISE.
- Describe the functionality of Windows PowerShell modules.
- Add the Exchange Management Shell module into a Windows PowerShell session.
- Explain how to use Windows PowerShell remotely.
- Analyze the benefits of using Windows PowerShell to manage Exchange Server 2013.

What's New in Windows PowerShell 3.0

Windows PowerShell is a command-line management interface that you can use to configure Windows Server 2012 and other products, such as System Center 2012, Exchange Server 2013, and Microsoft SharePoint® Server 2013. This management interface is an alternative to the GUI-based Exchange Administration Console, and you can use it to do the following:

- More than 260 core cmdlets
- Management of all Windows Server roles and features
- Windows PowerShell Workflow
- Windows PowerShell Web Access
- Scheduled Jobs
- Enhanced Online Help
- ISE IntelliSense
- Robust Session Connectivity

- Create automation scripts.
- Perform batch modifications.
- Access settings that might be unavailable or more difficult to configure in the GUI.

Windows PowerShell 3.0 has new features that facilitate managing larger groups of servers through better scaling, additional functionality, and better management. Windows PowerShell 3.0 includes the following new features:

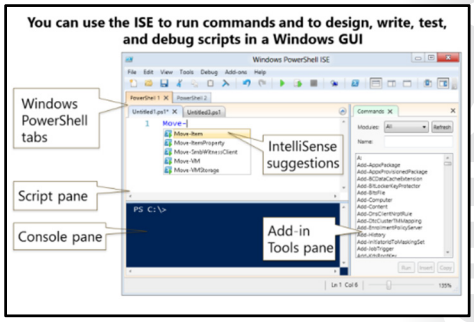
- Windows PowerShell Workflow. You can use this feature to coordinate complex parallel and sequenced commands.
- Windows PowerShell Web Access. You can use this feature to encrypt and authenticate access to Windows PowerShell by using a web browser on any device.
- Scheduled Jobs. You can use this feature to schedule Windows PowerShell commands and scripts to run administrative tasks automatically.

MCT USE ONLY. STUDENT USE PROHIBITED

- Enhanced Online Help. You can download or view online the most recent Help files from Microsoft by using the Update-Help cmdlet. This guarantees you have the most recent information about how to use Windows PowerShell.
- Windows PowerShell ISE IntelliSense(R) computer program. Windows PowerShell ISE provides hints for cmdlets, including parameters that make it easier to use Windows PowerShell than in the past.
- Robust Session Connectivity. You can use these connections to connect to a remote server. If connectivity is lost or if you intentionally disconnect, you can resume the connection at the point at which it was disconnected. Previously, if the connection to a session was lost, all the session data, variables, and command history were also lost.

What Is Windows PowerShell ISE?

Windows PowerShell ISE is a graphical Windows PowerShell development environment that includes debugging capabilities and an interactive console. The Windows PowerShell ISE is installed by default, and it includes the following features:



- *Integrated environment*. Edit, run, and debug scripts and perform interactive shell tasks from one development environment.
- *Syntax coloring*. Keywords, objects, properties, cmdlets, variables, strings, and other tokens appear in different colors, improving readability and reducing errors.
- *Selective invocation*. Select any portion of a Windows PowerShell script, run it, and view the results in the Console pane.
- *The Microsoft IntelliSense® feature*. Provides inline help for cmdlets, parameters, parameter values, and file paths as you type in either the Script or Console panes.
- *Add-on tools*. The ISE supports extending the interface with Windows Presentation Foundation (WPF) controls that display in either a horizontal or vertical pane. You can add as many as 20 tools at a time, each of which displays in a separate tab. The Commands add-on is an example add-on that is installed and enabled by default to provide help for each cmdlet.
- *Multiple sessions*. You can simultaneously use up to 32 independent sessions each on its own Windows PowerShell tab, within the ISE. This way, you can manage multiple servers, each in its own environment, from within one instance of ISE.
- *Script editor*. You can use the script editor to compose, edit, debug, and run functions, scripts, and script cmdlets. The script editor includes tab completion, automatic indenting, line numbers, search-and-replace, and go-to line.
- *Debugging*. You can use the integrated visual script debugger to set breakpoints, to step through the script, to check the call stack, and to hover over variables to inspect their value.
- *Object model*. The ISE comes with a complete object model, which you can use to write Windows PowerShell scripts to manipulate the ISE.
- *Customizability*. The ISE is customizable, from the size and placement of the panes, to text size and background colors.

MCT USE ONLY. STUDENT USE PROHIBITED

The Windows PowerShell ISE Profile

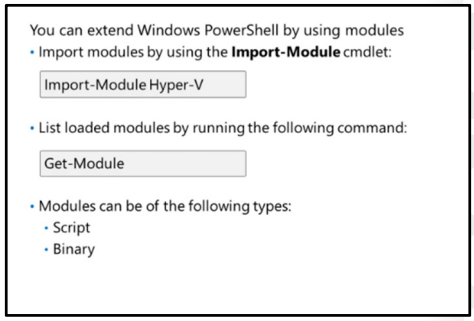
The Windows PowerShell ISE has its own Windows PowerShell profile: Microsoft.Windows PowerShellISE_profile.ps1. You can use this profile to store functions, aliases, variables, and commands that you use in the Windows PowerShell ISE.

Items in the Windows PowerShell AllHosts profiles (CurrentUser\AllHosts and AllUsers\AllHosts) are available in the Windows PowerShell ISE, just as they are in any Windows PowerShell host program. However, items in the Windows PowerShell console profile are not available in the Windows PowerShell ISE.

Instructions for moving and reconfiguring profiles are available in the Windows PowerShell ISE Help and the about_profiles help topic.

Windows PowerShell Modules

Windows PowerShell is extensible. You can add new cmdlets and functions in Windows PowerShell 3.0 in part by using modules. Windows PowerShell uses the Microsoft.PowerShell.Management module, which provides basic functionality. If you install additional roles, additional Windows PowerShell modules may be installed and registered.



Note: In earlier versions of Windows PowerShell, extensibility is provided by using snap-ins. For backward compatibility, Windows PowerShell 3.0 continues to support snap-ins.

For example, assume you install the Microsoft Hyper-V® role and also the Hyper-V module for Windows PowerShell. To manage Hyper-V from Windows PowerShell, you must import the Hyper-V module into the Windows PowerShell session. To import the Hyper-V module, run the following:

```
Import-Module Hyper-V
```

Run the following to list all imported modules:

```
Get-Module
```

You do not always need to import modules manually. For example, the Windows PowerShell module for Exchange Server 2013 is automatically imported if you open the Exchange Management Shell. However, you still can import the Exchange PowerShell module into a Windows PowerShell session to enable the Exchange cmdlets, if the Exchange Server 2013 management tools have been installed. In other cases, Windows PowerShell automatically loads modules if a cmdlet from the module is accessed. However, if you cannot run cmdlets for a specific application, you may need to import the appropriate Windows PowerShell module.

There are two basic module types:

- Binary. A binary module is created by using the Microsoft .NET Framework and is often included with a product to provide Windows PowerShell support. Binary modules often add cmdlets that consist of noun or subject types that are newly created in the Active Directory® Domain Services (AD DS) schema to support the product.

MCT USE ONLY STUDENT USE PROHIBITED

- Script. A script module consists of scripts, modules, and custom settings. These scripts can provide additional functions and variables to automate repetitive or tedious tasks. You may want to create your own module that includes functions or variables specific to your environment, as a way to save time or to manage configurations. The Exchange Management Shell is loaded by using a script module to connect to an Exchange server and to load the available cmdlets.

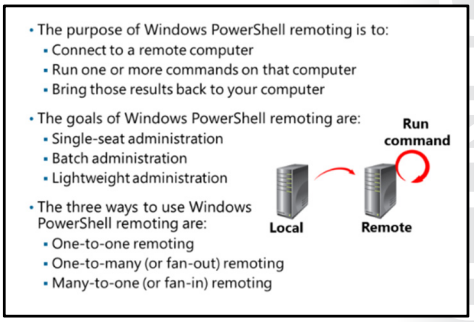
Additional Reading: For more information, see Windows PowerShell Modules at http://go.microsoft.com/fwlink/?LinkId=290688).

Demonstration: Adding Exchange Management Shell to Windows PowerShell ISE

Demonstration Steps

1. Open Windows PowerShell ISE as an administrator, and then review the Script pane and the Console pane.
2. Open **E:\Labfiles\Mod09\Democode\UsingEMSinISE.txt**, and then follow the instructions.

How Can You Use Windows PowerShell Remotely?



Windows PowerShell remoting is used to connect to remote computers, to run commands on those computers, and to direct the results back to your local computer. This enables single-seat administration, which means that you can manage the computers on the network from your workstation, laptop, or tablet, instead of needing to physically visit each computer. A key goal of Windows PowerShell remoting is to enable batch administration, which means that you can run commands on an entire set of remote computers concurrently.

There are three primary ways to use Windows PowerShell remoting:

- *One-to-one*. In this scenario, you connect to a single remote computer and run shell commands on it, exactly as if you had logged into the console and opened a Windows PowerShell window.
- *One-to-many*, or *fan-out*. In this scenario, you issue a command that is executed on one or more remote computers in parallel. You do not work with each remote computer interactively. Instead, your commands are issued and executed in a batch, and the results are returned to your computer for your use.
- *Many-to-one*, or *fan-in*. In this scenario, multiple administrators connect remotely to a single computer. Typically, those administrators have different permissions on the remote computer and might be working in a restricted runspace within the shell. This scenario usually requires custom development of the restricted runspace, a topic which this course does not cover.

Remoting requires that you have both Windows PowerShell and Windows Remote Management (WinRM) on your local computer and on any remote computers that you want to connect to. WinRM is a Microsoft implementation of Web Services for Management, which is a set of protocols that is widely adopted across different operating systems.

MCT USE ONLY STUDENT USE PROHIBITED

As their names imply, Web Services for Management and WinRM use web-based protocols. A benefit of these protocols is that they use a single, definable port. This characteristic makes it easier to pass the commands through firewalls than older protocols that randomly select a port. WinRM communicates by using HTTP. By default, WinRM and Windows PowerShell remoting use TCP port 5985 for incoming connections that are not encrypted and TCP port 5986 for incoming connections that are encrypted.

Applications that use WinRM, such as Windows PowerShell, can also apply their own encryption to the data that is passed to the WinRM service. WinRM supports authentication, and, by default, it uses the Active Directory native Kerberos version 5 protocol in a domain environment. Kerberos authentication does not pass credentials over the network, and it supports mutual authentication to help ensure that incoming connections are coming from authorized computers.

To establish a one-to-one remoting session by using the Windows PowerShell ISE, on the **File** menu, click the **New Remote PowerShell** tab. You also can establish a remote Windows PowerShell session by using the **Enter-PSSession** cmdlet. For example, to open a remote Windows PowerShell session on a computer named LON-MBX1, use the following syntax.

```
Enter-PSSession –ComputerName LON-MBX1
```

To establish a one-to-many remoting session, use the **Invoke-Command** cmdlet. To run the **Get-EventLog** cmdlet against the computers named LON-CAS1 and LON-MBX1, use the following.

```
Invoke-Command -ScriptBlock { Get-EventLog System -Newest 5 } -ComputerName LON-CAS1,
LON-MBX1
```

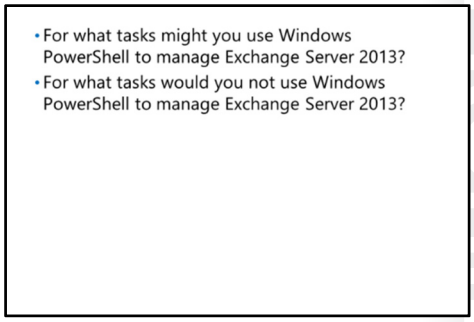
Note: Unlike in earlier versions, Windows Server 2012 enables Windows PowerShell remoting and WinRM by default.

When you load the Exchange Management Shell or the Exchange Management Shell module, a remote Windows PowerShell session is established with an Exchange server in the organization. If you are running the Exchange Management Shell on an Exchange server, you can establish the remote Windows PowerShell session with the local computer itself.

Discussion: Benefits of Using Windows PowerShell to Manage Exchange Server 2013

Discuss the following questions:

- For what tasks might you use Windows PowerShell to manage Exchange Server 2013?
- For what tasks would you not use Windows PowerShell to manage Exchange Server 2013?



MCT USE ONLY. STUDENT USE PROHIBITED

MCT USE ONLY. STUDENT USE PROHIBITED

Lesson 2

Managing Exchange Server Recipients by Using the Exchange Management Shell

You can use the Exchange Management Shell to manage all properties, settings, and objects within an Exchange Server organization. By using Windows PowerShell, you can save time and effort by automating many of your time-consuming or repetitive tasks. Automation also can help improve security and consistency, because it is less prone to human error than manual administration is.

This lesson examines how to use the Exchange Management Shell cmdlets. It also helps you develop the skills that you need to discover, explore, learn, and use other add-in commands, whether they are included with Exchange Server 2013 or with other software products.

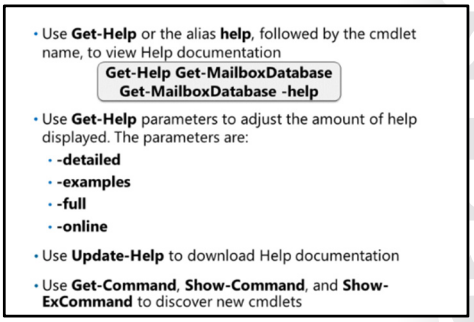
Lesson Objectives

After completing this lesson, students will be able to:

- Describe how to view the cmdlet help documentation.
- Describe how to use pipelines.
- Describe the options for formatting pipeline output.
- Manage Exchange recipients by using the Exchange Management Shell.
- Describe how to use loops and conditional expressions.
- Describe how to create Exchange Management Shell scripts.
- Analyze the scenarios for managing Exchange recipients with the Exchange Management Shell.
- Create an Exchange Management Shell script.

Accessing Help in Windows PowerShell

Whether you are experienced with or new to Windows PowerShell, the cmdlet Help documentation is a rich source of information. To access the Help documentation, use the **Get-Help** cmdlet or its alias, **help**, followed by the cmdlet name. **Get-Help** has parameters to adjust the Help content that is displayed. These parameters are: **Detailed**, **Examples**, **Full**, and **Online**. If you run **Get-Help** without a parameter, the default view is shown. The **Online** parameter opens the online version of the Help file. The following table summarizes the differences between the Help views.



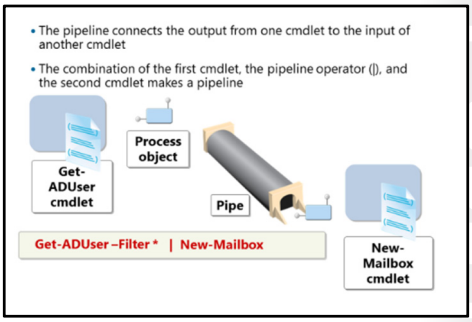
Details shown	Default	Detailed	Full	Examples
Synopsis	Yes	Yes	Yes	Yes
Syntax	Yes	Yes	Yes	Yes
Description	Yes	Yes	Yes	No

Details shown	Default	Detailed	Full	Examples
Inputs	No	No	Yes	No
Outputs	No	No	Yes	No
Errors	No	No	Yes	No
Examples	No	Yes	Yes	Yes
Related links	Yes	No	Yes	No
Remarks	Yes	Yes	No	No

The Exchange Management tools include help documentation, but the Windows role and feature Help documentation is not included by default. To download or update the Windows role and feature Help documentation locally, use the **Update-Help** cmdlet. Additional assistance is available with the **Show-Command** cmdlet. This cmdlet helps less experienced Windows PowerShell users interact with the input and output options of a specified cmdlet by using a GUI.

The **Get-Command** cmdlet returns a list of all locally available cmdlets, functions, and aliases. You can use it to discover new cmdlets by using wildcard searches. For example, to return a list of all cmdlets that include "Ex" in them, run **Get-Command *Ex***. You can also use the **Get-ExCommand** cmdlet to return the available Exchange cmdlets.

Windows PowerShell Pipelines



Windows PowerShell is an object-based environment, in that the input and outputs of the cmdlets are objects that can be manipulated. In some instances, you may want to take the output of one cmdlet and pass it to another cmdlet for additional actions. For example, if you need to create an Exchange Server mailbox for all AD DS accounts in a domain, you can manually list each user by using the **Get-ADUser** cmdlet. Then, by using Windows PowerShell, you can use the **New-Mailbox** cmdlet for each user account.

To simplify this process, you can pass the output data directly from one cmdlet into another cmdlet, which is called *piping*. Piping is performed by putting the pipeline operator (|) between cmdlets. Each cmdlet is executed from the left to the right, with each cmdlet passing its output to the next cmdlet in line. For example, you can get a list of all users in the domain and then pipe the list to the **New-Mailbox** cmdlet by running the following:

```
Get-ADUser –Filter * | New-Mailbox
```

Piping can be used extensively in Windows PowerShell, as with other shells. Windows PowerShell differs from typical shells because the data in the pipeline is an object instead of just simple text. By having an object in the pipeline, you can easily persist all the properties of the returned data. The data in the pipeline is assigned to a special variable named **$_**, which exists only while the pipeline is executing. For

MCT USE ONLY. STUDENT USE PROHIBITED

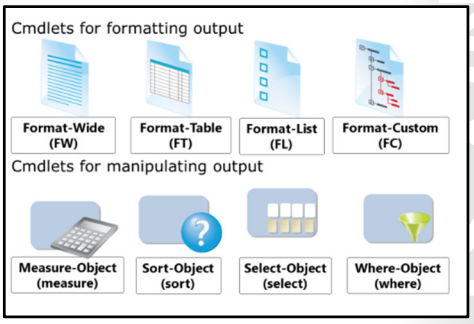
example, if you want to create mailboxes for all enabled accounts, you can use the **Where-Object** cmdlet to return only accounts that are enabled. To do this, run the following:

```
Get-ADUser –Filter * | Where-Object {$_.Enabled –eq $true} | New-Mailbox
```

By piping an object with a list of all the users, you can use the **Where-Object** cmdlet to filter the accounts that are disabled based on the **Enabled** property of the account.

Options for Formatting Windows PowerShell Output

When you work with Exchange Server data, you may need to retrieve lists of users, computers, or groups, and you may need to visualize the data by using a tool such as Microsoft Excel® spreadsheet software. Alternatively, you may need to view only the specific properties on the screen. Windows PowerShell enables both of these scenarios. It has several default cmdlets available to format data, as the following table describes.



Cmdlet	Description
Format-List	This cmdlet outputs data in a list format, with each property on its own line. You can specify the properties that you want displayed by using the **-Property** parameter. You can call this cmdlet by using the alias **FL**. This cmdlet is useful to view a small number of objects that have a large number of properties.
Format-Table	This cmdlet outputs data in a table format, with each property as its own column. You can specify the properties that you want to display by using the **-Property** parameter. You can call this cmdlet by using the alias **FT**. This cmdlet is useful to view a large number of objects that have a small number of properties.
Format-Wide	This cmdlet outputs data in a table format, with only one property for each object. You can specify the property that you want to display by using the **-Property** parameter, and you can specify the number of columns to display the data by using the **–Column** parameter. You can call this cmdlet by using the alias **FW**. This cmdlet is useful to view a large number of objects if you need to see only one property, such as the name, for each object.
Format-Custom	This cmdlet outputs data in a format previously defined by using a PS1XML file. The settings in this file can specify which properties to show, and how to arrange and group them. You can call this cmdlet by using the alias **FC**. This cmdlet is useful to view data that you access frequently if you also want to customize which properties to show.

MCT USE ONLY STUDENT USE PROHIBITED

The following table displays another set of cmdlets that enables complex formatting and reporting.

Cmdlet	Description
Measure-Object	This cmdlet takes the input object from the pipeline or variable and performs calculations on specified properties and on text in strings and files. Calculations include counting objects and determining the average, minimum, maximum, and sum of property values. This cmdlet can also count the number or occurrences of words and characters in a file or string. It is useful if you want to quickly calculate the number of users selected as part of a query or if you are determining how much memory a set of processes is using.
Select-Object	This cmdlet takes the input object from the pipeline or variable and outputs objects that have only the selected properties. It also can select a subset of items in each object by using the **-First**, **-Last**, **-Unique**, and **-Index** parameters, which is valuable if you work with large datasets.
Sort-Object	This cmdlet takes the input object from the pipeline or variable and sorts the data based on the selected properties. This option is helpful if you want to provide a sorted list of data.
Where-Object	This cmdlet takes the input object from the pipeline or variable and applies a filter that is based on a specified query. The queries that are used for filtering are enclosed in braces, and they include a comparison. This option is helpful if you want to select specific types of data.

You can use all of these cmdlets together to customize the output to the screen. You also can use the **Out-File** cmdlet to write the output to a text file or the **Export-Csv** cmdlet to export the data as a comma-separated values (CSV) file. For example, you can export Mailbox statistics for all mailboxes in Database1 by piping the results to the **Export-CSV** cmdlet by running the following:

```
Get-MailboxStatistics –Database Database1 | Export-Csv Stats.Csv
```

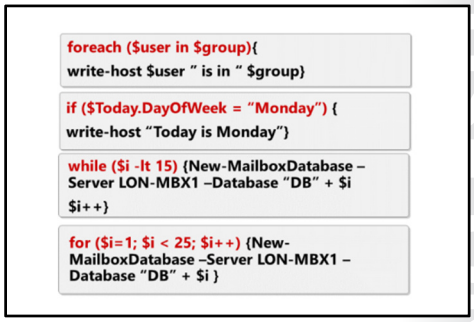
Demonstration: Managing Exchange Recipients by Using the Exchange Management Shell

Demonstration Steps

1. Sign in to LON-CAS1 as the domain administrator.
2. Open the Windows PowerShell ISE as an administrator, and then review the Script pane and the Console pane.
3. Open **E:\Labfiles\Mod09\Democode\ManagingExchange.txt**, and then follow the instructions.

MCT USE ONLY. STUDENT USE PROHIBITED

Using Windows PowerShell Variables and Loops



In advanced Windows PowerShell scripts, you may need to repeat commands a certain number of times, until a specific condition is met, or only if a specific condition is met. You can define these test conditions by using comparison statements.

Boolean Comparisons

Test or comparison statements are used as test conditions for loops and conditional constructs. These statements typically compare either two or more objects or two or more property values, and they result in a value of **True** or **False**. These comparisons are often called *Boolean comparisons,* because they can result only in one of the two Boolean values, **True** or **False**.

Using Boolean comparisons is quite common when designing a Windows PowerShell script. For example, you might compare two computer names to see whether they are equal, or you might compare a performance counter value to a predetermined threshold value to see which of the two is greater. The comparison operators sit between the two items that you want to compare. Recall simple mathematical comparisons that you learned as a child, such as 10 > 4, 5 < 10, and 15 = 15. Windows PowerShell performs comparisons the same way, although it has its own syntax. Some common comparison operators include the following:

- **-eq**. Equal to
- **-ne**. Not equal to
- **-le**. Less than or equal to
- **-ge**. Greater than or equal to
- **-gt**. Greater than
- **-lt**. Less than

Windows PowerShell defines two special variables for comparisons, $True, and $False, which represent the Boolean values True and False, respectively. If a comparison is true, the expression is evaluated as $True, and if the comparison is not true, the expression is evaluated as $False. For example, the comparison 4 is greater than 10 (4 –gt 10) produces $False as its result, but the comparison 10 is equal to 10 (10 –eq 10) produces $True.

Windows PowerShell enables you to execute comparisons on the command line itself. Type the comparison on the command line, and then press Enter to see the result of the comparison. You often use Boolean comparisons to control loops and conditional expressions.

Several Windows PowerShell constructs use Boolean comparisons to control the execution of code in a script. These constructs are **if**, **switch**, **for**, **while**, and **foreach**.

The foreach Statement

The **foreach** statement iterates through an array (collection), item by item, assigning a specifically named variable to the current item of the collection. Then it runs the code block for that element, as follows:

```
foreach (item in collection)
{
Code to complete for each item in the collection.
}
```

MCT USE ONLY. STUDENT USE PROHIBITED

Using the **foreach** statement can simplify batch modifications. Consider, for example, setting a description for all users who are members of a specific group, as shown in the following example:

```
# Get a list of the members of the Exchange Organization Management group
$DAdmins =  Get-ADGroupMember "Organization Management"
# Go through each member and set the Description
foreach ($user in $DAdmins)
{
Set-ADUser $user -Description “In the Exchange Organization Management Group”
}
```

The if Statement

You can use the **if** statement to execute a block of code if the specified criteria are met. The basic functionality of an **if** statement is shown in the following example:

```
if (Boolean comparison)
{
Code to complete if test expression is true
}
```

Another option is to use **else** and **elseif** statements. If you want to execute special code if a condition exists, or if you want to execute other code if a condition does not exist, you can use **else**. If there are additional conditions that you want to test for, use the **elseif** statement. Consider the following example:

```
$Users = Get-ADUser –Filter * -Properties physicalDeliveryOfficeName
# For each user go through the if statements
foreach ($user in $Users)
{
if ($User.physicalDeliveryOfficeName –eq “London”)
{
New-Mailbox $User –Database “London Database 1”
}
elseif ($User.physicalDeliveryOfficeName –eq “Swindon”)
{
New-Mailbox $User –Database “Swindon Database 1”
}
else
{
New-Mailbox $User –Database “Mailbox Database 1”
}
}
```

The switch Statement

The **switch** statement works similar to **ifelse** statements. The **switch** statement enables a single condition statement to have multiple options for execution. The switch statement has the following syntax:

```
switch (Value Testing)
{
Value 1 { Code run if value 1 condition exists}
Value 2 { Code run if value 2 condition exists}
Value 3 { Code run if value 3 condition exists}
default { Code run if no other condition exists}
}
```

Using the previous example, you can achieve the same functionality in fewer lines, as shown in this example:

```
$Users = Get-ADUser –Filter * -Properties physicalDeliveryOfficeName
# For each user go through the if statements
foreach ($user in $Users)
{
switch ($User.physicalDeliveryOfficeName)
```

MCT USE ONLY STUDENT USE PROHIBITED

```
{
“London” { New-Mailbox $User –Database “London Database 1”}
“Swindon” { New-Mailbox $User –Database “Swindon Database 1”}
default { New-Mailbox $User –Database “Mailbox Database 1”}
}
}
```

If a larger number of false statements are needed, the switch statement may be an easier option to use and debug.

The for Loop

You can use a **for** loop to execute a block of code a specific number of times. For example, you may want to use a **for** loop to request or create multiple items. The **for** statement syntax is as follows:

```
for (setup loop variables ; Boolean comparison ; action after each loop)
{
Code to complete while Boolean comparison is true
}
```

The **for** loop begins with settings to configure variables, the Boolean comparison, and an action to complete after each loop. Consider the following example, which creates five mailbox databases with unique names, by using a **for** statement:

```
# Create a variable named $i and assign it a value of 1
# Execute the for loop for as long as $i is less than 6
# After each loop add 1 to the value of $i
for ($i = 1 ; $i –lt 6 ; $i++)
{
# Create a variable with the name of the Server and the database
$ServerName = “LON-MBX1”
$DatabaseName = “Database” + $i
#Create Mailbox Database
New-MailboxDatabase –name $DatabaseName –server $ServerName
}
```

The while Loop

You can use a **while** loop to execute a block of code while a specific condition exists. A **while** loop resembles a **for** loop, except that a **while** loop does not have built-in mechanisms to set up variables and actions to run after each loop. This differentiation means that the **while** statement continues to execute until a condition is met, rather than executing a set number of times. The **while** statement syntax is as follows:

```
while (Boolean comparison)
{
Code to complete while Boolean expression is true
}
```

The following creates mailbox databases until there are 15 mailbox databases. The value of the $i and $c variables must be set before the **while** loop executes, so that the **while** loop executes as follows:

```
$i = 100
$c = Get-MailboxDatabase | Measure-Object
$ServerName = “LON-MBX1”
while ($c.Count -lt 16)
{
 $DatabaseName = “Database” + $i
New-MailboxDatabase –name $DatabaseName –Server $ServerName
$c = Get-MailboxDatabase | Measure-Object
 $i++
}
```

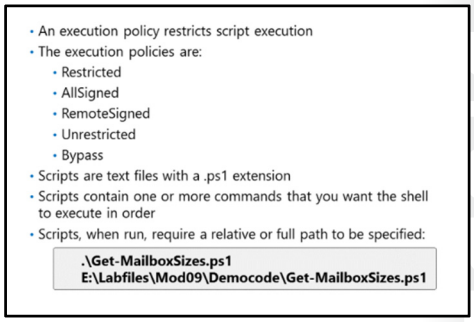
MCT USE ONLY. STUDENT USE PROHIBITED

Also available is the **do/while** loop, which is similar to the **while** loop, except that the Boolean expression is evaluated at the end of the loop instead of the beginning. This approach means that the code block in a **do/while** loop is always executed at least one time. The value of $c does not need to be set before the **do/while** loop, because $c is evaluated at the end of the loop. The following example shows a **do/while** loop:

```
$i = 100
$ServerName = "LON-MBX1"
do {
 $DatabaseName = "Database" + $i
 New-MailboxDatabase –name $DatabaseName –Server $ServerName
 $c = Get-MailboxDatabase | Measure-Object
} while ($c.Count -lt 16)
```

Creating Exchange Management Shell Scripts

You can perform complicated multi-step tasks by using a pipeline and multiple cmdlets. There may be times when you need to run multiple functions, make choices, wait for tasks to complete, or run the same code repeatedly. In these cases, you can use a Windows PowerShell script to put all the steps together. A script is a text-based file that includes at least one Windows PowerShell command and that is saved with a .PS1 file name extension. You can create scripts that take input from the command line, so you can customize how the script executes.



Execution Policy

By default, the execution policy does not enable Windows PowerShell scripts to be executed automatically. This constraint helps prevent unattended scripts from running without the administrator's knowledge. You can set any of the following execution policies:

- Restricted. This setting is the default policy for Windows Server 2012. With this policy, configuration files cannot load, and scripts cannot run. The Restricted execution policy is best for a computer that you do not run scripts on, or that you run scripts on only rarely. You can open the shell manually, if you need to, with a less restrictive execution policy.
- AllSigned. This policy requires that all scripts and configuration files be signed by a trusted publisher, including scripts that are created on your local computer. This execution policy is useful for environments where you do not want to run any script unless it has a trusted digital signature. This policy requires additional effort, because it requires you to digitally sign every script that you write and to re-sign each script every time that you make any changes to it.
- RemoteSigned. This policy requires that all scripts and configuration files downloaded from the Internet be signed by a trusted publisher. This execution policy is useful because it assumes that local scripts are ones that you create yourself, and that you trust them. It does not require local scripts to be signed. However, scripts that are downloaded from the Internet or are received through email are not trusted unless they carry an intact, trusted digital signature. You can still run those scripts—for example, by running the shell under a lesser execution policy, or even by signing the script yourself. But because you must take these additional steps, it is unlikely that you can run such a script accidentally or unknowingly.

MCT USE ONLY. STUDENT USE PROHIBITED

- Unrestricted. With this policy, you can load all configuration files and run all scripts. If you run a script that was downloaded from the Internet, you are warned about potential dangers and must grant permission for the script to run. The Unrestricted execution policy usually is not appropriate for production environments, because it provides little protection against accidentally or unknowingly running untrusted scripts.
- Bypass. With this policy, you can load all configuration files and run all scripts. If you run a script that was downloaded from the Internet, the script runs without any warnings. This execution policy is not usually appropriate for production environments, because it provides no protection against accidentally or unknowingly running untrusted scripts.

You can view the execution policy that is in effect for a particular computer by using the **Get-ExecutionPolicy** cmdlet. To configure the execution policy, open an elevated Windows PowerShell window, and then run the **Set-ExecutionPolicy** cmdlet. After the execution policy is configured, you can run a script by typing the name of the script.

Simple Scripts

Scripts are text files that have a .PS1 file name extension. These files contain one or more commands that you want the shell to execute in a particular order. You can edit scripts by using Notepad, but the Windows PowerShell ISE provides a better editing experience. In the ISE, you can type commands interactively, obtain hints about the correct command syntax, and immediately see the results. You then can paste those results into a script for long-term use. Or you can type commands directly into a script, highlight each command, and then press F8 to execute only the highlighted command. If you are pleased with the results, save the script and you are finished.

Generally, there are few differences between what you can do in a script and what you can do at the command line. Commands work in the same way in a script, which means that you can create a script simply by pasting commands that you have already tested at the command line. The following is a simple script in a text file that is named **Get-MailboxSizes.ps1**.

```
# This script returns the five largest mailboxes
Get-Mailbox | Get-MailboxStatistics | `
Sort-Object -Property TotalItemSize -Descending| `
Select-Object -First 5 | `
Format-Table displayName, TotalItemSize, ItemCount -AutoSize
```

Although this script contains a single pipeline statement, it is broken up by using the backtick (`) character. You can break up long lines of code by using the backtick to make the script easier to read. You do not need to use a backtick immediately after a pipe or a comma, as shown in the example. Notice that the first line of this script starts with a number sign (#). A line that begins with a number sign is not processed. Therefore, you can start a line with a number sign, and then write notes and comments about the script. To run a script, type either the full or the relative path of the script. For example, to run the **Get-MailboxSizes.ps1** script, you can use either of the following options if the script is in your current directory or search path:

```
.\Get-MailboxSizes.ps1
E:\Labfiles\Mod09\Democode\Get-MailboxSizes.ps1
```

If the script name or path contains spaces, enclose the name with single or double quotation marks, and echo the name to the console by using an ampersand (&) character. The following example shows how to do this by using both the relative and the full path.

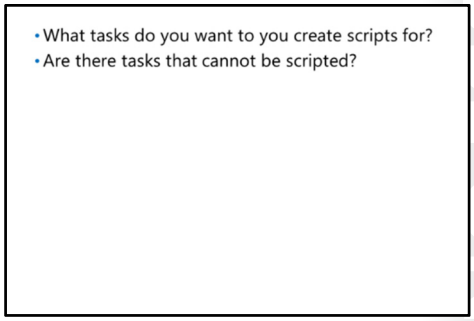
```
& '.\Get Mailbox Sizes.ps1'
& 'E:\Labfiles\Mod09\Democode\Get Mailbox Sizes.ps1'
```

MCT USE ONLY. STUDENT USE PROHIBITED

Discussion: Managing Exchange Recipients with Exchange Management Shell Scripts

Discuss the following questions:

- What tasks do you want to you create scripts for?
- Are there tasks that cannot be scripted?



Demonstration: Creating an Exchange Management Shell Script

Demonstration Steps

1. Sign in to LON-CAS1 as the domain administrator.
2. Open the Windows PowerShell ISE as an administrator, and then review the Script pane and the Console pane.
3. Open and execute **E:\Labfiles\Mod09\Democode\Add Users And Mailboxes.ps1**, and then follow the instructions.

MCT USE ONLY. STUDENT USE PROHIBITED

Lesson 3
Using Windows PowerShell to Manage Exchange Server

As you become familiar with Windows PowerShell, you can perform administrative and management tasks more easily. Windows PowerShell 3.0 has many Exchange Management Shell cmdlets and advanced features that you can use to perform numerous management tasks. This lesson introduces some of the Exchange Management Shell cmdlets and advanced features of Windows PowerShell 3.0 and discusses how you might use the features to manage servers in your environment.

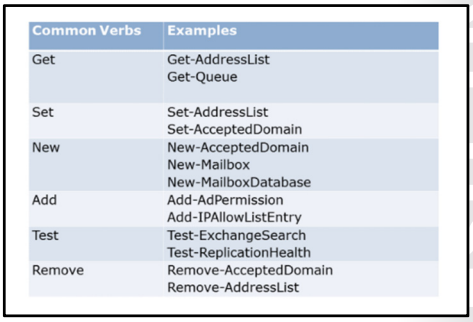
Lesson Objectives

After completing this lesson, students will be able to:

- Describe common Exchange Management Shell cmdlets.
- Manage server configuration by using Exchange Management Shell.
- Test Exchange Server 2012 by using the built-in test cmdlets.
- Describe Windows PowerShell jobs.
- Manage Exchange Server by using scheduled jobs and test cmdlets..
- Monitor Exchange Server by using the Exchange Management Shell.
- Describe the built-in Windows PowerShell Scripts.

Overview of Exchange Management Shell Cmdlets

In addition to managing Exchange Server recipients, you can also use Windows PowerShell for all Exchange Server configuration tasks. You may choose to create a script that automates the configuration of a new Client Access server, to create reports, or to perform other management activities. All of the names of the Exchange Management Shell cmdlets are standardized verb-noun pairs. The verb in the cmdlet name is the action the cmdlet performs. The noun is the object that the action is performed on or with. For example, the **Mount-Database** cmdlet mounts the specified database object.



Common Verbs	Examples
Get	Get-AddressList Get-Queue
Set	Set-AddressList Set-AcceptedDomain
New	New-AcceptedDomain New-Mailbox New-MailboxDatabase
Add	Add-AdPermission Add-IPAllowListEntry
Test	Test-ExchangeSearch Test-ReplicationHealth
Remove	Remove-AcceptedDomain Remove-AddressList

To better understand the Exchange Management Shell cmdlets, you must be familiar with the common verbs used. These common verbs are listed in the following table.

Common verbs	Cmdlet examples	Description
Get	**Get-AddressList** **Get-Queue**	Retrieves a resource or object. Use this cmdlet to view the properties or settings of most objects.
Set	**Set-AddressList** **Set-AcceptedDomain**	Configures settings or makes changes to an object. Use this cmdlet to change settings for most objects.
New	**New-AcceptedDomain**	Creates new objects or resources.

MCT USE ONLY. STUDENT USE PROHIBITED

Common verbs	Cmdlet examples	Description
	New-Mailbox **New-MailboxDatabase**	
Add	**Add-AdPermission** **Add-IPAllowListEntry**	Adds a value, resource, or object to another object or container.
Test	**Test-ExchangeSearch** **Test-ReplicationHealth**	Tests functionality.
Remove	**Remove-AcceptedDomain** **Remove-AddressList**	Removes objects, values, and resources.

Standardized names help you easily determine what a cmdlet does, and they help you find a cmdlet that accomplishes a specific task. Therefore, if you know the verb, you can find the **Get-Command** cmdlet to list all of the cmdlets that are related. For example, to list all of the cmdlets that use the Test verb, run the following from the Exchange Management Shell:

```
Get-Command –Name “Test-”
```

If you know the name of the noun and you need to know the verbs that you can use with it, you can use **Get-Command** to find the available cmdlets. For example, to list all of the cmdlets that use the **AddressList** noun, run the following from the Exchange Management Shell:

```
Get-Command –Name “*-AddressList“
```

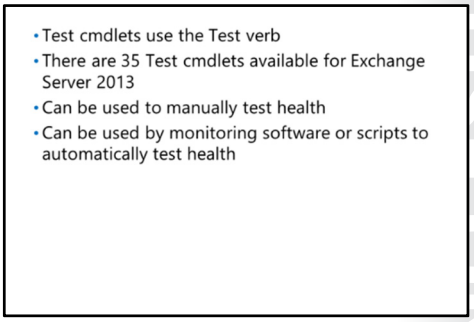
Demonstration: Managing Server Configuration by Using Exchange Management Shell

Demonstration Steps

1. Sign in to LON-CAS1 as the domain administrator.
2. Open the Windows PowerShell ISE as an administrator, and then review the Script pane and the Console pane.
3. Open and execute **E:\Labfiles\Mod09\Democode\Managing Exchange Settings.txt**, and then follow the instructions.

MCT USE ONLY. STUDENT USE PROHIBITED

Overview of Exchange Server 2013 Test Cmdlets

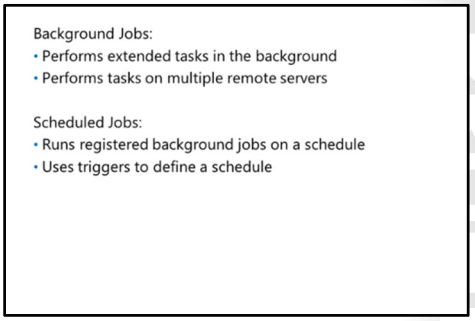


There are 35 cmdlets available for performing tests within an Exchange Server organization. These tests can identify mail flow, connectivity, and replication problems. You can use these test cmdlets to test Exchange Server functionality manually. Also, the cmdlets can be used by monitoring tools, such as Microsoft System foreach (item in collection) Center 2012 Operations Manager, to test Exchange Server functionality automatically.

A few of the more common test cmdlets are the following:

- **Test-ReplicationHealth** reports on the status of the database copies within a database availability group (DAG). For example, a monitoring system can proactively check replication, the availability of Active Manager, and the health of a cluster service.
- **Test-OutlookConnectivity** tests the end-to-end client connectivity of the Microsoft Office Outlook® messaging client to the Exchange Server organization. For example, a monitoring system could run this cmdlet automatically to determine whether users can connect by using Outlook.
- **Test-ServiceHealth** reports the status of the Windows Server services that relate to Exchange Server on the server specified.
- **Test-Mailflow** tests all aspects of mail transport by verifying that each Mailbox server can successfully send itself a message.
- **Test-SmtpConnectivity** diagnoses whether an SMTP connection can be established with the specified server.
- **Test-WebServicesConnectivity** tests the functionality of the Exchange Web Services. Service functionality is tested by executing commands to list, create, and delete items.

What Are Windows PowerShell Jobs?



A Windows PowerShell background job runs a command or a set of commands without interacting with the current Windows PowerShell session. You can start a background job by using the **Start-Job** cmdlet, and then you can continue to work in the session. Jobs can help you perform tasks that take a long time to finish. You can also use jobs to perform the same task on several computers. The following example shows how to create a job on the local computer:

```
Start-Job -ScriptBlock {Test-
ServiceHealth}
```

You can see the status of the job by using the **Get-Job** cmdlet, and you can use the **Wait-Job** cmdlet to be notified when the job finishes. If you want to remove a job that has not run yet, use the **Remove-Job** cmdlet. These jobs run in the background, so they do not return results to your Windows PowerShell

MCT USE ONLY. STUDENT USE PROHIBITED

session. If you output data to the console in a background job, you can return that data by using the **Receive-Job** cmdlet.

Windows PowerShell 3.0 introduces an improvement to background jobs, called scheduled jobs. Scheduled jobs can be triggered to start automatically, or they can be performed on a recurring schedule. When a scheduled job is created, it is stored on disk and then registered in Task Scheduler. When a scheduled job runs, it creates an instance of the job that then can be managed by using the common job management cmdlets. The only difference between scheduled jobs and background jobs is that scheduled jobs save the results on disk.

To create a scheduled job, use the **Register-ScheduledJob** cmdlet. You can specify the **ScriptBlock** parameter to run a Windows PowerShell command, or you can specify a script by using the **FilePath** parameter. The following example shows how to register a scheduled job to run the **Get-MailboxSizes.ps1** script.

```
Register-ScheduledJob –Name MailboxSizes –FilePath \\LON-
CAS1\Scripts\Mod09\democode\Get-MailboxSizes.ps1
```

To enable the scheduled job to run, you must define a schedule or a trigger. To create a trigger, use the **New-JobTrigger** cmdlet. Then use the **Add-JobTrigger** cmdlet to add the trigger to an already registered scheduled job, or to assign a trigger when a new scheduled job is registered. You can schedule triggers once, daily, weekly, at server startup, or when you sign in. The following example shows how to create a trigger that runs every Monday and Friday at 9:00 A.M., and that registers the new scheduled job together with the trigger:

```
$Trigger = New-JobTrigger –Weekly –DaysOfWeek Monday,Friday –At 9:00AM `
Register-ScheduledJob –Name ScheduledMailboxSizes –FilePath `
\\LON-CAS1\Scripts\Mod09\democode\Get-MailboxSizes.ps1 -Trigger $Trigger
```

You can also use the **Add-JobTrigger** cmdlet to modify an existing scheduled job, as the following the example shows:

```
Add-JobTrigger -Name MailboxSizesJob -Trigger `
(New-JobTrigger -Daily -At 9:00AM)
```

You can use scheduled jobs to automatically create reports, monitor service health, verify configuration settings, perform user and group maintenance, and many other tasks.

Demonstration: Managing Exchange Servers by Using Test Cmdlets and Scheduled Jobs

Demonstration Steps

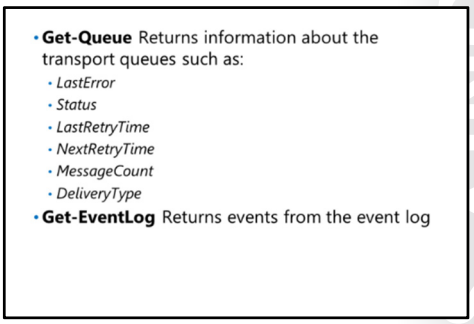
1. Start virtual machines LON-DC1, LON-CAS1, and LON-MBX1, and then sign in to LON-CAS1 as the domain administrator.
2. Open the Exchange Management Shell.
3. Execute the **Test-ReplicationHealth** cmdlet and review the results.
4. Execute the **Test-ServiceHealth** cmdlet for **LON-MBX1** and view the results.
5. Start a new job to test the Exchange service health, by using the **Start-Job** cmdlet to run **E:\Labfiles\Mod09\democode\Health.ps1**.
6. Obtain the status of the job by running **Get-Job**.
7. Create a new scheduled job by running the following commands, each followed by pressing Enter.

MCT USE ONLY. STUDENT USE PROHIBITED

```
$Trigger = New-JobTrigger –Weekly –DaysOfWeek Monday,Friday –At 9:00AM
Register-ScheduledJob –Name ScheduledJob1 –FilePath
E:\Labfiles\Mod09\democode\health.ps1 -Trigger $Trigger
```

8. Run the scheduled job immediately by using the **Start-Job** cmdlet.

Monitoring Exchange Servers by Using the Exchange Management Shell



Sometimes, you need to look at specific Exchange Server components to address potential issues. You can use the Test cmdlets to quickly access a wide variety of standard tests. Some of these cmdlets are as follows:

- **Get-Queue**. Views the status of the transport queues.
- **Get-EventLog**. Views events in the specified Event log.

The **Get-Queue** cmdlet returns information about the transport queues. You can use this information to monitor the overall transport health. Some of the returned information is as follows:

- *LastError*. If the queue is in a failed state, the error listed occurred the last time delivery was attempted for the queued messages.
- *Status*. The current status of the queue.
- *LastRetryTime*. The time when delivery was last attempted.
- *NextRetryTime*. The time then the next delivery will be attempted.
- *MessageCount*. The number of messages in the queue.
- *DeliveryType*. The type of queue; for example, when the queue is for delivering shadow redundancy messages, delivering messages to an external recipient, or delivering to another mailbox within the Exchange Server organization.

You can use the **Get-Queue** cmdlet to easily return queues that have a large number of queued messages. For example, to return queues that have more than 500 queued messages, run the following:

```
Get-Queue –Filter {MessageCount –gt 500}
```

You can use this command in a monitoring script or in a script that is scheduled as a Windows PowerShell job.

The **Get-EventLog** cmdlet is not specific to Exchange Server, but you can use it to gather information about Exchange Server from the Event Logs. For example, to return the most recent ten Event Log entries that have a source of MSExchange Common, run the following command.

```
Get-EventLog –LogName Application –Source “MSExchange Common” –EntryType Error –
Newest 10
```

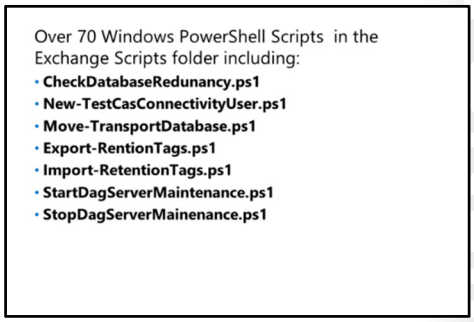
You can modify the command to return events that contain specific error numbers or words. You can also return data from multiple computers. The following command returns the most recent ten events from multiple servers.

MCT USE ONLY. STUDENT USE PROHIBITED

```
Get-EventLog –LogName Application –Source “MSExchange Common” –EntryType Error –
Newest 10 –Computer LON-CAS1, LON-MBX1
```

Using Provided Exchange Management Scripts

Over 70 Windows PowerShell scripts are included in the Scripts folder of the Exchange install. Some of these scripts can be useful when performing administrative functions. Some of the more important scripts available are the following:



- **CheckDatabaseRedundancy .ps1** verifies if databases have enough configured and healthy copies. Use this to view a report on the redundancy of the copies of the mailbox databases.
- **New-TestCasConnectivityUser.ps1** configures the mailbox needed for several of the Test cmdlets. You use this to configure the mailbox settings for the test cmdlets manually.
- **Move-TransportDatabase.ps1** moves the locations of the Transport database. This is useful when you need to move the Transport database to new storage for performance or capacity reasons.
- **Export-RentionTags.ps1** and **Import-RetentionTags.ps1** are used to export the retention tags from one Exchange organization and into another Exchange organization. This is useful to keep retention policies synchronized between two organizations.
- **StartDagServerMaintenance.ps1** prepares a DAG member for maintenance by suspending all database copies and activations, pausing the node in Failover Clustering, and moving all active databases to other DAG members. Use this script before performing maintenance on a DAG member.
- **StopDagServerMainenance.ps1** removes the settings made by StartDagServerMainenance.ps1.However it does not reactivate resources back to the server. Use this script after you complete maintenance on a DAG member that you ran StartDagServerMaintenance.ps1 on.

Use the **Get-Help** cmdlet to view the documentation before trying to use any of the scripts. Do not modify any of the provided scripts. If you do customize the scripts, save the modifications with a new name in another folder so that the provided scripts remain unmodified. This also ensures that your modified scripts will not be overwritten by an Exchange update that includes updates to the provided scripts.

MCT USE ONLY STUDENT USE PROHIBITED

Lab: Managing Microsoft Exchange Server 2013 by Using Exchange Management Shell

Scenario

As the A. Datum organization grows in size and complexity, it is becoming increasingly apparent that some of the Information Technology (IT) management processes need to be streamlined. This requirement includes managing the Exchange Server deployment. Because A. Datum adds new branch offices and acquires other companies frequently, it is important that recipient management is quick and efficient. Also, the server management tasks need to be more consistent and efficient.

To address these management issues, you need to be familiar with the Exchange Management Shell and how to use it to manage the Exchange Server organization. You need to understand how to run simple and complex commands and how to create scripts that automate many of the regular management tasks.

Objectives

After completing this lab, you will be able to:

- Identify and use the key functionalities of the Exchange Management Shell.
- Manage recipients by using the Exchange Management Shell.

Lab Setup

Estimated Time: 60 minutes

Virtual machines	**20342B-LON-DC1** **20342B-LON-CAS1** **20342B-LON-MBX1**
User Name	Adatum\Administrator
Password	Pa$$w0rd

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

1. On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In Microsoft Hyper-V® Manager, click **20342B-LON-DC1**, and then in the Actions pane click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Log on by using the following credentials:
   - User name: **Adatum\Administrator**
   - Password: **Pa$$w0rd**
5. Repeat steps 2 and 3 for the following virtual machines: **20342B-LON-CAS1** and **20342B-LON-MBX1**.

MCT USE ONLY. STUDENT USE PROHIBITED

Exercise 1: Exploring the Exchange Management Shell

Scenario

You want to explore the Windows PowerShell ISE interfaces and the available cmdlets in order to become familiar with the Exchange Management Shell interface.

The main tasks for this exercise are as follows:

1. Import the Exchange Management Shell module into the Windows PowerShell ISE

2. Generate a Table view of the Deleted Item retention settings of all mailbox databases in the Exchange organization

3. Create a Windows PowerShell job to return the five most recent events from the Application Event log on LON-CAS1 and LON-MBX1

▶ Task 1: Import the Exchange Management Shell module into the Windows PowerShell ISE

1. Sign in to the LON-CAS1 virtual machine with the user name of **Adatum\Administrator** and the password **Pa$$w0rd**.
2. Open the Windows PowerShell ISE as an administrator, and then review the Script pane and the Console pane.
3. Import the Exchange Management Shell Module by typing Import-Module 'C:\Program Files\Microsoft\Exchange Server\V15\bin\RemoteExchange.ps1'; Connect-ExchangeServer -auto -ClientApplication:ManagementShell.

▶ Task 2: Generate a Table view of the Deleted Item retention settings of all mailbox databases in the Exchange organization

1. Use the **Get-MailboxDatabase** cmdlet to list all of the databases.
2. Pipe the output of the **Get-MailboxDatabase** cmdlet to the **Format-Table** cmdlet. Display the output of the **Name** and the **DeletedItemRetention** properties.

▶ Task 3: Create a Windows PowerShell job to return the five most recent events from the Application Event log on LON-CAS1 and LON-MBX1

1. Create a job to return the five most recent events from the Application Event Log, and then assign the job to a variable named **$job**.
2. Check the status of the created job by using the **Get-Job** cmdlet.
3. Receive the job results by using the **Receive-Job** cmdlet.

Results: After completing this lab, you will have used the Exchange Management Shell and performed basic management tasks.

Exercise 2: Using Exchange Management Shell to Manage Recipients

Scenario

One of the primary goals in using Exchange Management Shell is to manage recipients more efficiently. You want to explore the various Exchange Management Shell options for managing recipient objects.

The main tasks for this exercise are as follows:

MCT USE ONLY. STUDENT USE PROHIBITED

1. Create a list of all of the cmdlets that are available to manage mailbox objects
2. Modify a script to create new mailboxes from a CSV file
3. Configure settings on mailboxes by using the Exchange Management Shell
4. To prepare for the next module

▶ Task 1: Create a list of all of the cmdlets that are available to manage mailbox objects

1. Use the **Get-Command** cmdlet to list all available cmdlets.
2. Use the **Get-Command** cmdlet to list all cmdlets that include **Mailbox**.

▶ Task 2: Modify a script to create new mailboxes from a CSV file

1. On LON-CAS1, use Notepad.exe to edit **E:\Labfiles\Mod09\labfiles\ AddUsers.csv**.
2. Modify the entry for Jim by removing the extra column, and then save the updated file as **E:\labfiles\Mod09\labfiles\AddConsultants.csv**.
3. Use the Windows PowerShell ISE to open the script **E:\Labfiles\Mod09\labfiles\AddUsers.ps1**.
4. Modify the **$OU** variable to read: **$OU = "Consultants"**.
5. Modify the line beginning with $password to read: **$password = ConvertTo-SecureString "Pa$$w0rd!" –AsPlainText –Force**.
6. Modify the line beginning Import-CSV to read: Import-CSV E:\Labfiles\Mod09\Labfiles\AddConsultants.CSV | Where-Object {New-Mailbox –Alias $_.SamAccountName `.
7. Modify the line beginning -OrganizationalUnit to read: -OrganizationalUnit $OU –UserPrincipalName $_.UserPrincipalName –Password $password `.
8. Save the updated script as **AddConsultants.ps1**.
9. Run the AddConsultants.ps1 script.
10. Use Get-ADUser –Filter * –SearchBase "OU=consultants,DC=Adatum,DC=com" to verify that accounts for Darren Waite, Ioannis Xylaras, and Marko Zajc are created.

▶ Task 3: Configure settings on mailboxes by using the Exchange Management Shell

1. Using the **Set-Mailbox** cmdlet, set the office name for all users located in the Research OU to **Research**.
2. Using the **Set-Mailbox** cmdlet, set the Mail Tip for all users in the IT distribution group to be **If you require IT assistance please contact the Help Desk**.

▶ Task 4: To prepare for the next module

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

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20342B-LON-CAS1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20342B-LON-MBX1** and **20342B-LON-DC1**.
5. In Hyper-V Manager, click **20342B-LON-DC1**, and in the Actions pane, click **Start**.

MCT USE ONLY. STUDENT USE PROHIBITED

6. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
7. Sign in using the following credentials:
   a. User name: **Adatum\Administrator**
   b. Password: **Pa$$w0rd**
8. Repeat steps 5 to 7 for **20342B-LON-MBX1, 20342B-LON-CAS1, 20342B-LON-CAS2, 20342B-LON-CL1, 20342B-TREY-DC1,** and **20342B-TREY-EX1**. Sign in to the TREY virtual machines as **TreyResearch\Administrator** with the password of **Pa$$w0rd**.

Results: After completing this lab, you will have used the Exchange Management Shell and performed basic recipient management tasks.

Question: What happens if you try to run an Exchange Server cmdlet and do not have the Exchange Management Shell module imported?

Question: Why do you need to specify the **DeletedItemRetention** property for the **Format-Table** cmdlet?

MCT USE ONLY. STUDENT USE PROHIBITED

Module Review and Takeaways

Best Practice

Set a goal to spend time learning how to use the Exchange Management Shell and Windows PowerShell for your common tasks. As you become more familiar with Windows PowerShell, you will become better prepared to use it for more complex tasks and for resolving certain problems.
Save the commands that you use to resolve problems in a script file for later reference.
Use the Windows PowerShell ISE to help you write scripts and to ensure that you are using the correct syntax.

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
You cannot find the correct Windows PowerShell cmdlet for a task.	

Review Question(s)

Question: Which cmdlet creates a new mailbox?

Question: Which cmdlet mounts a database?

Tools

You can use the tools in the following table to work with Windows PowerShell.

Tool	Description
Windows PowerShell Integrated Script Editor (ISE)	A simple, powerful interface to create and test scripts, and to discover new cmdlets.
Microsoft Visual Studio Workflow Designer	A development tool that you can use to create Windows PowerShell workflows.
Powershell.exe	The Windows PowerShell executable.
Exchange Management Shell	A shortcut that automatically imports the Exchange Management PowerShell module.

MCT USE ONLY. STUDENT USE PROHIBITED

MCT USE ONLY. STUDENT USE PROHIBITED

Module 10

Designing and Implementing Integration with Microsoft Exchange Online

Contents:

Module Overview **10-1**

Lesson 1: Planning for Exchange Online **10-2**

Lesson 2: Planning and Implementing the Migration to Exchange Online **10-10**

Lesson 3: Planning to Coexist with Exchange Online **10-15**

Lab: Designing Integration with Exchange Online **10-25**

Module Review and Takeaways **10-27**

Module Overview

Increasingly, Exchange administrators who run Exchange servers on premises are considering migrating to Microsoft® Exchange Online to help reduce operational, licensing, and setup costs.

Exchange Online is part of the Microsoft Office 365™ hosted productivity software, which in turn is part of Microsoft Online Services. Office 365 provides cloud-based versions of specific Microsoft products. This module examines the features of Office 365 and of Exchange Online, and it helps you plan an Exchange Online solution.

Microsoft Exchange Server 2013 can connect an existing Exchange Server organization to Exchange Online. You can use a hybrid deployment to allow collaboration between users of Exchange Server mailboxes and users of Exchange Online mailboxes. You can also use federated delegation to enhance collaboration by sharing information between Exchange on-premises organizations and Exchange Online organizations.

Objectives

After completing this module, you will be able to:

- Plan for Exchange Online.
- Plan and Implement a migration to Exchange Online.
- Plan coexistence with Exchange Online.

MCT USE ONLY. STUDENT USE PROHIBITED

Lesson 1
Planning for Exchange Online

If your organization currently does not have an email messaging system, you can set up Exchange Online as the messaging system. If your organization already has a messaging system, you need to understand how Exchange Online can coexist with the existing messaging system.

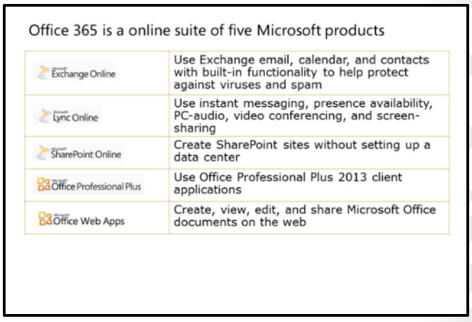
Lesson Objectives

After completing this lesson, you will be able to:

- Describe Office 365.
- Describe the new features available in Exchange Online.
- Describe Exchange Online Protection (EOP).
- Discuss your Office 365 experiences and expectations.

What Is Office 365?

Office 365 is an online suite of the following Microsoft services: Exchange Online, Microsoft Lync® Online, Microsoft SharePoint® Online, Microsoft Office Professional Plus, and Microsoft Office Web Apps. It is a subscription-based service with a range of pricing options.



Exchange Online

Exchange Online provides Exchange Server email, calendar, and contacts, plus antivirus and anti-spam features. You can connect an existing Exchange Server 2013 organization to Exchange Online to provide a hybrid deployment, which includes free/busy information and MailTips for users. Exchange Online features are described in the next topic.

Lync Online

Lync Online provides users with instant messaging (IM), presence availability, online meeting infrastructure, audio and video calling, and screen sharing. You can connect an organization's existing servers running Microsoft Office Communications Server 2007, Microsoft Lync Server 2013, or Microsoft Lync Server 2010 to Lync Online.

SharePoint Online

You can use SharePoint Online to create and manage SharePoint sites directly from the cloud. Because you can share documents or keep teams updated by using a common SharePoint team site, you do not need to set up SharePoint in the organization's data center. You can also share a SharePoint site between organizations if you do not want to set up servers in a perimeter data center.

Office Professional Plus

Office Professional Plus is an online version of Microsoft Office 2013. Office Professional Plus includes all Office desktop applications, including Microsoft Word and the Microsoft Excel® spreadsheet software. Office Professional Plus also uses an on-demand, per-user license model, and it connects to the cloud.

Additionally, Office Professional Plus has the following features:

MCT USE ONLY. STUDENT USE PROHIBITED

- Microsoft Office Professional Plus 2013 client applications are available as a monthly subscription.
- Per-user licensing with as many as five simultaneous installations.
- Support for 32-bit and 64-bit installations.
- Easy access and management through the Microsoft Office 365 Portal.

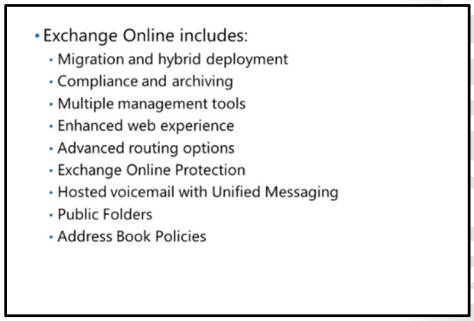
Note: Office Professional Plus is not a streaming client application. Office Professional Plus provides the full Office Professional 2013 feature set on the local machines, but it differs from Office Professional 2013 in license management.

Office Web Apps

Users can use Office Web Apps to create, view, share, and edit Microsoft Office documents directly on the web. They do not need a locally installed version of Microsoft Office in order to work on their content.

What's New in Exchange Online?

Exchange Online includes most of the features that are available in Exchange Server 2013, plus additional features such as the ability to migrate mailboxes in the Exchange Administration Center (EAC). Other key Exchange Online features include the following:



- Migration and hybrid deployment. Exchange Online provides migration tools, which you can use to automatically move users to Exchange Online. Alternatively, you can connect a Microsoft Exchange Server 2003, Microsoft Exchange Server 2007, or Exchange Server 2013 environment to the cloud and then use hybrid deployment features. In a hybrid deployment, users can share calendar free/busy data between the cloud and on-premises, and you can migrate the users between the two whenever necessary.
- Compliance and archiving. Exchange Online provides the archiving and eDiscovery capabilities of Exchange Server 2013, including built-in personal archives, multi-mailbox search, retention policies, transport rules, and optional legal holds to preserve email.
- Multiple management tools. Exchange Online includes management tools, such as the EAC, the Windows® PowerShell® command-line interface, and the Office 365 Administration Center. The web-based EAC in Exchange Server 2013 is closely integrated with Exchange Online, so you can manage policies, security, user accounts, and groups. You can also use Windows PowerShell to remotely manage all aspects of a hosted Exchange Server environment across the Internet.
- Enhanced web experience. The Microsoft Outlook® Web App experience is available through the Windows Internet Explorer® browser, Firefox, and Safari. Instant messaging is integrated, so users can chat from within Outlook Web App.
- Advanced routing options. You can use Exchange Online to route outbound email through the on-premises infrastructure. This feature means that you can perform custom post-processing of outbound email, use non-Microsoft data loss prevention (DLP) appliances, and deliver email to business partners through private networks.

MCT USE ONLY. STUDENT USE PROHIBITED

- Exchange Online Protection. Exchange Online Protection is included for automatic anti-spam and antivirus scanning.
- Hosted voicemail with Unified Messaging. You can replace your on-premises voicemail system by integrating your on-premises private branch exchange (PBX) with hosted voicemail provided by Exchange Online.
- Public Folders. Exchange Online supports Public Folders and Public Folder migration from an on-premises environment.
- Address Book Policies. Address book policies are available in Exchange Online to fine tune address lists.

Note: When referring to the local Exchange Server organization, we use the term *on-premises* to differentiate it from Exchange Online.

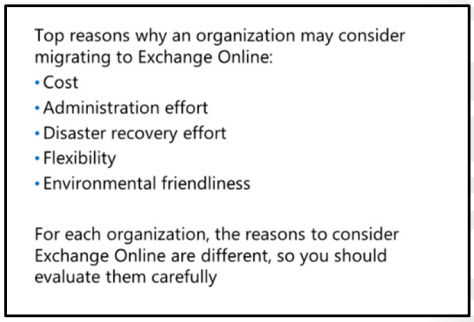
Exchange Online supports the following messaging clients:

- Outlook Web App (Internet Explorer, Firefox, and Safari).
- Microsoft Office 2007, Office 2010, and Office 2013.
- Microsoft Entourage® 2008 for Mac Exchange Web Services Edition.
- Microsoft Outlook for Mac 2011 (without Personal Archives).
- Any non-Microsoft Internet Message Access Protocol 4 (IMAP4) or Post Office Protocol 3 (POP3) client.

Note: Exchange Online features are subject to change. For update feature lists, see "Exchange Online for Enterprises Service Description" (http://go.microsoft.com/fwlink/?LinkId=290681).

Why Migrate to Exchange Online?

Discussions about the advantages and disadvantages of Exchange Online have been ongoing since the introduction of an Exchange Server version as a cloud service. The most common reasons for an organization to consider moving to Exchange Online include the following:



- Cost. The main reason organizations consider moving to Exchange Online is to minimize the cost of a regular mailbox, because Exchange Online subscription fees cover all capital and operational, licensing, patching, and setup costs.
- Administration effort. Because Exchange Online is managed as a service, the messaging administrators do not need to plan or implement system patches or plan for or manage server failures, so they have more time for other projects.
- Disaster recovery effort. Exchange Online provides standard disaster recovery mechanisms, including data center failovers. Messaging administrators do not need to build a test environment or regularly

MCT USE ONLY. STUDENT USE PROHIBITED

train for disasters, because the hosted service manages disaster recovery if all mailboxes are on Exchange Online.

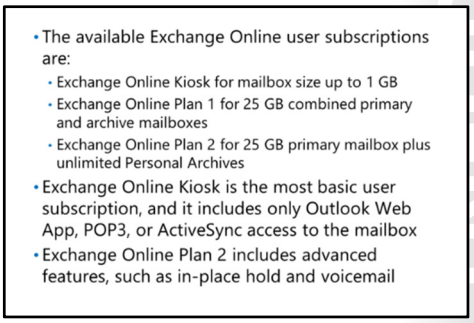
- Flexibility. You can use Exchange Online to meet demands quickly as business requirements change. You can increase or decrease the number of mailboxes almost immediately without needing to plan for or build additional hardware. For example, if an organization merges with another company, all mailboxes can be available almost immediately, because Exchange Online has sufficient resources available.
- Environmental friendliness. An on-premises Exchange Server deployment requires that you have a certain number of physical servers available to satisfy your messaging requirements. With Exchange Online, you save the physical space and power needed for your messaging servers, so you can decrease your organization's carbon footprint.

Note: The advantages of Exchange Online are not the same for all organizations. You might find other reasons specific to your organization as you start to consider migrating to Exchange Online.

Exchange Online User Subscriptions

To provide Exchange Online to your users, you need a user subscription license for each user. You can subscribe to only Exchange Online, or to Exchange Online along with other features in Office 365. Office 365 is available in the following service plans:

- Office 365 for small businesses.
- Office 365 for midsize businesses and enterprises.



The Office 365 for small businesses service plan provides Exchange Online Kiosk subscriptions. The Office 365 for enterprises service plan includes all subscription options. One of the components of the Office 365 for enterprises service plan is Exchange Online.

When you subscribe to Exchange Online, you can choose one of the following service plans:

- Exchange Online Kiosk (part of Office 365 Enterprise K1).
- Exchange Online Plan 1 (part of Office 365 Enterprise E1/E2).
- Exchange Online Plan 2 (part of Office 365 Enterprise E3/E4).

MCT USE ONLY. STUDENT USE PROHIBITED

The following table describes the features that are available in each Exchange Online user subscription.

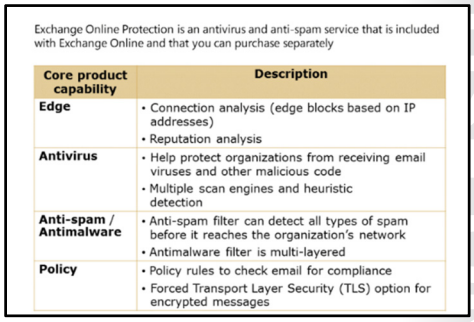
Feature	Exchange Online Kiosk	Exchange Online Plan 1	Exchange Online Plan 2
Mailbox size	1 gigabyte (GB)	25 GB shared between the primary mailbox and archive mailbox	25 GB for the user's mailbox, plus unlimited archive mailbox storage
Outlook Web App (regular and light versions)	Yes	Yes	Yes
POP3	Yes	Yes	Yes
IMAP4	No	Yes	Yes
Microsoft Outlook	No	Yes	Yes
Microsoft Exchange ActiveSync® technology	Yes	Yes	Yes
Exchange Web Services + Macintosh Clients	No	Yes	Yes
Blackberry services	No	Yes	Yes
Inbox rules	Yes	Yes	Yes
Delegate access	Yes	Yes	Yes
Lync interoperability	Yes	Yes	Yes
Personal Archives	No	Yes	Yes
Voicemail (Unified Messaging)	No	No	Yes
In-place hold	No	No	Yes
Office 365 admin center access	No	Yes	Yes
EAC access	Yes	Yes	Yes

Note: Exchange Online subscription options are subject to change. For updated information, see the Office 365 website (http://go.microsoft.com/fwlink/?LinkId=290682).

MCT USE ONLY. STUDENT USE PROHIBITED

What Is Exchange Online Protection?

Current messaging environments require a robust antivirus and anti-spam solution to minimize the impact of malicious messaging. EOP is an antivirus, anti-spam service that is included with Exchange Online and that can be purchased separately for an on-premises Exchange Server environment. EOP is a hosted service, so it requires no hardware or software installation.


Exchange Online Protection is an antivirus and anti-spam service that is included with Exchange Online and that you can purchase separately

Core product capability	Description
Edge	• Connection analysis (edge blocks based on IP addresses) • Reputation analysis
Antivirus	• Help protect organizations from receiving email viruses and other malicious code • Multiple scan engines and heuristic detection
Anti-spam / Antimalware	• Anti-spam filter can detect all types of spam before it reaches the organization's network • Antimalware filter is multi-layered
Policy	• Policy rules to check email for compliance • Forced Transport Layer Security (TLS) option for encrypted messages

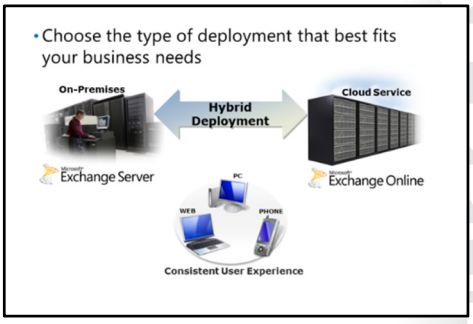
Exchange Online Protection includes the following functionality:

- Incoming, outgoing, and internal email messages are scanned. This scanning helps protect your organization from malicious content that originates behind your firewall.
- Multiple antivirus engines help catch email-borne viruses and other malicious code.
- Proprietary anti-spam technology is used to achieve high accuracy rates.
- All functionality is built in to the service. No configuration is necessary to start or to maintain the filtering technology. EOP requires only an EOP Send connector in the on-premises Exchange Server environment, so that messages are sent to the EOP domain for scanning. If you use only Exchange Online, you do not need to do any additional configuration.
- Customizable filters help you comply with corporate policies and with government regulations.

If you register for Exchange Online or Office 365, you automatically use EOP for any message that is received in or sent from your online tenant. You do not need to do any extra configuration. The Hybrid Configuration Wizard in Exchange Server 2013 configures EOP automatically by customizing the Send and Receive connectors.

Exchange Online Deployment Scenarios

When considering Exchange Online, decide between the following Exchange Online deployment scenarios by considering your business needs:



- Maintain only an on-premises Exchange Server organization. Do not use Exchange Online.
- Migrate your organization's on-premises Exchange Server installation to Exchange Online, so that all users are hosted in the cloud.
- Connect your organization's on-premises Exchange Server installation to Exchange Online in a hybrid deployment.

Exchange Server On-Premises

In an Exchange Server on-premises installation, you maintain a local installation of Exchange Server in your data center. This arrangement means that your organization manages its messaging environment in its own data center.

MCT USE ONLY. STUDENT USE PROHIBITED

In the on-premises scenario, you can perform maintenance, upgrades, and customization at your convenience. However, this scenario requires considerable upfront capital for hardware, software, licenses, IT personnel for maintenance, and physical building space.

Exchange Online

In an Exchange Online environment, all of the mailboxes are hosted in the cloud. You do not host any Exchange servers in your data center. Instead, you purchase the Exchange Online service from Microsoft. The advantages of this scenario are rapid deployment and easy scalability. You also receive automatic upgrades to the most recent technology, helping to ensure an easy and seamless upgrade experience.

Hybrid Deployment

An environment that mixes Exchange on-premises and Exchange Online is called a *hybrid deployment*. In a hybrid deployment, features such as free/busy information and calendar sharing are available between the on-premises and online mailbox users. This scenario provides features of both implementations, such as hosting the primary mailboxes on-premises and moving the archive mailboxes to Exchange Online. Additionally, with a hybrid deployment, you can migrate to Exchange Online in stages.

In a hybrid deployment, you can do the following:

- Manage Exchange Server on-premises and Exchange Online from a single instance of the Exchange Admin Center or the Exchange Management Shell.
- Move mailboxes between the Exchange Server on-premises and Exchange Online by using the EAC or the Exchange Management Shell.
- Share calendaring, including free/busy information, between on-premises and Exchange Online users.
- Resolve addresses for internal users against the GAL.
- Use MailTips, anti-spam scanning, and out-of-office auto-replies that count on-premises and Exchange Online recipients as internal.
- Use delivery reports to track messages across Exchange Server on-premises and Exchange Online.
- Perform multi-mailbox searches across all mailboxes.

Considerations

To decide on the most suitable scenario for your organization, consider the following questions:

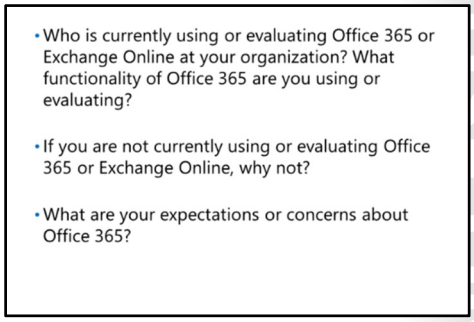
- Do you want to move all mailboxes to Exchange Online, only a subset of mailboxes, or no mailboxes?
- Do you want to move just some of the functionality—such as mailbox archiving—to the cloud?
- Does your organization often use mailbox delegation? If yes, ensure that both the mailboxes and the mailboxes with delegation rights to those mailboxes are hosted either online or on-premises.
- Is it important to have full control over the features and functionality of your messaging system?
- Does your organization have organizational policies, governmental regulations, or compliance requirements that govern whether and how to store messaging data outside the organization's local area network (LAN)?
- Does your organization satisfy the client requirements for Exchange Online?
- Does your organization have a reliable connection to the Internet with sufficient bandwidth to move all mailboxes to the cloud?
- Does your organization have many mobile users or users who work outside the corporate offices and would benefit from a connection to the cloud rather than to the corporate data center?

MCT USE ONLY. STUDENT USE PROHIBITED

Discussion: Office 365 Experience and Expectations

Discuss the following questions related to your organization's Office 365 and Exchange Online experience.

- **Question:** Who is currently using or evaluating Office 365 or Exchange Online? What functionality of Office 365 are you using or evaluating?
- **Question:** If you are not currently using Office 365 with Exchange Online, why not?
- **Question:** What are your expectations or concerns about Office 365?



MCT USE ONLY. STUDENT USE PROHIBITED

Lesson 2

Planning and Implementing the Migration to Exchange Online

If you are planning to move from an Exchange on-premises deployment to Exchange Online, you must consider how to move the existing data, such as the user accounts and the mailbox content.

You use the same tools to manage Exchange Online users as you do to manage the on-premises users. This lesson describes your migration options and the tools you can use to manage the mailboxes both during and after migration.

Lesson Objectives

After completing this lesson, you will be able to:

- Plan migrations to Exchange Online.
- Migrate to Exchange Online.
- Describe Exchange Online management tools.

Planning Migrations to Exchange Online

Exchange Online offers various migration options and built-in tools to fit the migration needs of your organization. All migration options and built-in tools have the same management experience across tools and application programming interfaces (APIs) in Exchange Server 2010 and newer.

Migration	IMAP migration	• Supports a wide range of email platforms • Email only (no calendar, contacts, or desks)	
	Cutover Exchange migration	• Good for fast, direct migrations • Does not need a server on-premises	
Coexistence	Staged Exchange migration	• Supports staged migrations • Requires Directory Synchronization	
	Hybrid deployment	• You can manage users on-premises and online • Enables cross-premises calendaring, smooth migration, and easy off-boarding • Automatically redirects clients	

	IMAP migration	Cutover	Staged	Exchange 2010 hybrid	Exchange 2013 hybrid
Exchange Server 5.5	X				
Exchange 2000 Server	X				
Exchange Server 2003	X	X	X	X	
Exchange Server 2007	X	X	X	X	X
Exchange Server 2010	X	X	X	X	X
Exchange Server 2013	X	X	X	X	X
Lotus Notes/ Domino	X				
GroupWise/ Other	X				

IMAP Migration

The most common way to migrate from non-Microsoft messaging systems, such as Lotus Notes or GroupWise, to Exchange Online is to use IMAP migration. To use IMAP migration, do the following:

1. Ensure that, in the existing messaging system, you can access the mailboxes by using IMAP4.
2. Create a comma-separated values (.csv) file to list the users you want to migrate.
3. Use the EAC to migrate mailbox contents to the respective online mailboxes.

This migration option supports the widest range of email platforms, including Microsoft Exchange Server 5.5 and Microsoft Exchange 2000 Server.

This option has the following limitations:

- Only email messages migrate to the online mailbox. Calendar and contacts information does not migrate.
- Coexistence is not possible. You need to migrate all mailboxes at the same time to help ensure that you do not lose data.
- You can move at most 1,000 mailboxes at a time. The Office 365 Portal can read only .csv files that have a maximum of 1,000 rows per file. If you need to move more than 1,000 mailboxes, you must create additional .csv files, each containing a maximum of 1,000 mailboxes, and then import each file into Exchange Online.

MCT USE ONLY. STUDENT USE PROHIBITED

Cutover Exchange Migration

Cutover Exchange migration migrates all mailboxes from an Exchange on-premises installation to Exchange Online at the same time. This migration method does not support coexistence. You must migrate all mailboxes or none. For example, use cutover Exchange migration if you want to migrate all mailboxes in a short period of time, such as a few days or over a weekend.

Cutover Exchange migration uses Outlook Anywhere, a feature of Microsoft Exchange, to connect to the source mailboxes, and it copies all contents to the online mailboxes.

Cutover Exchange migration includes the following features:

- The migration service provisions new mailboxes in the cloud-based organization. It creates a cloud-based mailbox for each user account in the Exchange on-premises organization. It also synchronizes on-premises distribution groups and contacts to the cloud.
- After the migration service creates the new cloud-based mailboxes, it migrates all mailbox items, such as messages, contacts, and calendar items, from the Exchange Server on-premises mailboxes to the corresponding cloud-based mailboxes.
- After the initial migration, the Exchange Server and cloud-based mailboxes are synchronized every 24 hours. In the synchronization, new email messages that are sent to an Exchange Server on-premises mailbox are copied to the corresponding cloud-based mailbox. The synchronization is necessary until you finalize the migration process and change the Domain Name System (DNS) mail exchanger (MX) resource record so that all new messages go directly to the cloud-based mailbox.

You do not need any servers running Exchange Server 2013 on-premises to perform a cutover Exchange migration. However, similar to IMAP migration, you can use a cutover Exchange migration to migrate at most 1,000 mailboxes in total. If you have more than 1,000 mailboxes, you need to use a staged migration. New to the cutover migration is also that you can migrate Public Folders as well. However, this requires at least Exchange 2007 SP3 RU10 running in your environment.

Staged Exchange Migration

Staged Exchange migration is similar to cutover Exchange migration except that staged Exchange migration allows for some coexistence, which means that you can migrate mailboxes in stages. You can use staged Exchange migration if you cannot migrate quickly, if the organization requires a longer coexistence phase, or if a hybrid deployment is not an option.

Staged Exchange migration uses Outlook Anywhere for the connection, and it requires a .csv file. After a mailbox is migrated, Directory Synchronization updates the information, and the user is automatically reachable in Exchange Online at their original email address as well as in the Exchange on-premises environment through a mail-enabled user.

Staged Exchange migration is available for Exchange Server 2007 and newer. Before you use staged Exchange migration, you need to configure and install the Directory Synchronization tool.

Hybrid Deployment

Hybrid deployment is the smoothest migration method, and it has the lowest impact on the users. With this option, you use the EAC or the Exchange Management Shell to migrate users to or from Exchange Online. Hybrid deployment also provides full coexistence, so that users can exchange free/busy information or MailTips. No other migration option provides full coexistence. Starting with Exchange Server 2013, you can also move Public Folders between on-premises and Exchange Online environments in a hybrid deployment.

Use hybrid deployment if you require long-term coexistence or if you do not plan to move all mailboxes to Exchange Online. Also, hybrid deployment is the only option if you need to preserve Outlook .ost files on the client. If you preserve the Outlook .ost files when you move a mailbox from the on-premises

MCT USE ONLY. STUDENT USE PROHIBITED

environment to Exchange Online, a full .ost synchronization is no longer triggered when the user opens Outlook for the first time after the mailbox move.

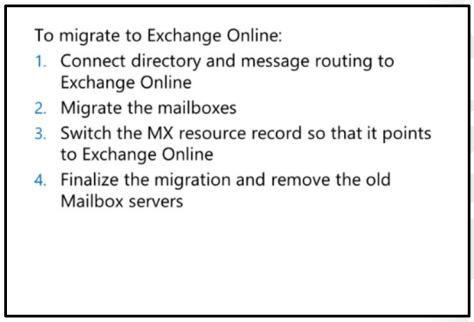
The principal benefit of hybrid deployment is that mailbox moves occur over the Internet by using the Mailbox Replication Service proxy. The Client Access servers that communicate between Exchange on-premises and Exchange Online perform the mailbox moves. You do not need to create .csv files. Also, in this approach the mailbox stays online during the move. You need to restart Outlook only when the move is complete.

To use this migration method, you must configure your Exchange Server organization for hybrid deployment, in order to have features such as free/busy information available for both on-premises mailboxes and cloud-based mailboxes. You also need at least one Exchange Server 2013 machine in your Exchange Server organization, and you need to configure Directory Synchronization and Exchange Federated Delegation. You can use the Hybrid Configuration Wizard to configure a hybrid deployment, which this module describes later.

Note: Exchange Server 2013 hybrid does not work with Exchange Server 2003. However, Exchange Server 2003 customers can deploy Exchange Server 2010 hybrid with Exchange Online in order to have a smoother experience migrating to the cloud if other options are not suitable for their business requirements.

Migrating to Exchange Online

Migrating an existing messaging system to Exchange Online is a complex task that includes many variables, such as the client protocol and messaging system that are used, as well as the size of your system. However, if you decide to migrate all of the mailboxes to Exchange Online, most migrations follow this general pattern:



1. Connect directory and message routing to Exchange Online. As mentioned previously, if you establish a connection to Exchange Online, your existing email directory synchronizes to Exchange Online. Depending on your source directory, you can use the Directory Synchronization tool to synchronize the Exchange Online directory, or, if you use legacy systems, you can use other tools such as a .csv file. Be sure that all existing Simple Mail Transfer Protocol (SMTP) addresses in the source messaging system exist in Exchange Online. Otherwise, you will lose messages when you configure the MX resource record after the migration. Also, verify that the existing messaging system can send messages to Exchange Online.
2. Migrate the mailboxes. Choose a migration method that uses either Microsoft tools or non-Microsoft tools. You can perform a staged migration, or you can migrate everything at the same time. Which option is better depends on the organization's size, the existing messaging environment, and other factors.
3. Switch the MX resource record so that it points to Exchange Online. By switching the resource record, you cause all inbound message traffic to flow directly to Exchange Online. After you make this change, you will no longer see many messages in your local messaging system. If messages continue to be sent to your local messaging system, investigate why the sending messaging system is not using the updated MX resource record.
4. Finalize the migration and remove the old Mailbox servers. Shut down everything in the on-premises messaging system. Check for the following:

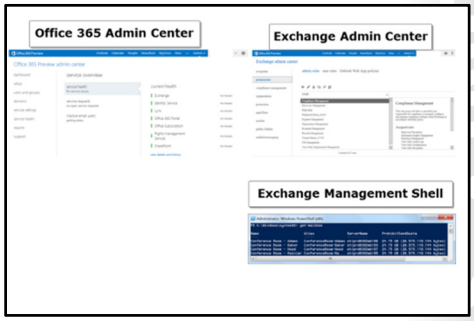
MCT USE ONLY. STUDENT USE PROHIBITED

- Any inbound or outbound messages flowing through the system.
- Any mailbox access after you switched over to Exchange Online.

5. After you shut down everything, you can remove the old mail servers from the data center and retire them.

Exchange Online Management Tools

Exchange Online provides several tools that you can use to manage your organization. You can choose between Office 365 Admin Center, Exchange Admin Center, and the Exchange Management Shell. Depending on your configuration, you can manage Exchange Online users by using the same tools that you use to manage Exchange on-premises users. The benefit of this type of configuration is that you do not need to consider where the mailboxes are hosted.



Office 365 Admin Center

You can use the Office 365 Admin Center to manage Office 365 accounts. You can perform tasks that are common across the Office 365 services within the Office 365 Admin Center, and you can follow links to the Exchange Admin Center, where you can manage settings that are specific to Exchange Online.

Generally, you use the Office 365 Admin Center to do the following:

- Provision new mailboxes and security groups.
- Manage common user properties.
- Create and manage service requests.
- Add and manage SMTP domains.

You can do the following only in the Office 365 Admin Center:

- Reset passwords.
- Manage service subscriptions.
- Assign licenses.

Exchange Admin Center

The Exchange Admin Center in Exchange Online is almost the same as the version of Exchange Admin Center that is available in an Exchange Server 2013 on-premises installation.

In Exchange Online, the Exchange Admin Center is the central management platform for creating and managing user mailboxes, distribution groups, and contacts. You also can configure organization-wide settings such as Unified Messaging IP gateways and Exchange ActiveSync access settings. The Exchange Admin Center has the following high-level categories:

- Recipients. Mailboxes, distribution groups, external contacts, and email migration.
- Permissions. Administrator roles, user roles, and Outlook Web App policies.
- Compliance management. Rules, journaling, eDiscovery, and delivery reports.
- Organization. Organization, sharing, and apps.

MCT USE ONLY. STUDENT USE PROHIBITED

- Protection. Filtering, spam, and quarantine.
- Mail flow. Rules, delivery reports, message trace, accepted domains, and connectors.
- Mobile. Mobile device access, and mobile device mailbox policies.
- Public folders. Public folders and public folder mailboxes.
- Unified messaging. Unified Messaging dialing plans and Unified Messaging gateways.

As in Exchange Server 2013, you can provide access to the Exchange Admin Center features by using role-based access control (RBAC).

Exchange Management Shell

You can use the Exchange Management Shell with remote Windows PowerShell to connect to Exchange Online. This way, you can perform management tasks by using cmdlets and scripts.

Exchange Online uses almost the same Windows PowerShell cmdlets as Exchange Server 2013 with Service Pack 1 (SP1) or newer. However, some cmdlets and parameters are disabled in Exchange Online because they do not apply in a data center environment.

MCT USE ONLY. STUDENT USE PROHIBITED

Lesson 3
Planning to Coexist with Exchange Online

If you plan to create a hybrid deployment by connecting your existing Exchange Server organization to Exchange Online, it is crucial that you plan the coexistence with Exchange Online.

This lesson focuses on a hybrid deployment and explains what you need to do to plan, implement, and manage a hybrid deployment.

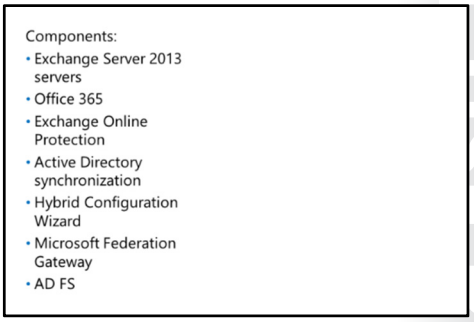
Lesson Objectives

After completing this lesson, you will be able to:

- Plan a hybrid deployment.
- Choose a mail flow option for a hybrid deployment.
- Use federated delegation for Exchange Online.
- Explain the hybrid deployment process.
- Implement a hybrid deployment.
- Use the Hybrid Configuration Wizard in Exchange Server 2013.
- Migrate an existing hybrid deployment to Exchange Server 2013.
- Manage a hybrid deployment.
- Use best practices for implementing a hybrid deployment.

Planning a Hybrid Deployment

To plan for a hybrid deployment, you first need to understand the components of a hybrid deployment and the prerequisites for configuring a hybrid deployment.



A hybrid deployment has the following components:

- Exchange Server 2013 servers. The on-premises Exchange Server organization requires at least the Exchange Server 2013 Client Access and Mailbox server roles.
- Office 365. The Office 365 service includes Exchange Online as a part of its subscription service. If you plan a hybrid deployment, you must create and configure a cloud-based Exchange Online organization.
- Exchange Online Protection. By default, the Exchange Online Protection (EOP) service is included in all Office 365 subscriptions for enterprise tenants. EOP works with on-premises Exchange Server 2013 Client Access servers to help secure message delivery between Exchange Server on-premises and Exchange Online. Depending on your configuration, EOP may also route incoming email from external recipients for Exchange Server on-premises and Exchange Online.
- Active Directory® synchronization. To support the unified GAL, synchronization of Active Directory directory services replicates information about mail-enabled objects from on-premises Active

MCT USE ONLY. STUDENT USE PROHIBITED

Directory to Office 365. You must deploy Active Directory synchronization on a separate, on-premises server before you can configure a hybrid deployment.

- Hybrid Configuration Wizard. Exchange Server 2013 includes the Hybrid Configuration Wizard, which you can use to configure a hybrid deployment between on-premises Exchange Server and Exchange Online.
- Microsoft Federation Gateway. The Microsoft Federation Gateway is a free, cloud-based service that acts as the trust broker between on-premises Exchange Server 2013 and Exchange Online. If you are configuring a hybrid deployment, you must have a federation trust with the Microsoft Federation Gateway.
- Active Directory Federation Services (AD FS). As an option, you can use AD FS to configure single sign-on (SSO) and centralized user management.

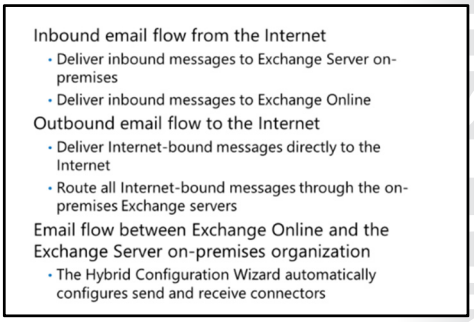
Before you can configure a hybrid deployment in Exchange Server 2013, either manually or by using the Hybrid Configuration Wizard, you must meet the following prerequisites:

- Exchange Server on-premises. You can configure a hybrid deployment for an on-premises organization that is based on Exchange Server 2007 or newer. For Exchange Server 2007 and Exchange Server 2010, you must install at least one Exchange Server 2013 Client Access and Mailbox server role in the on-premises organization. You must install the most recent service packs on all on-premises Exchange servers.
- Office 365 for enterprises. You must have an Office 365 for enterprises tenant administrator account and user licenses available on the tenant service. The version of the Office 365 tenant must be 15.0.000.0 or higher.
- Register custom domains. You must register any custom domains that you want to use in the hybrid deployment with Office 365. You can register the domains by using the Office 365 Portal.
- Active Directory synchronization. You must deploy the Directory Synchronization tool in the on-premises organization in order to synchronize Active Directory to Office 365.
- Autodiscover DNS records. You must configure the Autodiscover DNS records for your existing SMTP domains on the Internet to point to an on-premises Exchange Server 2013 Client Access server. For this reason, you need to install an Exchange Server 2013 Client Access server role in Exchange Server 2007 or Exchange Server 2010 environments.
- Trusted Digital Certificate. You must install and assign Exchange services to a valid digital certificate that you purchase from a trusted public certification authority (CA). The easiest way to verify that Exchange Online trusts your certificate is to run the Microsoft Remote Connectivity Analyzer against the Exchange Server on-premises environment. You cannot use self-signed certificates for Exchange services in a hybrid deployment.
- Office 365 organization in the Exchange Admin Center. Before you can configure the hybrid environment, you must connect the Exchange Admin Center to the Office 365 organization by using your Office 365 tenant administrator credentials. This way, you can manage both the on-premises and Exchange Online organizations from a single management console.
- Edge Synchronization (for Edge Transport). If the on-premises organization has Edge Transport servers and you want to configure the Edge Transport servers for hybrid secure mail transport, you must configure Edge Synchronization before you configure the hybrid environment. With Edge Synchronization, you can automatically configure the Edge Transport servers from the Exchange Admin Center.

MCT USE ONLY. STUDENT USE PROHIBITED

Mail Flow Options for a Hybrid Deployment

When you plan a hybrid deployment, you need to consider how to organize the email flow to and from the Internet, and between Exchange Online and the Exchange Server on-premises organization.



Inbound Email Flow from the Internet

Inbound email can be delivered to Exchange Server on-premises, or you can target the MX resource record to Exchange Online. You should choose a method depending on your organization's requirements. The following are some considerations for each scenario:

- Inbound to Exchange Server on-premises. Use this option if you want to keep full control of your email domains, such as message tracking and journaling of messages in your company. Configure the MX resource record to point to your organization's SMTP smart host. Then, the hybrid deployment automatically forwards all messages for mailboxes that are located on Exchange Online.
- Inbound to Exchange Online. Use this option if you want Microsoft to handle your email domains, and you want to automatically take advantage of the antivirus and anti-spam scanning engines from Exchange Online Protection. To configure this, point the MX resource record to Exchange Online Protection. Exchange Online automatically delivers messages that are located on Exchange Server on-premises by using the Exchange Server 2013 Hub Transport server that you define when you run the Hybrid Deployment Wizard.

Outbound Email Flow to the Internet

Similar to inbound email flow, you can configure outbound email flow from your domain to the Internet, by using one of the following options:

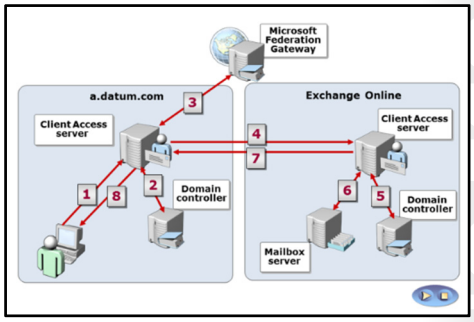
- Deliver Internet-bound messages directly. Use this option to send any outbound message that is targeted to the Internet directly from either Exchange Online or Exchange Server on-premises. If the mailbox is located on Exchange Online, the Internet messages are delivered directly to the target SMTP domain without passing through the Exchange Server on-premises environment. Messages sent from on-premises mailboxes are routed directly to Internet recipients without passing through Exchange Online. The benefit of this option is that the message traffic is optimized, but the drawback is that it is harder to track messages to the Internet because not every message flows through the on-premises Exchange servers.
- Route all Internet-bound messages through your on-premises Exchange servers. This option forces Exchange Online to send any message that is targeted to the Internet through the Exchange Server on-premises environment first. The Exchange servers then route the message to the Internet and deliver the message. The benefit of this option is that all messages pass through the Exchange servers, so you can use message tracking, journaling, and other compliance features.

Email Flow Between Exchange Online and the Exchange Server On-Premises Organization

Email flow between Exchange Online and an Exchange Server on-premises organization uses SMTP send and receive connectors that the Hybrid Configuration Wizard configures automatically. The connectors enforce the requirement that messages be encrypted by using the Transport Layer Security (TLS) protocol.

MCT USE ONLY STUDENT USE PROHIBITED

How Federated Delegation Works for Exchange Online



To share information between Exchange Server on-premises and Exchange Online, use federated delegation. If you configure federated delegation on both sides, users can view availability information and MailTips, and they can track messages by using delivery reports.

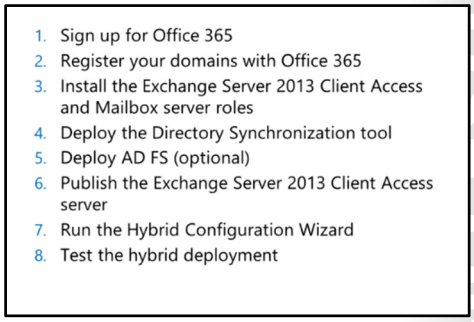
If federated delegation is implemented, the following steps describe the communication flow if an on-premises user invites a user who is hosted on Exchange Online to a meeting:

1. A user in adatum.com invites an Exchange Online user to a meeting. This meeting request is sent to the Exchange Web Service on the Client Access server at A. Datum.
2. The A. Datum Client Access server checks with an adatum.com domain controller to verify both that the user has permission to see availability information and that an organization relationship is configured with Exchange Online. If both verifications succeed, the Client Access server continues to the next step.
3. The A. Datum Client Access server connects to the Microsoft Federation Gateway and requests a security token for the A. Datum user. Because you configured adatum.com in the organization identifier, the Microsoft Federation Gateway issues the token.
4. The A. Datum Client Access server sends a request for the user's availability information to the Exchange Online Client Access server. The request uses the Autodiscover endpoint entry that is configured in the organization relationship to contact the remote server. The request also includes the security token.
5. The Exchange Online Client Access server validates the security token, and then the Client Access server checks with a domain controller in Exchange Online to verify that the organization has an organization relationship with adatum.com.
6. The Exchange Online Client Access server retrieves the user's availability information from the user's Mailbox server.
7. The Exchange Online Client Access server sends the availability information to the A. Datum Client Access server.
8. The A. Datum Client Access server provides the availability information to the A. Datum user.

MCT USE ONLY. STUDENT USE PROHIBITED

The Hybrid Deployment Process

If you are configuring a new hybrid deployment, you should plan to spend at least a couple of days on this process, depending on the complexity of your organization's deployment due to factors such as DNS replication and certificate verifications.



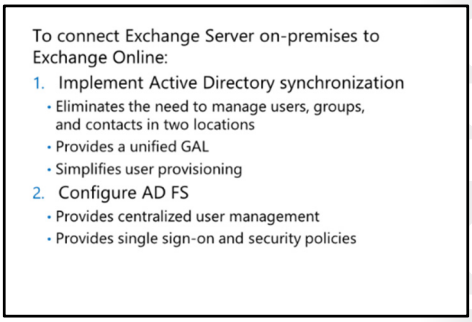
Follow these high level steps to configure a hybrid deployment:

1. Sign up for Office 365. Register your Office 365 enterprise tenant if you do not have one yet. If you already have an Office 365 tenant, make sure the tenant version is 15.0.000.0 or higher.
2. Register your domains with Office 365. Register with Office 365 the SMTP domains that you want to use for Exchange Online. Each domain needs to be verified with a DNS service (SRV) resource record, so this step might take a while.
3. Install the Exchange Server 2013 Client Access and Mailbox server roles. If your Exchange Server on-premises organization does not run on Exchange Server 2013, you need to deploy at least one server that runs Exchange Server 2013 with the Client Access and Mailbox server roles. The Exchange Server 2013 schema updates that happen during the server deployment are required to run the Directory Synchronization tool.
4. Deploy the Directory Synchronization tool. Activate directory synchronization in Office 365, and then deploy the Directory Synchronization tool. Activating directory synchronization may take up to a day to replicate the information throughout Office 365. Therefore, you should first install the Directory Synchronization tool after Office 365 is activated correctly.
5. Deploy AD FS. If you want, you can deploy AD FS for SSO. This is an optional step and not a requirement.
6. Publish the Exchange Server 2013 Client Access server. Make sure that the correct certificates are installed, that the Exchange Server 2013 Client Access server role is correctly published in the firewall, and that Autodiscover is working. The easiest way to verify the Autodiscover and Client Access server configuration from the Internet is to use the Microsoft Remote Connectivity Analyzer, which is available at http://go.microsoft.com/fwlink/?LinkId=290683.
7. Run the Hybrid Configuration Wizard. The Hybrid Configuration Wizard configures Exchange Server on-premises and Exchange Online for a hybrid deployment. Verify in the log files that all configurations are completed successfully. Additional information about the Hybrid Configuration Wizard is provided later in this lesson.
8. Test the hybrid deployment. Test the hybrid deployment by moving a non-productive mailbox to Exchange Online and then checking that free/busy information and MailTips are working as expected.

MCT USE ONLY. STUDENT USE PROHIBITED

Optimizing User Access to Exchange Online

You can run Exchange Online independently from the existing messaging infrastructure, but the functionality that is available in a hybrid deployment can help you manage mailbox migrations for on-premises and online users more efficiently.



Implement Active Directory Synchronization

After you implement AD FS, you should also implement Active Directory synchronization between your organization's Active Directory forest and Exchange Online. To do this, use the Directory Synchronization tool.

The Directory Synchronization tool simplifies management by synchronizing the local Active Directory forest with Exchange Online. As a result, you do not need to administer the organization's objects from two locations.

The Directory Synchronization tool updates the Microsoft online environment whenever changes occur in Active Directory Domain Services (AD DS). This means that changes such as adding a new employee, deleting an employee, and changing contact information automatically propagate to Exchange Online, so you do not need to update Exchange Online manually. These synchronized items are read-only in Exchange Online, and you continue to manage them with the AD DS tools.

The Directory Synchronization tool synchronizes changes every three hours. To help protect your security, the tool does not update sensitive information such as domain passwords. The tool also updates distribution groups and the GAL, and it plays an important role during coexistence between your on-premises organization and Exchange Online.

Configure AD FS

You can configure AD FS to allow SSO and centralized user management. You do not need to configure AD FS, but we recommend using it to improve user satisfaction.

With AD FS, users can access online services with the same domain credentials that they use to access on-premises applications through the process of SSO. There is no need for a client-side sign-in tool.

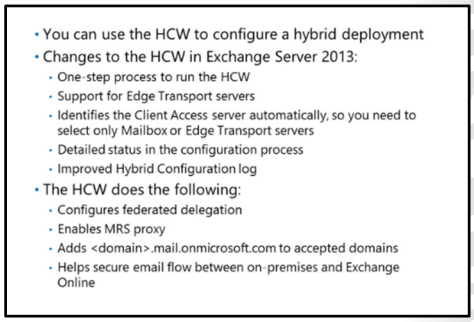
AD FS provides the following benefits:

- Improved manageability and lower total cost of ownership (TCO).
- Passwords stay within the organization. Microsoft does not see credentials and passwords, because they are not synchronized to the cloud.
- Organizations retain security control over user accounts and password expiration.
- Configuration and management are simpler. AD FS does not require changes to the Active Directory code or alterations of the enterprise's Active Directory deployment.

With AD FS, you can deploy a multi-factor authentication system, which can include soft certificate and smartcard authentication from out-of-the-box products such as RSA and Swivel. You can customize the login page for Exchange Online and for other federated web applications, such as SharePoint Online.

MCT USE ONLY. STUDENT USE PROHIBITED

The Hybrid Configuration Wizard in Exchange Server 2013



Exchange Server 2010 with Service Pack 2 (SP2) introduced the Hybrid Configuration Wizard to make it easier to configure a hybrid deployment for Exchange Server on-premises. The Hybrid Configuration Wizard provides a wizard and cmdlets that you can use to establish and manage Exchange Server on-premises and Exchange Online in a hybrid deployment, by using the Exchange Admin Center.

Exchange Server 2013 includes the following changes to the Hybrid Configuration Wizard:

- The configuration is one step in Exchange Server 2013. One wizard creates the HybridConfiguration Active Directory object and configures the hybrid deployment properties and services. In Exchange Server 2010, configuration is two steps.
- The Hybrid Configuration Wizard automatically selects the Client Access server. You need to select only the Mailbox or Edge Transport servers to configure the hybrid deployment email flow.
- You can configure the Edge Transport servers in the Hybrid Configuration Wizard.
- The Hybrid Configuration Wizard shows a detailed status during the configuration process.
- The Hybrid Configuration log is improved, and it separates each hybrid configuration step into a clearly delineated section. This improvement simplifies review and troubleshooting. The new log identifies where each hybrid configuration task is performed, either in the on-premises Exchange Server organization or in Exchange Online.

Before you run the Hybrid Configuration Wizard, you must satisfy all prerequisites, such as setting up Active Directory synchronization between AD DS and Exchange Online, as explained in the Planning a Hybrid Deployment topic.

The Hybrid Configuration Wizard does the following:

- Federated delegation. The wizard checks to see whether a federation trust exists with the Microsoft Federation Gateway for your organization. If the trust exists, it is used to support the hybrid deployment. If the trust does not exist, the wizard creates it and adds to it the domains that you select.
- Enables the Mailbox Replication Service proxy. The wizard enables the Mailbox Replication Service proxy on all Client Access servers that you select. This enables mailbox moves from Exchange Server on-premises to Exchange Online and vice versa.
- Adds <domain>.mail.onmicrosoft.com to accepted domains. The wizard adds a coexistence domain to the accepted domains list of the on-premises organization. By default, this domain is <domain>.mail.onmicrosoft.com. This coexistence domain is used for email flow between the on-premises organization and the Exchange Online tenant, and the domain is added as a secondary proxy domain to any email address policy of your organization.
- Helps secure email flow between on-premises and Exchange Online. The wizard configures selected Hub Transport servers and EOP in Office 365 to help secure email routing. The wizard creates or updates existing Send and Receive connectors in the on-premises organization and Inbound and Outbound connectors in EOP. The wizard prompts you to decide whether you want the Exchange Online tenant to send the messages directly to the Internet or to forward all external messages to the on-premises environment first before routing them outside the organization.

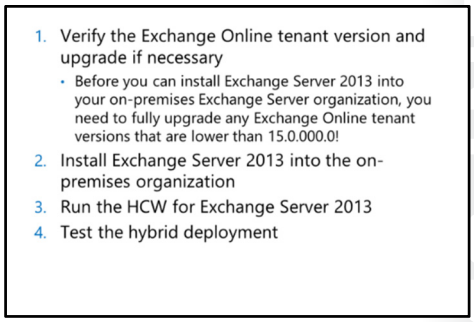
MCT USE ONLY. STUDENT USE PROHIBITED

You can use the Hybrid Configuration Wizard to manage the following:

- Free/busy sharing. You can allow on-premises users and Exchange Online users to view free/busy information.
- Mailbox moves. You can move mailboxes from Exchange Server on-premises to Exchange Online and from Exchange Online to Exchange Server on-premises. You can also preserve the users' Outlook profiles and offline .ost files.
- Message tracking. You can use delivery reports to track messages between Exchange Server on-premises and Exchange Online.
- MailTips. You can allow users to retrieve information while they are composing a message, such as an Out-of-Office notification.
- Online archiving. You can store personal archives in the Exchange Online tenant.
- Outlook Web App redirection. You can use this feature to provide a single URL to users when you move their mailbox from on-premises to Exchange Online.
- Secure email. You can help secure message delivery between the on-premises and cloud organizations by using the TLS protocol. All messages that are transferred between the on-premises organization and Exchange Online are encrypted and transferred directly, without any other server involvement.

Migrating an Existing Hybrid Deployment to Exchange Server 2013

If you already have an existing hybrid deployment with Exchange Server 2010 on-premises, you need to consider additional requirements in order to migrate to Exchange Server 2013.



To migrate an existing hybrid deployment to Exchange Server 2013, follow these steps:

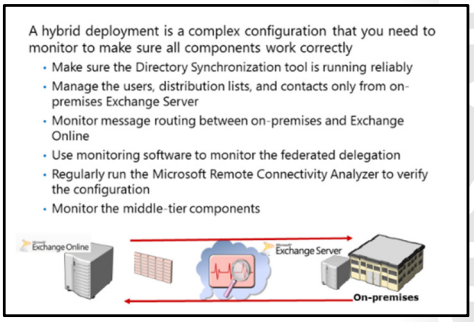
1. Verify the versions of the Exchange Online tenant, and upgrade if necessary. Exchange Server 2013 is supported only if it is running against a tenant version of 15.0.000.0 or higher. To verify your current Exchange Online tenant version, connect Windows PowerShell to Exchange Online, and then run the following cmdlet: **Get-OrganizationConfig | Format-List AdminDisplayVersion,IsUpgradingOrganization**. If the tenant version is not 15.0.000.0 or higher, create a service request with Exchange Online to upgrade your tenant version.
2. Install Exchange Server 2013 into the on-premises organization. After the Exchange Online tenant is upgraded, you can install Exchange Server 2013 into the on-premises environment. Set up at least one Client Access and Mailbox server role to take over the communication to Exchange Online. Then, move the Autodiscover service to the Exchange Server 2013 Client Access server role, and make sure the Internet firewall correctly publishes the Client Access server. As usual, verify the functionality by using the Microsoft Remote Connectivity Analyzer.
3. Run the Hybrid Configuration Wizard for Exchange Server 2013. Run the Hybrid Configuration Wizard to update the existing Hybrid Configuration Wizard configuration and to change communication to the Exchange Server 2013 Client Access server.

MCT USE ONLY. STUDENT USE PROHIBITED

4. Test the hybrid deployment. To test the new hybrid deployment, create a new mailbox, move it to Exchange Online, and then make sure the client features such as free/busy information and MailTips are working correctly.

Note: Before you can install Exchange Server 2013 into your on-premises Exchange Server organization, you need to fully upgrade any Exchange Online tenant versions that are lower than 15.0.000.0.

Managing a Hybrid Deployment



A hybrid deployment includes some special management and monitoring requirements, due to its complex configuration that automatically exchanges information between the Exchange Server on-premises organization and Exchange Online. The complexity of this configuration is not visible to end users, but you need to monitor specific areas of the deployment to ensure that it functions properly.

To monitor a hybrid deployment, do the following:

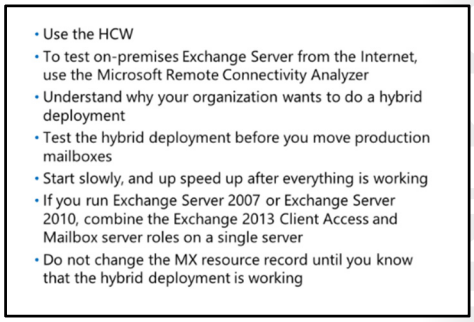
- Make sure the Directory Synchronization tool is running reliably. The Directory Synchronization tool is the essential tool that a hybrid deployment needs to synchronize the Exchange Server on-premises environment with Exchange Online. For example, if you configure a personal archive for a mailbox that is stored in Exchange Online, the Directory Synchronization tool synchronizes the properties of the mailbox so that Exchange Online recognizes the archive. If Directory Synchronization is not running, Exchange Online does not recognize the change, and the user cannot use their archive. Office 365 automatically monitors your Directory Synchronization activity, and it sends a message to the technical account if Directory Synchronization does not occur for a day.
- Use the Exchange Admin Center of the on-premises Exchange Server environment to manage Exchange Server 2013 on-premises, the Exchange Online tenant, the hybrid settings, and the mailbox migrations so that Directory Synchronization synchronizes them correctly. If you use the Exchange Admin Center to synchronize users, distribution lists, and contacts, keep in mind that synchronization occurs in one direction only—from the Exchange Server on-premises organization to Exchange Online. For example, if you create an on-premises user mailbox, Directory Synchronization creates the user mailbox in Exchange Online. But, if you create a user mailbox in Exchange Online, Directory Synchronization does not synchronize or create the user mailbox in AD DS.
- Monitor message routing between on-premises and Exchange Online. Message routing between Exchange Server on-premises and Exchange Online is one of the most important factors that makes a hybrid deployment successful. Make sure that the messages flow successfully and do not queue somewhere. For this reason, we recommend that you monitor the queues in the Exchange Server on-premises environment so that you can react quickly if messages queue for too long.
- Use monitoring software to monitor the federated delegation. Federated delegation is the basis for the information exchange between Exchange Server on-premises and Exchange Online. If federated delegation does not work correctly, users cannot retrieve any free/busy information, MailTips, or other information between the on-premises and cloud deployments. Consider testing federated delegation with the monitoring software, so you are notified immediately if federated delegation does not work. Also consider using the following test cmdlets:

MCT USE ONLY. STUDENT USE PROHIBITED

- **Test-FederationTrust**
- **Test-FederationTrustCertificate**
- **Test-OrganizationRelationship**
- Regularly run the Microsoft Remote Connectivity Analyzer to verify the configuration. The Microsoft Remote Connectivity Analyzer is a Microsoft tool that can verify your configuration, such as the Exchange Web Services or Exchange ActiveSync settings, and ensure that all settings are configured properly. This tool helps prevent issues that you did not find previously. Because a hybrid deployment uses those services to communicate between Exchange Online and on-premises, we recommend that you occasionally run these tests in order to verify that the configuration did not change in any way.
- Monitor the middle-tier components. A hybrid deployment involves not only Exchange servers, but also other components, such as firewalls, so you must ensure that these components function correctly. Therefore, consider monitoring any middle-tier component that is involved in the deployment. These components include Microsoft Forefront® Threat Management Gateway, AD FS, and other products.

Best Practices for Implementing a Hybrid Deployment

When implementing a hybrid deployment, remember the following best practices:



- Use the Hybrid Configuration Wizard. You configure a hybrid deployment in two ways: either manually or by using the Hybrid Configuration Wizard. We recommend that you use the Hybrid Configuration Wizard, because the wizard handles all necessary configuration settings for you.
- To test the Exchange Server on-premises environment from the Internet, use the Microsoft Remote Connectivity Analyzer, which is available at http://go.microsoft.com/fwlink/?LinkId=290683.
- Understand why your organization wants to implement a hybrid deployment. Do not try to use "everything" if your organization only wants to move archive mailboxes to Exchange Online.
- Test the hybrid deployment before you move production mailboxes. Always move test mailboxes first, and then consider moving production mailboxes.
- Start slowly, and then speed up when everything works. At the beginning, move mailboxes only for people who can live with a short outage. After you gain confidence that the hybrid deployment works reliably, move the other mailboxes.
- If you run Exchange Server 2007 or Exchange Server 2010, combine the Exchange Server 2013 Client Access and Mailbox server roles on a single server.
- Do not change the MX resource record at first. Change it only after you know that the hybrid deployment works.

MCT USE ONLY. STUDENT USE PROHIBITED

Lab: Designing Integration with Exchange Online

Scenario

A. Datum is considering moving some mailboxes to Exchange Online. A. Datum wants to ensure not only that the users can still use their internal adatum.com domain accounts to authenticate, but also that they can access their mailboxes online. The mailbox location should be transparent to users when they access their mailbox, send email messages, or book meetings.

Objectives

Students will be able to design coexistence with Exchange Online.

Lab Setup

Estimated Time: 45 minutes

This lab does not require any virtual machines.

Exercise 1: Designing Integration with Microsoft Exchange Online

Scenario

A. Datum currently has the following Exchange Server 2013 servers deployed:

Berlin site:

- BER-CAS1 (Client Access role)
- BER-MBX1 (Mailbox role)

London site:

- LON-CAS1 (Client Access role)
- LON-MBX1 (Mailbox role)

Additionally, the following information is available:

- The London site is connected to the Internet and uses Forefront Threat Management Gateway as a firewall solution.
- Autodiscover from the Internet is currently not configured and not working.
- The Exchange Server organization has a non-Microsoft email relay to send and receive messages to and from the Internet, and it does not use Exchange Server Edge Transport servers.
- You have a subscription to a test implementation of Exchange Online. The Exchange Online tenant has a version of 14.16.190.13.

Your manager asks you not only to evaluate Exchange Online on its own, but also to connect the existing Exchange Server organization with Exchange Online. You have the following requirements:

- Free/busy information and MailTips must be available on both sides.
- Mailbox moves must be done in either the Exchange Admin Center or in Windows PowerShell.
- All new and existing user mailboxes must be managed from A. Datum's internal AD DS domain, adatum.com.
- The user must sign in only once, when they start their computer.
- The user must not store their passwords in Exchange Online or anywhere that Microsoft can access.

Design a solution that enables A. Datum to move some mailboxes to Exchange Online while the users can still use their domain accounts.

MCT USE ONLY. STUDENT USE PROHIBITED

The main tasks for this exercise are as follows:

1. Read and analyze the scenario requirements
2. Design a solution
3. Discuss your solution with the class

▶ Task 1: Read and analyze the scenario requirements

1. Read the exercise scenario, and then analyze the requirements from an integration perspective. Identify the configurations needed to satisfy the requirements.

▶ Task 2: Design a solution

Propose a solution that satisfies all of the requirements. Use the following questions as a guideline:

- What components do you need to install and configure in order to satisfy the requirements?
- What existing Exchange Server 2013 server can you use for message transport and Autodiscover? What additional areas do you need to plan in order to run the Hybrid Configuration Wizard successfully?
- What would be different if A. Datum were running only Exchange Server 2010, and not Exchange Server 2013?

▶ Task 3: Discuss your solution with the class

1. Present your proposed solution to the class. Discuss alternative solutions with other students and with the instructor.

Results: After completing this exercise, you will have successfully:

- Read and analyzed the scenario requirements.
- Designed a solution.
- Discussed your solution with the class.

Question: Before you can run the Hybrid Configuration Wizard in the Exchange Admin Center, what do you need to do?

Question: You run Exchange Server 2010 in a hybrid deployment. Your current Exchange Online tenant is version 14.16.190.13. What do you need to do before you can install Exchange Server 2013?

MCT USE ONLY. STUDENT USE PROHIBITED

Module Review and Takeaways

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
Free/busy information is not available for mailboxes in Exchange Online when accessing the information about mailboxes that are stored on-premises.	
Moving mailboxes between Exchange Online and on-premises fails.	

Review Question(s)

Question: How can you deploy Exchange Online?

Question: What additional functionality does Exchange Online Protection provide?

Question: What functionality does federated delegation provide?

Question: When planning a hybrid deployment for your organization, what components do you need to consider implementing, and what is their purpose?

Question: You created a new mailbox in Exchange Online, and now the on-premises users complain that they cannot see the new mailbox. What can you do?

Tools

You can use the following tools to monitor and test a hybrid deployment.

Tool	Use for	Where to find it
Microsoft Remote Connectivity Analyzer	Troubleshooting your on-premises SSO, Exchange ActiveSync®, or Exchange Web Service.	http://go.microsoft.com/fwlink/?LinkId=290683
Test-FederationTrust cmdlet	Make sure the federation trust is working correctly.	Exchange Management Shell
Start-OnlineCoexistenceSync cmdlet	Start a manual directory synchronization to immediately synchronize AD DS with Office 365.	Microsoft Online Directory Sync Shell

MCT USE ONLY. STUDENT USE PROHIBITED

MCT USE ONLY. STUDENT USE PROHIBITED

Module 11

Designing and Implementing Messaging Coexistence

Contents:

Module Overview 11-1

Lesson 1: Designing and Implementing Federation 11-2

Lesson 2: Designing Coexistence Between Exchange Server Organizations 11-8

Lesson 3: Designing and Implementing Cross-Forest Mailbox Moves 11-15

Lab: Implementing Messaging Coexistence 11-23

Module Review and Takeaways 11-30

Module Overview

Microsoft® Exchange Server 2013 provides options to integrate with other messaging systems, with other organizations that are using Exchange Server, and with Microsoft Exchange Online. You can achieve integration by using coexistence and by using federation. If you integrate with federated partners that are also using Exchange Server 2013, you can share information with partner organizations. If you integrate with Exchange Online, you can expand the messaging system in your organization without adding additional servers.

Objectives

After completing this module, you will be able to:

- Design and implement federation.
- Design coexistence between Exchange organizations.
- Design and implement cross-forest mailbox moving.

MCT USE ONLY. STUDENT USE PROHIBITED

Lesson 1
Designing and Implementing Federation

If you integrate with federated partners, you can share calendaring information and contacts between organizations. To configure federated partners, you must know how to create a federated trust, and then you must implement an organization relationship or a sharing policy.

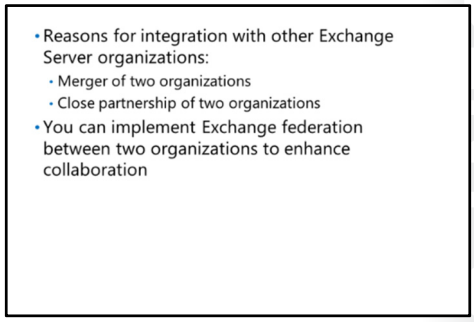
Lesson Objectives

After completing this lesson, you will be able to:

- Describe scenarios for integrating with other Exchange Server organizations.
- Describe federation.
- Describe federation deployment components.
- Describe considerations for designing federated trusts and certificates.
- Design organization relationships.
- Design sharing policies.

Scenarios for Integrating with Other Exchange Server Organizations

After a merger of organizations, it is common for multiple Exchange Server installations to coexist. Under these circumstances, coexistence is typically temporary, lasting only until the two messaging systems can be merged.



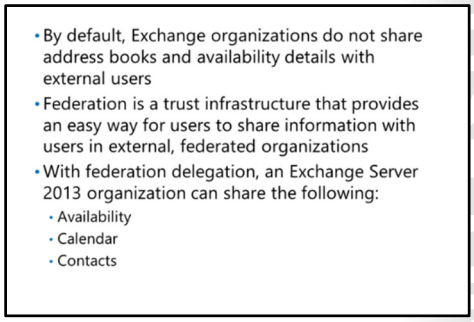
During an upgrade from Exchange Server 2007 or Exchange Server 2010 to Exchange Server 2013, a single global address list (GAL) is maintained, and the calendar data is shared between the Exchange Server versions. However, if two separate Exchange Server organizations coexist, the calendar data and GALs are not automatically synchronized between the two organizations. As a result, collaboration between the two organizations can be difficult. You can implement Exchange Server federation to make this type of collaboration easier and more functional.

You can configure message delivery between the Exchange Server organizations by using Send connectors. Send connectors enable you to apply specific configuration settings to messages that are being transferred between the organizations and to provide additional security to Simple Mail Transfer Protocol (SMTP) traffic.

Integration with Microsoft Office 365™ also requires coexistence with another Exchange Server organization. However, in the case of Office 365, coexistence is often long term. Office 365 also includes tools for directory synchronization.

MCT USE ONLY. STUDENT USE PROHIBITED

What Is Federation?



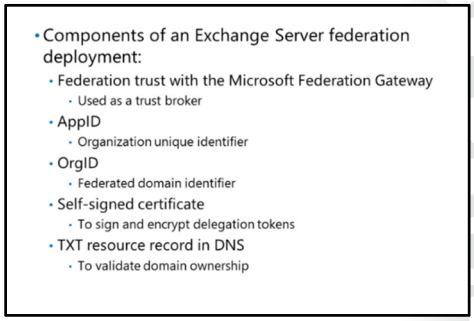
Users in one company often need to collaborate intensively with users in another company who are on external networks, such as vendors, partners, or customers. In addition, users also often need to share their address lists, availability information, and calendar with the external partners. By default, sharing address lists or availability data with users who are outside an Exchange Server organization is not possible. However, you can use federation technology in Exchange Server 2010 or Exchange Server 2013 to enable collaboration with users who are in another Exchange Server organization.

Federation refers to the underlying trust infrastructure that supports federated sharing. Federated sharing is an easy way for users to share calendar and contact information with people in external, federated organizations. With federation, both organizations want authentication assertions from one organization to be recognized by the other.

You can use federated delegation to configure your Exchange Server 2010 or Exchange Server 2013 organization to share information with other Exchange Server 2010 or Exchange Server 2013 organizations. This shared information can include availability information, calendar information, and contacts.

To configure federated delegation, you must create a federation trust for your organization, and then you must configure organization relationships or sharing policies. This process is much simpler than other ways to share information between organizations. However, this method does not synchronize all GAL information. It shares only user contacts. To participate in federated delegation, user mailboxes must be on an Exchange Server 2013 Mailbox server. Organization relationships or sharing policies define the information that is shared.

Federation Deployment Components



A federation scenario between two Exchange Server organizations has several components. The most important component is the Microsoft Federation Gateway. This is a cloud-based service, offered for free by Microsoft, and it acts as a trust broker between two federated Exchange Server organizations. Instead of establishing a direct trust, like with forest trusts in Active Directory® Domain Services (AD DS), you establish a trust in an Exchange Server federation scenario by using Microsoft Federation Gateway. Before starting to share data with the other Exchange Server organization, each organization must establish a trust relationship with Microsoft Federation Gateway.

After you configure a trust with Microsoft Federation Gateway, the Microsoft Federation Gateway service issues a Security Assertion Markup Language (SAML) delegation token to each user who AD DS authenticates. This token enables the authenticated user to access shared resources within the federated Exchange Server organization. With the Microsoft Federation Gateway acting as the trust broker,

MCT USE ONLY. STUDENT USE PROHIBITED

organizations do not need to establish multiple, individual trust relationships with other organizations. In addition, users can access external resources with their AD DS credentials by using single sign-on (SSO).

When you establish a trust with Microsoft Federation Gateway, your organization exchanges digital certificates with the Microsoft Federation Gateway certificate as well as the federation metadata. To establish this trust, you can use the Exchange admin center or the Exchange Management Shell. Either start a wizard to create the trust in the Exchange admin center or run the **New-FederationTrust** cmdlet in the Exchange Management Shell. A self-signed certificate is created on Exchange Server. This certificate is used to sign and encrypt delegation tokens from the Microsoft Federation Gateway that allow users to be trusted by external federated organizations.

To enable federation with Microsoft Federation Gateway in the Exchange admin center, navigate to the organization node, and then, on the **sharing** tab, click **enable federation**.

When you create a federation trust with Microsoft Federation Gateway, an object call application identifier (AppID) is also automatically created. You can reach this object by executing the **Get-FederationTrust** cmdlet. AppID uniquely identifies your Exchange Server organization on the Microsoft Federation Gateway side when establishing relationships with another Exchange Server organization. Another purpose of AppID is to provide valid proof that an organization is the owner of the domain that is being used for federation. This proof of ownership is achieved by creating a text (TXT) resource record with the AppID in the public Domain Name System (DNS) zone for each federated domain. To get the content of the TXT record, execute the following cmdlet:

```
Get-FederatedDomainProof –domainname <NameOfDomain>
```

This cmdlet returns the content of the TXT resource record that you should place in DNS, in the field **DnsRecord**. After that, you can use DNS Manager to create a TXT resource record in your public DNS that contains content for the domain proof.

When configuring federation, you must also define which of the authoritative accepted domains is used and enabled for federation. This parameter is defined by a federated organization identifier (OrgID). It is important to define this parameter because only users who have email addresses on the domain that is configured in OrgID can use features that the Exchange Server federation provides. OrgID is a combination of a pre-defined string and the accepted domain. The domain that is in OrgID is the domain that is selected as the primary shared domain in the wizard for creating a federated trust.

For example, if you specify the federated domain adatum.com as the primary shared domain in your organization, the FYDIBOHF25SPDLT.adatum.com account namespace is automatically created as the OrgID for the federation trust for your Exchange Server organization. You can set OrgID by executing the following cmdlet:

```
Set-FederatedOrganizationIdentifier –Enabled $true
```

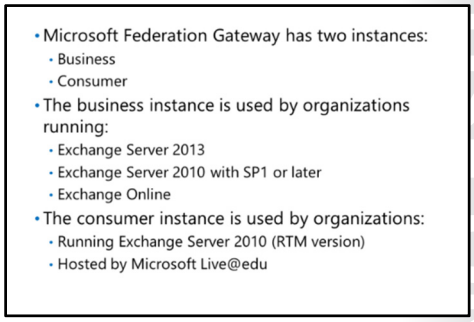
Note: The name of the accepted domain that you select to federate can have a maximum of 32 characters.

To enable or disable all federation sharing features in your organization, all you need to do is enable or disable the OrgID for the federation trust.

Note: If you change the OrgID, the accepted domains, or the AppID used for the federation trust, all federation sharing features are affected, both in your organization and in all federated organizations.

MCT USE ONLY. STUDENT USE PROHIBITED

Considerations for Designing Federated Trusts and Certificates



To perform federated delegation, you need to configure a federation trust with Microsoft Federation Gateway, which is a central point for federation trusts between Exchange Server organizations. Two instances of the Microsoft Federation Gateway are available: consumer and business. All federated partners must use the same instance of the Microsoft Federation Gateway.

The following types of Exchange Server organizations use the business instance of the Microsoft Federation Gateway by default:

- Exchange Server 2013 organizations, by using the Enable Federation Trust Wizard and self-signed certificates for a federation trust.
- Organizations that use Exchange Server 2010 with Service Pack 1 (SP1) or later, by using the New Federation Trust Wizard and self-signed certificates for a federation trust.
- Exchange Server organizations that are hosted by Microsoft Online Services, such as the Exchange Online service in Office 365.

The following types of Exchange Server organizations use the consumer instance of the Microsoft Federation Gateway by default:

- Organizations that use the Release to Manufacturing (RTM) version of Exchange Server 2010, with certificates that third-party certification authorities (CAs) issue.
- Exchange Server organizations that are hosted by Microsoft Live@edu.

We recommend that all Exchange Server 2013 organizations use the business instance of the Microsoft Federation Gateway.

Before you start to configure federation between the two Exchange Server organizations, you need to verify which Microsoft Federation Gateway instance each Exchange Server organization is using, to determine whether a federation trust already exists. Run the following cmdlet:

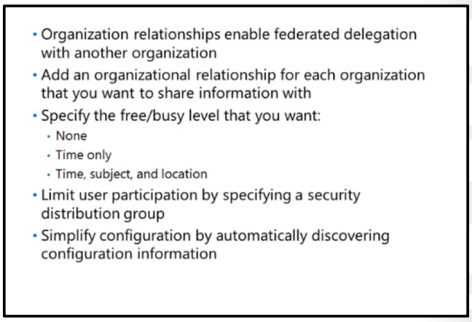
```
Get-FederationInformation -DomainName <hosted Exchange domain namespace>
```

The business instance returns a value of **<uri:federation:MicrosoftOnline>** for the **TokenIssuerURIs** parameter. The consumer instance returns a value of **<uri:WindowsLiveID>** for the **TokenIssuerURIs** parameter.

To establish a federation trust with the Microsoft Federation Gateway, you should use a self-signed certificate or an X.509 certificate signed by a CA. The certificate must be created and installed on the Exchange Server 2013 server that was used to create the trust. We recommend that you use a self-signed certificate, which is automatically created and installed when you use the Enable Federation Trust Wizard in the Exchange Administration Center (EAC). This certificate is used only to sign and encrypt delegation tokens that are used for federated sharing, and not to identify your organization. Therefore, the subject name in the certificate is not relevant. Exchange Server 2013 automatically distributes the certificate that is created by the Enable Federation Trust Wizard to all other Exchange Server 2013 servers in the organization. You can also choose to use an external x.509 certificate for federation.

MCT USE ONLY. STUDENT USE PROHIBITED

Configuring Organization Relationships



You can use organization relationships to enable federated delegation with an external Exchange Server organization. The external Exchange Server organization must also have a federation trust in place with the Microsoft Federation Gateway. Each organization relationship is for a single external organization, which is identified by its domain name and application identity.

When you create an organization relationship in the Exchange admin center, you configure the following options:

- Relationship name. This can be any name you choose.
- Domains to share with. Type the fully qualified domain name (FQDN) of the domain that you want to establish federation with.
- Enable calendar free/busy information sharing. With this option, you turn on information sharing. If you enable this option, choose one of the following options:
  - **Calendar free/busy information with time only**
  - **Calendar free/busy information with time, subject and location**
  - **Share calendar free/busy information for**
    - **Everyone in your organization**
    - **A specified security group**

When you configure the organization relationship, Microsoft Federated Gateway service checks your DNS zone and searches for the appropriate TXT resource record with the domain proof content. Create this record before you create the organization relationship.

Note: Even if an organization relationship specifies that all user calendars are shared, users can override this setting. Users can configure the default permissions for their own calendars to prevent sharing. However, changing the default permission also affects sharing with internal users.

To identify the external organization that you want to create the organization relationship with, you typically use the domain name of the external organization to automatically populate the necessary information into the organization relationship. If you specify the domain name, all of the necessary configuration information is obtained from the Microsoft Federation Gateway.

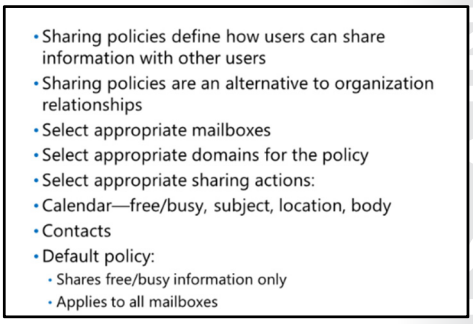
If you use the Exchange Management Shell to create the organization relationship, use the **Get-FederationInformation** cmdlet to obtain the federation information for the external organization. You can pipe this information to the **New-OrganizationRelationship** cmdlet when you create the organization relationship.

You can obtain the URL for the Availability Web Service of the external organization by using Autodiscover. If the external organization does not have Autodiscover configured for access from the Internet, enter the URL manually.

MCT USE ONLY. STUDENT USE PROHIBITED

Sharing of availability information performs best if users are using Microsoft Outlook® 2010, Office Outlook 2013, or Microsoft Outlook Web App on an Exchange Server 2013 Client Access server. Users of Outlook 2007 can view availability information for external users, but the users must be selected from the GAL, which means that GAL synchronization must be in place. Users who have mailboxes on Exchange Server 2007 with Service Pack 2 (SP2) can use Microsoft Office Outlook Web Access to view availability information for external users.

Designing Sharing Policies



For controlling federated delegation, sharing policies are an alternative to organization relationships. You can assign a sharing policy to specific mailboxes, or you can use a sharing policy to determine what information a user can share with external users. Instead of information being automatically available for users in an external organization, the user in your organization sends a sharing invitation to the external user to share the calendar or contacts. Although the organization containing the external user's mailbox does not need to have a federation trust, you should configure a federation trust to enable a two-way sharing relationship.

When you create a sharing policy, you can control the calendar information that your organization shares on a per-domain basis. You can choose whether to allow sharing of only free/busy information, or you can include the subject and location, or you can include the body. You also have the option to allow the sharing of contacts.

For a sharing policy to take effect, you must apply it to mailboxes. You can do this by using the properties of the sharing policy or the properties of the recipient. You can apply only a single sharing policy in each mailbox.

After installation, a sharing policy, called the Default Sharing Policy, is created. This policy automatically applies to all Exchange Server 2013 mailboxes, and it allows the sharing of free/busy information with all domains. The Default Sharing Policy enables users to share their free/busy information with external users immediately after a federation trust is created.

You can use the Exchange admin center or the Exchange Management Shell to create sharing policies and to assign them to specific mailboxes.

Only Outlook 2010 or newer and Outlook Web App can create sharing invitations. In addition, an Exchange Server 2013 Mailbox server must host the user mailbox.

MCT USE ONLY STUDENT USE PROHIBITED

Lesson 2

Designing Coexistence Between Exchange Server Organizations

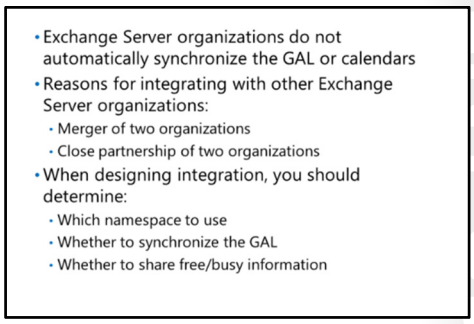
When you upgrade from a legacy messaging system to Exchange Server 2013, you might need the legacy messaging system and Exchange Server 2013 to coexist. You can use several configurations to accomplish this goal. When you plan the coexistence of the two messaging systems, you must consider several factors, such as message routing, address list synchronization, and calendar interoperability.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe multi-forest Exchange Server 2013 deployments.
- Design message routing.
- Design GAL synchronization.
- Design calendar interoperability.
- Design administration between Exchange Server organizations.

Multi-Forest Exchange Server 2013 Deployments



Each AD DS forest can support only a single Exchange Server organization. But, if you upgrade an Exchange Server organization from a previous version to Exchange Server 2013, both versions exist in the same organization at the same time. Exchange Server 2013 supports interoperability between a previous version of Exchange Server and Exchange Server 2013.

Typically, you need to plan for coexistence and integration with a second Exchange Server organization after two organizations merge. For example, say a large company that has an Exchange Server messaging system buys another company that also has an Exchange Server messaging system. Until the two messaging systems are merged, the two Exchange Server organizations need to coexist. Integration with another Exchange Server organization can also occur between partner organizations.

If you integrate two Exchange Server organizations, you need to determine the following:

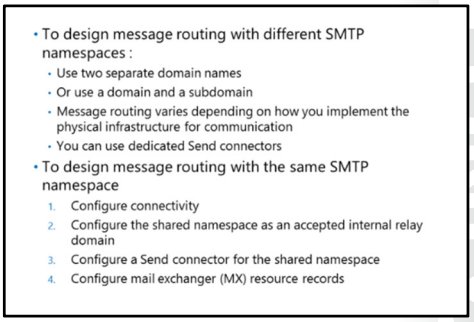
- Which namespace to use. If a smaller organization merges with a larger organization, typically the users in the smaller organization need an email address that is in the domain of the larger organization. If the organizations will share a single namespace, determine how messages will be routed to the appropriate mailbox. Alternatively, the two organizations can use completely separate domain names.
- Whether to synchronize the GAL. In most cases, you should synchronize the GAL between the two organizations. This configuration makes it easier for users in each organization to address messages to the appropriate people. However, if the integration is for only a short time, for example before a full migration, you might not want to make the effort required to configure GAL synchronization.

MCT USE ONLY. STUDENT USE PROHIBITED

- Whether to synchronize free/busy information. If your organization uses calendars extensively to book meetings, you might want to synchronize free/busy information between the two organizations.

Designing Message Routing

When you design message routing in a coexistence scenario, determine whether to use a unique SMTP namespace or the same SMTP namespace for each organization.



Designing Message routing with a unique SMTP namespace

By properly configuring message routing, you can ensure that messages are delivered to the intended recipient. If each organization has a unique SMTP namespace, the message routing is easier to understand and to implement. However, you might not want to use unique SMTP namespaces for business purposes, because it creates the appearance of separate organizations.

Note: If it is implemented, an Edge Transport server can rewrite addresses in order to make multiple messaging systems that have separate namespaces appear as a single namespace. To use address rewriting, the email names for each email account must be unique across organizations. Keep in mind that Exchange Server 2013 now includes an Edge Transport role in Service Pack 1 (SP1), or you can use an Edge Transport server from Exchange Server 2010.

If you use separate SMTP namespaces, the email address for a user changes when the user's mailbox is moved between the two messaging systems. This rewriting can be a problem, because the user does not receive messages sent to the old address in the new mailbox. Users may not receive important messages from customers or internal staff, because the senders are unaware of the new email address. You can mitigate this problem by forwarding messages from the old mailbox to the new Exchange Server 2013 mailbox.

You can create unique SMTP namespaces by using either of the following:

- Two separate domain names. You can use two separate domain names if two organizations are merging. For example, in a merger between Contoso, Ltd., and A. Datum Corporation, the two domains can be contoso.com and adatum.com.
- A domain and a subdomain. You can use a domain name and a subdomain name if one organization is a subsidiary of another. For example, if Contoso, Ltd., is a subsidiary of A. Datum Corporation, the domain names can be adatum.com and contoso.adatum.com.

The configuration of message routing varies depending on how you implement the physical infrastructure for communication. If the two organizations have completely separate data centers and no direct link between the two locations, you can use standard SMTP delivery over the Internet for messages.

If there are two data centers but there is a direct link between them, you can place messaging traffic on the direct link instead of the Internet. To do this, create Send connectors in each organization to direct messages to the appropriate IP address for delivery. Each Send connector is configured with the domain name for the other organization. If there are multiple locations with direct links, you can create multiple Send connectors to optimize delivery.

MCT USE ONLY. STUDENT USE PROHIBITED

If there is a single physical location, you can configure both domains as accepted domains on the Exchange Server 2013 organization. The second domain is configured as an external relay domain. Exchange Server 2013 does not host any mailboxes for an external relay domain, but it does accept messages for a forward relay domain. The messages for an external relay domain are forwarded from Exchange Server 2013 to the external messaging system by using a Send connector. If you centralize message delivery by using Edge Transport servers, you simplify antivirus scanning and you can enforce messaging policies, such as the application of a corporate disclaimer.

Designing message routing with the same SMTP namespace

You can use a single namespace for two messaging organizations. The second messaging organization can be an Exchange Server organization or a different SMTP messaging system.

You typically use a single namespace for two messaging systems temporarily, such as while two organizations are merging. During the transition, you should also configure the recipients in the smaller organization to accept email in both their old domain and the new domain during the migration.

To use the same namespace for multiple organizations, all messages are delivered first to the Exchange Server 2013 organization. The Exchange Server 2013 organization determines whether the recipient is in the Exchange Server 2013 organization or in the second messaging system. If the recipient is in the second messaging system, the Exchange Server organization forwards the message to that system for delivery.

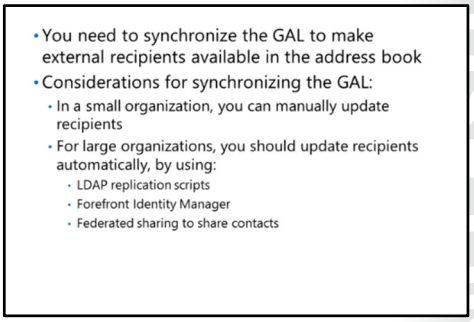
To use a single namespace with two messaging organizations, perform the following configuration steps:

1. Configure connectivity between the two messaging systems. The connectivity can be a direct link between the two systems, or it can be over the Internet.
2. Configure the shared namespace as an accepted internal relay domain. This way, Exchange Server 2013 can relay messages that have no matching recipient in the Exchange Server 2013 organization.
3. Configure a Send connector for the shared namespace. Exchange Server 2013 uses this Send connector to forward messages to the other messaging system. This Send connector is used only if there are no matching recipients in the Exchange Server 2013 organization.
4. Configure mail exchanger (MX) resource records for the Exchange Server 2013 organization. Internet messaging systems use the MX resource records to locate the Edge Transport servers of the Exchange Server 2013 organization.

In addition to configuring the Exchange Server 2013 organization, you must also configure the other messaging system to accept messages from the Exchange Server 2013 organization. In most cases, outgoing messages from the other messaging system are also relayed through the Exchange Edge Transport servers to centralize the management of external message delivery.

MCT USE ONLY. STUDENT USE PROHIBITED

Designing GAL Synchronization



GAL synchronization is an important part of maintaining two separate messaging systems. If you do not configure GAL synchronization, users in each messaging organization have only recipients from their own messaging organization available in the GAL when they address messages. By synchronizing GALs, you can ensure that all recipients are available for addressing in both organizations.

When you synchronize the GAL of an external messaging system into Exchange Server 2013, the external users are created as contacts. If you have only a small number of external users, you can create the contacts manually in the Exchange Server 2013 organization.

When you migrate mailboxes from the external messaging system to the Exchange Server 2013 organization, you need to synchronize the address lists. Before you migrate each mailbox to the Exchange Server 2013 organization, you need to remove the contact for that user. When you migrate the mailbox, the mailbox replaces the contact in the GAL. On the external messaging system, you must remove the mailbox and replace it with a contact containing the email address for that user in the Exchange Server 2013 organization. If you plan to move a large number of mailboxes, you should automate this process.

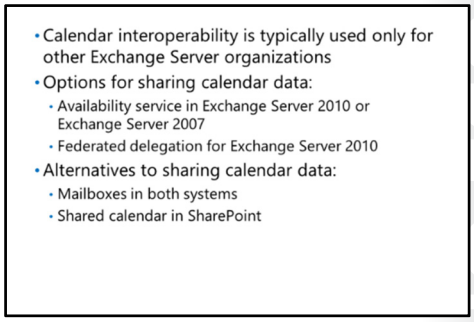
To automate GAL synchronization, you can use the following:

- Lightweight Directory Access Protocol (LDAP) replication scripts. To use LDAP replication scripts, the external messaging system must support the use of LDAP to query recipient information and to create contacts. Although this type of synchronization is possible for other Exchange Server organizations, it might not be possible with other messaging systems. You must run LDAP replication scripts manually, or you must schedule them to run periodically.
- Microsoft Forefront® Identity Manager 2010. This is a flexible tool for synchronizing information between directories. Forefront Identity Manager has additional capabilities for synchronizing information compared to LDAP, so it can synchronize data between a wider range of systems. It can also perform dynamic updates based on events such as the creation of new users and mailbox moves.

Federated delegation is another alternative for sharing contact information between organizations. You can implement federated delegation to allow specific users in your Exchange Server 2013 organization to share contacts with specific users in another Exchange Server 2013 organization. This strategy does not synchronize the GAL between the two Exchange Server organizations, but it can be useful for organizations that want only limited integration, such as partners or subsidiaries.

MCT USE ONLY. STUDENT USE PROHIBITED

Designing Calendar Interoperability



If you implement calendar sharing between messaging systems, users can view the schedules of users in the other organization when they draft and send meeting requests. The level of importance for this capability depends on how your organization uses meeting requests and how long you anticipate coexistence of the two messaging systems to last. For example, this capability is important for your organization if you configure all meeting rooms in your organization as resources and users in both messaging systems need to book those rooms. Typically, you configure calendar interoperability only between Exchange Server organizations.

You have the following options for sharing calendar data:

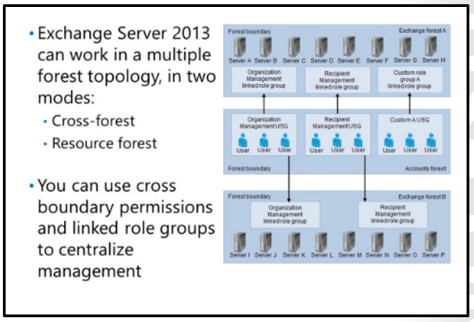
- The Availability service in Exchange Server 2013 or Exchange Server 2010. You can configure a Client Access server in one Exchange Server organization to use the Exchange Server Availability service on the Client Access server in the other Exchange Server organization. This approach gives the first organization the ability to read calendar information in the second organization.
- Federated delegation for Exchange Server 2013. This solution is designed for ongoing interoperability between Exchange Server organizations. One feature of federated delegation is the ability to share calendar information in a selective and controlled way. However, both organizations must be using Exchange Server 2010.

Sharing calendar information can be complex to implement. In some cases, it may be preferable to use an alternative, such as one of the following:

- Mailboxes in both systems. If only a few users need to access calendars in the second Exchange Server organization, the simplest method may be to give those few users a second mailbox in the second Exchange Server organization. Those users now have two mailboxes that you need to maintain. However, you can configure a forwarding address on one of the mailboxes to centralize all messages in a single mailbox.
- Shared calendar in Microsoft SharePoint® services. SharePoint is a web-based solution designed for collaboration, and it provides shared calendars that multiple users can access. This approach can be useful for organizational event calendars and for booking resources, such as meeting rooms.

MCT USE ONLY. STUDENT USE PROHIBITED

Designing Administration Between Exchange Organizations



In some scenarios, usually where security is a top business priority, organizations may create multiple forests in order to provide security boundaries within one organization. In some scenarios, administrators can use the multiple forests to increase security for the environment as well as to use least privileges.

Exchange Server 2013 can work in a multiple forest topology in the following modes:

- Cross-forest. Two AD DS forests exist, each with its own Exchange Server installation and organization.
- Resource forest. One Exchange-enabled AD DS forest exists, and one or more accounts in the AD DS forest exists. The account forest hosts only user accounts, not mailboxes.

Like Exchange Server 2010, Exchange Server 2013 also uses a role-based access control (RBAC) permissions model to determine what each administrator and end user can do. RBAC configuration in each forest is configured independently of all other forests. Permissions that are defined in one forest do not propagate in any way to other forests. If multiple Exchange Server–enabled AD DS forests exist and you want to configure an identical security model in all of them, you must explicitly apply the same configuration in each forest.

Cross-boundary permissions

If you grant permissions by using RBAC, users can only view or modify Exchange related objects within a specific forest and specific management scope. However, there is a way that you can grant permissions so that users can also view and modify Exchange objects in another AD DS forest. By using this approach, called *cross-boundary permissions*, you can centralize Exchange management in a single forest.

The base technology for cross-boundary permissions is *linked role groups*. Linked role groups are used in organizations that install Exchange Server 2013 in a dedicated resource forest and place users in other, trusted foreign forests. Linked role groups actually create a link between a role group in the Exchange-enabled AD DS forest and a universal security group in a foreign forest. Linked role groups can only be associated with one foreign universal security group in one specific AD DS forest.

For example, if an administrator in a foreign forest is a member of the Organization Management linked role group that is located in ForestA, this administrator can only manage Exchange objects in ForestA. A user must be a member of linked role groups in each Exchange forest in order to have permissions to manage each forest.

By using cross-boundary permissions, you can apply role assignment policies to the mailboxes of users who have their mailboxes in an Exchange-enabled AD DS forest, but who have their user accounts in another account forest.

A linked role group is linked to a universal security group in another forest. That group can be any of the following:

- A universal security group dedicated to the specific use of the linked role group.
- A universal security group that is linked to by linked role groups in multiple Exchange Server 2013 forests.
- A role group universal security group in another Exchange Server 2013 forest.

MCT USE ONLY STUDENT USE PROHIBITED

- A universal security group that is associated with an Exchange Server 2007 administrative role or with an Exchange Server 2010 role group.

A linked role group must be linked to a universal security group in another forest. You cannot link a linked role group to a universal security group in the same forest.

You can associate universal security groups in AD DS account forests with role groups in one or more Exchange Server 2013 resource forests. The members of the universal security groups in the AD DS account forest effectively become members of the role groups through this membership. Roles are assigned to the linked role group only in the Exchange enabled AD DS forest. You manage membership in the linked role groups by managing membership of the universal security group in the AD DS account forest. When you add members to the universal security group in the account forest, they are granted the permissions assigned to the linked role group in the Exchange Server 2013 forest. You cannot manage membership of the linked role group from the Exchange Server 2013 forest. Essentially, you manage membership of linked role groups in the account forest, and roles are assigned in the Exchange forest.

Creating linked role groups

You must use Exchange Management Shell to create linked role groups. For example, you want to use a universal security group called Recipient Management Administrators in the domain contoso.com as a linked role group that performs tasks and has permissions equal to the Exchange Recipient administrators role group. Run the following commands:

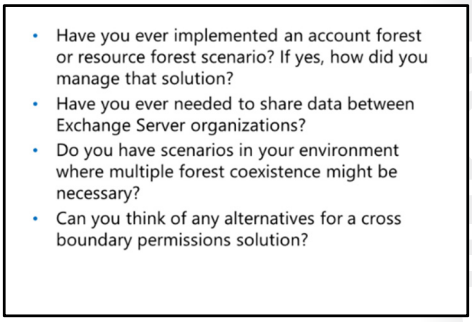
```
$ForeignCredential = Get-Credential – this stores the credentials from contoso.com
forest
$RoleGroup = Get-RoleGroup "Recipient Management" – this defines RoleGroup as
variable with details of Recipient Management role group
New-RoleGroup "Recipient Management - Linked" -LinkedForeignGroup "Recipient
Management Administrators" -LinkedDomainController DC01.users.contoso.com -
LinkedCredential $ForeignCredential -Roles $RoleGroup.Roles – this is cmdlet to
create linked role group
```

You can use one other approach to assign administrative permissions across forest boundaries. Instead of using linked role groups, you can use linked mailboxes. Linked mailboxes work similarly, but you are using users and mailboxes instead of universal security groups and role groups. When a linked mailbox becomes a member of a role group, that linked mailbox, and, in turn, the user in the accounts forest associated with the linked mailbox, is granted the permissions provided by the role group.

Discussion: Experience with Configuring Coexistence

Within your classroom, discuss the following questions related to your organization's experience with configuring coexistence.

- Have you ever implemented an account forest or resource forest scenario? If yes, how did you manage that solution?
- Have you ever needed to share data between Exchange Server organizations?
- Do you have scenarios in your environment where multiple forest coexistence might be necessary?
- Can you think of any alternatives for a cross boundary permissions solution?



MCT USE ONLY. STUDENT USE PROHIBITED

Lesson 3
Designing and Implementing Cross-Forest Mailbox Moves

In several business scenarios, you may need to move mailboxes across forests. However, because Exchange Server works with only one AD DS forest, moving mailboxes is not a trivial process and you must carefully plan it before you start. You should be aware of all of the prerequisites and limitations of this procedure before you implement it, and you should also know how to choose the right approach. This lesson explains how to design and implement cross-forest mailbox moves.

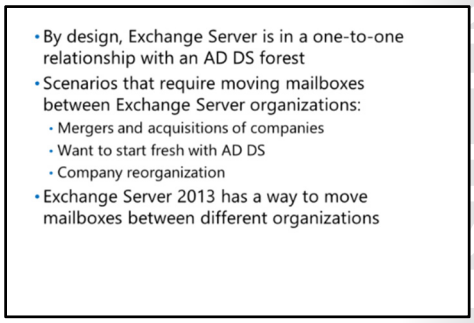
Lesson Objectives

After completing this lesson, you will be able to:

- Describe cross-forest mailbox move scenarios.
- Describe options for implementing cross-forest mailbox moves.
- Describe the prerequisites for implementing cross-forest mailbox moves.
- Prepare for and implement cross-forest mailbox moves.
- Describe considerations for cross-forest mailbox moves.
- Move mailboxes between forests.
- Describe recommendations for implementing cross-forest mailbox moves.

Cross-Forest Mailbox Move Scenarios

Sometimes, organizations need to move mailboxes from one Exchange Server organization to another Exchange Server organization. Each Exchange Server organization can span only one AD DS forest, and this scenario involves multiple AD DS forests.



There are several reasons why organizations choose to have multiple AD DS forests or find themselves with multiple forests. Very often, if one company acquires another, two AD DS and Exchange Server infrastructures exist, and IT departments usually need to merge these infrastructures into one. In addition, if the organization wants to start cleanly with AD DS and Exchange Server, it might choose to implement a completely new forest and then migrate the necessary resources, while leaving any issues in the old system. Also, the organization may deliberately choose to implement separate forests, usually for security reasons. For example, it might want to isolate a business critical department from the rest of company. In such cases, you can maintain coexistence between two parallel infrastructures, or you can merge them.

Whatever scenario is in place, if you have two Exchange Server organizations and you want to merge them into one, you need to move mailboxes from one Exchange Server organization to another.

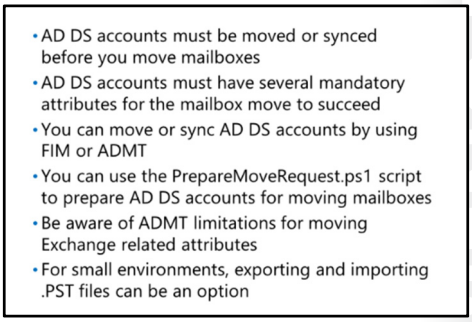
By default, Exchange Server 2013, like previous versions, supports working in only one AD DS forest. However, there is a way to move mailboxes from one Exchange organization to another. This lesson explains how.

Note: At the time of writing this course, supported scenarios for moving mailboxes include:

MCT USE ONLY. STUDENT USE PROHIBITED

- Moving mailboxes between Exchange Server 2013 organizations.
- Moving mailboxes between Exchange Server 2010 and Exchange Server 2013 organizations.

Options for Implementing Cross-Forest Mailbox Moves



Moving mailboxes from one Exchange organization to another involves more than just moving mailbox objects. Because Exchange Server is deeply integrated with AD DS, you also need to move or recreate AD DS objects and attributes in the destination organization. Each mailbox is associated with a specific user account in AD DS, so, when you move a user's mailbox, you also need to move the corresponding user account, or you need to pre-create that account in the destination forest before you initiate the mailbox move. Moving user accounts and mailboxes is not a single-phase process. You should first prepare user accounts and then initiate the mailbox moves for one or more mailboxes.

The mailbox move process in Exchange Server requires that you have an appropriate mail-enabled user account in the destination forest, so the process can attach the mailbox that is being moved to this account. For the migration process to work, this user account must have these mandatory attributes defined:

- **displayName**
- **Mail**
- **mailNickname**
- **msExchArchiveGUID and msExchArchiveName**
- **msExchMailboxGUID**
- **msExchRecipientDisplayType**
- **msExchRecipientTypeDetails**
- **msExchUserCulture**
- **msExchVersion**
- **cn**
- **proxyAddresses**
- **sAMAccountName**
- **targetAddress**
- **userAccountControl**
- **userPrincipalName**

To move or sync an object from one AD DS forest to another, together with attributes defined in the source forest, you can use Microsoft Forefront Identity Manager 2010 (FIM) or Active Directory Migration Tool (ADMT). If you use FIM, you can provision a user object in the destination forest based on a user object in the source forest, and then you can synchronize all or selected attributes. This provisioning is not a move process, but it gives you a lot of flexibility in the provisioning process, and you can make

MCT USE ONLY. STUDENT USE PROHIBITED

selections based on attributes. However, implementing FIM might be too complex and expensive for most organizations. Therefore, Microsoft provides a script in the Windows® PowerShell® command-line interface that prepares the AD DS target object and synchronizes the required attributes for cross-forest moves to work. The script creates the mail enabled user account in the target forest if necessary, or it synchronizes an existing user when possible. This script is called Prepare-MoveRequest.ps1, and it is in the Program Files\Microsoft\Exchange Server\V15\Scripts folder.

You can run this script with appropriate parameters to target user objects to move mailboxes. The script does not actually move mailboxes—it just creates or updates user objects with the attributes that are required for migration in the target organization.

For a local recipient, such as a mail enabled user or a mail enabled contact that already exists, the script does the following:

- If the local forest object is a mail contact, the script converts it to a mail enabled user and persists the existing Exchange-related attributes of the contact.
- If the local forest object is a mail enabled user, the script reuses this mail user and stamps the essential attributes on the local mail user object.

This script uses the existing target accounts if the following are true:

- The target account has a value in **proxyAddresses** that matches one of the **proxyAddresses** of the source account.
- The target account is a mail enabled user. For this to succeed, the target account needs to have mail attributes such as **mail** or **targetAddress**.
- You need to specify the **-UseLocalObject** parameter in the script.

If you choose to use ADMT to move user accounts from one forest to another, be aware that ADMT does not migrate any Exchange attributes, including the mandatory attributes listed earlier. This is because, when ADMT transfers Exchange Server attributes, the target user looks like a legacy mailbox in the target domain. This leaves the target account in an invalid state, which is unexpected for the PrepareMoveRequest.ps1 script. To prevent this, Exchange Server attributes are excluded from ADMT.

 **Note:** At the time of writing this course, a version of ADMT that supports AD DS in Windows Server® 2012 is not available. Information about ADMT and moving mailboxes might change in a new version of ADMT.

The PrepareMoveRequest script in Exchange Server 2013 supports a new parameter, **OverwriteLocalObject**, for user objects created by ADMT. The script copies the mandatory Exchange Server attributes from the source forest user to the target user.

 **Note:** You can choose to not use ADMT at all in the migration process. If you just run the PrepareMoveRequest.ps1 script in the target organization, a new mail enabled user is created in the target AD DS, but the user account is disabled because the password is not migrated with this script. You can move the mailbox, but you need to manually set the password and enable the account.

After user objects are prepared in the target AD DS forest, you can start moving the mailboxes by using the EAC or the Exchange Management Shell.

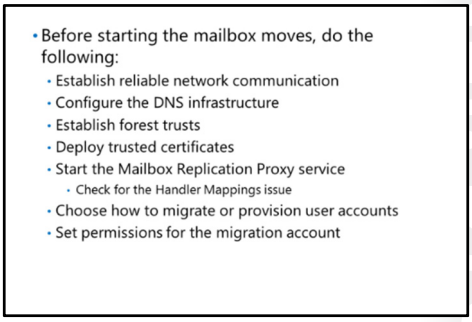
MCT USE ONLY. STUDENT USE PROHIBITED

Alternative to Exchange Server Cross-Forest Mailbox Moves

If you want to avoid synchronizing or moving user account objects between two AD DS forests, you can take an alternative approach. That alternative is to export each mailbox from the source Exchange Server organization to a .PST file, and then to import the .PST files to an appropriate user's mailbox in the destination forest. However, this approach also requires that you manually create AD DS accounts and mailboxes in the destination organization before you export and import the .PST files. This process can be very time consuming, and it is appropriate only for smaller organizations.

Prerequisites for Implementing Cross-Forest Mailbox Moves

To implement a cross-forest mailbox move, you must first prepare the infrastructure in both the source and the destination AD DS forests. Before you start the move, do the following:



1. Establish reliable communication between the source and destination AD DS infrastructures. This communication includes secure and fast network links.
2. Configure the DNS infrastructure in both forests so it supports name resolution for resources in another forest. You can do this by configuring conditional forwarders or stub zones.
3. Establish forest trusts. Although you do not need a forest trust to perform or prepare the mailbox move, you need to establish a trust if you choose to move user accounts with ADMT. To establish a forest trust, you should have the forest functional level on Windows Server 2008.
4. Deploy trusted certificates on the source and destination Exchange servers. If you deploy certificates from internal CAs, you should establish a cross-CA trust. Or you can use public, globally trusted certificates on Client Access servers, in which case you do not need to implement a cross-forest trust.
5. Start the Mailbox Replication Proxy (MRSProxy) service on the Client Access server in the source Exchange Server organization. By default, this service is disabled. To enable it, run the following cmdlet in the Exchange Management Shell:

```
Set-WebServicesVirtualDirectory -Identity "EWS (Default Web Site)" -MRSProxyEnabled
$true
```

   You can also use the **MaxMRSConnections** parameter with this cmdlet. The value of this parameter establishes how many mailbox moves you can do simultaneously. The default value is 100. You should reduce this number if the mailbox move is going across a slow link. Be aware that you need to restart the MSExchangeMailboxReplication service if you change this value. When you restart the service, a database availability group (DAG) is affected, because this service is responsible for copying the log files to the servers hosting the passive copies of the mailbox databases.

Note: If you enable the Mailbox Replication Proxy service on the source Client Access servers, the mailbox move endpoint becomes MrsProxy.svc. In some cases, the Internet Information Services (IIS) configuration is missing the svc-Integrated handler mapping, which results in an error, such as "(405) method not allowed," when you start moving mailboxes. To resolve this issue, navigate to **C:\Windows\Microsoft.Net\Framework\v3.0\Windows**

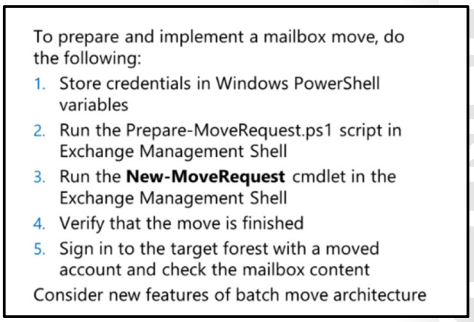
MCT USE ONLY. STUDENT USE PROHIBITED

Communication Foundation, and then execute the following command: **ServiceModelReg.exe –r**. This command reinstalls the handler mappings in IIS. To check for existing handler mappings in IIS, start the IIS console and then, in the center pane, double-click **Handler Mappings**, while the virtual directory or website is selected.

6. Choose how to migrate or provision user accounts and mandatory attributes in the destination forest. As mentioned earlier, you can use ADMT, identity management software such as FIM, or the Prepare-MoveRequest.ps1 script. The script creates new user account objects in the target AD DS, but without password migration and with a limited set of attributes.
7. Set permissions for the migration account. You should establish migration accounts in both AD DS forests, which you use during the move. In the source Exchange Server organization, the migration account must have the privileges of Recipient Administrators. In the target forest, you must delegate the following RBAC roles to the migration account:
   - Move Mailboxes role.
   - Mail Recipients role.
   - Mail Recipient Creation role.

Preparing for and Implementing Cross-Forest Mailbox Moves

After you perform all of the prerequisites, perform the preparation and implementation steps for moving the mailboxes.



First, execute the Prepare-MoveRequest.ps1 script. During execution, this script connects to both the source and the destination Exchange Server organizations, and also to AD DS. Therefore, you must provide credentials for both AD DS forests. To assign a specific authentication credential for both the source and the destination AD DS forests, run the **Get-Credential** cmdlet and store the user input in a temporary variable.

Note: When you run the **Get-Credential** cmdlet, the cmdlet asks for the user name and password. These credentials can be stored in a Windows PowerShell variable and then used in another cmdlet.

Before you actually run the script, execute **Get-Credential** twice—once to store the credentials for the source AD DS, and once to store the credentials of the target AD DS.

For example, execute the following:

```
$Local = Get-Credential
$Remote = Get-Credential
```

In the first command, your credentials for the local (target) forest are stored in a variable called Local, and in the second command, the credentials for the remote (source) forest are stored in a variable called Remote.

MCT USE ONLY. STUDENT USE PROHIBITED

After you store AD DS credentials in these two variables, run the script. For example, to migrate a mailbox that has the alias AidanD@contoso.com from Contoso to Adatum, run the script as follows:

```
Prepare-MoveRequest.ps1 -Identity AidanD@contoso.com -RemoteForestDomainController
DC01.contoso.com -RemoteForestCredential $Remote -LocalForestDomainController
DC01.Adatum.com -LocalForestCredential $Local
```

This command creates a disabled user object, with an email address, in the Adatum.com domain.

Next, run the actual move request. You can do it from the Exchange admin center, or you can run the following cmdlet in the target domain:

```
New-MoveRequest -Identity AidanD@contoso.com -Remote -TargetDatabase "Mailbox
Database Name" -RemoteGlobalCatalog "FQDN of Source DC" -RemoteCredential $Remote -
TargetDeliveryDomain "Target domain name"
```

After the move finishes, the **proxyAddresses** and **targetAddress** attributes are changed in the target forest. If the account is disabled in the target forest, enable it, set a password, and then sign in to Outlook Web App to ensure that the mailbox content is moved.

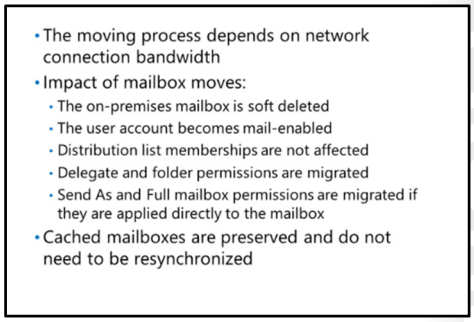
Batch moving in Exchange Server 2013

Exchange Server 2013 introduces a new batch move architecture, which is different from the architecture in earlier versions of Exchange Server. This new architecture improves how the Mailbox Replication service performs moves, with enhanced management capability. The batch-moves architecture in Exchange Server 2013 includes the following features:

- Ability to move multiple mailboxes in large batches.
- Email notification during the move with reporting.
- Automatic retry and automatic prioritization of moves.
- Primary and personal archive mailboxes can be moved together or separately.
- An option for manual move request finalization, so you can review the move before you complete it.
- Periodic incremental synchronizations to update migration changes.

Considerations for Cross-Forest Mailbox Moves

When you move mailboxes between two Exchange Server organizations, consider the following:



- The speed of mailbox moves is limited by the speed and latency of the network connection between organizations. Keep in mind that moving a mailbox is slower than copying a file of the same size. You should perform tests to determine the throughput of your move process, to determine how quickly you can move mailboxes.
- The source mailbox is soft deleted when the move is complete. If there are problems with the moved mailbox, you can recover the soft deleted mailbox for the time period that is set in the deleted mailbox retention limit for the mailbox database.

MCT USE ONLY. STUDENT USE PROHIBITED

- The user account in source forest becomes a mail-enabled user account after the mailbox is moved. This way, the account remains in the GAL of the source Exchange Server organization.
- Distribution list memberships are not affected. In each Exchange Server organization, the user account is already a member of the distribution list. In the source Exchange Server organization, the group member changes from being a mailbox user to being a mail-enabled user.
- The delegate and folder permissions are migrated. When you move a resource mailbox, the delegates for the mailbox are preserved. However, the permissions are not valid unless the delegate and the resource mailbox are both migrated. If you move a resource mailbox first and the delegate later, the delegate has proper permissions after the delegate mailbox move is complete.
- Send As and Full mailbox permissions are migrated if they are applied directly to the mailbox. Similar to delegate and folder permissions, mailbox permissions are migrated, but they are valid only if the recipient that is assigned permissions has also been migrated. These permissions are not migrated if they are inherited.

After a mailbox is moved, the Outlook profile is reconfigured to start accessing the mailbox in the target organization. The cached mailbox for the Outlook profile remains valid and does not need to be resynchronized. This is an important benefit for large mailboxes.

Demonstration: Moving Mailboxes Between Forests

This demonstration shows you the process of mailbox and account preparation and moving. You must first prepare the infrastructure, which this demonstration does not show. You will do the entire procedure in the lab.

Demonstration Steps

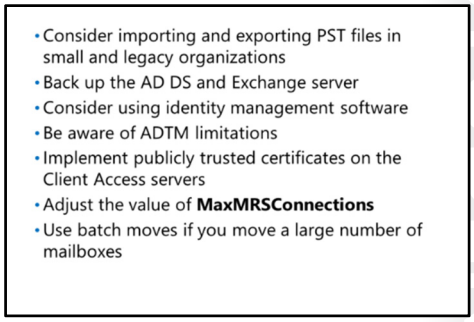
1. On LON-CAS1, open Exchange Management Shell.
2. Change the path to **C:\Program Files\Microsoft\Exchange Server\v15\scripts**.
3. Type **$Local = Get-Credential**, and then press Enter.
4. In the Windows PowerShell Credential window, for **User name** type **Adatum\Administrator** and for **Password** type **Pa$$w0rd**. Click **OK**.
5. Type **$Remote= Get-Credential,** and then press Enter.
6. In the Windows PowerShell Credential window, for **User name,** type **Treyresearch\Administrator**, and for **Password,** type **Pa$$w0rd**. Click **OK**.
7. Type **.\Prepare-MoveRequest.Ps1 -Identity Cindy@treyresearch.net -RemoteForestDomainController trey-dc1.treyresearch.net -RemoteForestCredential $Remote -LocalForestDomainController lon-dc1.adatum.com -LocalForestCredential $Local -TargetMailUserOU "OU=IT,dc=adatum,dc=com"**, and then press Enter.
8. Ensure that you receive the message **1 mailbox(es) ready to move**.
9. Switch to LON-DC1. Open Active Directory Users and Computers.
10. Click the **IT OU**.
11. Ensure that there is an object called **Cindy White** and that it is disabled.
12. Close Active Directory Users and Computers.
13. On LON-CAS1, in the Exchange Administration Center, navigate to **recipients**, and then click the **migration** tab.

MCT USE ONLY. STUDENT USE PROHIBITED

14. Click **Move to this forest**.
15. Click **Cindy White**.
16. Type **Treyresearch\administrator** for the Domain\user name and **Pa$$w0rd** for the password on Windows user account credential prompt.
17. In **Remote MRS Proxy Server**, type **trey-ex1.treyresearch.net**.
18. Select target database to be Mailbox Database 1.
19. In **bad and large mail threshold**, type **10**.
20. Start the batch job.
21. After the job reaches the status **Synced**, click **Complete this migration batch**.
22. Verify that the migration finishes without errors.

Recommendations for Implementing Cross-Forest Mailbox Moves

When you design and implement a cross-forest mailbox move, we recommend that you do the following:



- Ensure that users thoroughly clean their mailboxes by removing all unnecessary messages and then emptying the Junk Mail and Deleted Items folders.
- For smaller organizations, with few mailboxes, consider moving mailboxes by exporting and then importing .PST files. Also, use this method if you want to move mailboxes between an Exchange Server 2003 or Exchange Server 2007 organization and an Exchange Server 2013 organization.
- Back up the AD DS and Exchange servers in both the source and destination organization before you start to move the mailboxes.
- Consider using identity management software to provision accounts and to synchronize attributes between the two AD DS forests.
- Be aware that ADMT currently supports only Windows Server 2008 domains and forests.
- Implement publicly trusted certificates on Client Access servers in Exchange Server organizations.
- Adjust the value of the **MaxMRSConnections** parameter when you set up the MRSProxy service according to the available network bandwidth between the Exchange Server organizations.
- Use batch moves if you move a large number of mailboxes at the same time.

MCT USE ONLY. STUDENT USE PROHIBITED

Lab: Implementing Messaging Coexistence

Scenario

A. Datum has purchased Trey Research and is exploring options for implementing coexistence with Trey Research's messaging organization. Trey Research is currently running Exchange Server 2010 in a separate Exchange Server organization. The A. Datum management team has not yet finalized how to integrate the business units, but it wants to explore how the messaging organizations can be integrated. As a proof of concept, you need to configure messaging coexistence between the two Exchange Server organizations. You also need to evaluate the process for migrating mailboxes from Trey Research to the A. Datum Exchange Server 2013 servers.

Objectives

After completing this lab, you will be able to:

- Implement message routing coexistence.
- Migrate user mailboxes.

Lab Setup

Estimated Time: 90 minutes

Virtual machines	20342B-LON-DC1 20342B-LON-CAS1 20342B-LON-MBX1 20342B-LON-CL1 20342B-TREY-DC1 20342B-TREY-EX1
User Name	Adatum\Administrator
Password	Pa$$w0rd

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

1. On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In Microsoft Hyper-V® Manager, click **20342B-LON-DC1**, and then, in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Log on by using the following credentials:

- User name: **Adatum\Administrator**
- Password: **Pa$$w0rd**

5. Repeat steps 2, 3, and 4 for **20342B-LON-CAS1**, **20342B-LON-MBX1**, and **20342B-LON-CL1**.
6. In Hyper-V Manager, click **20342B-TREY-DC1**, and then, in the Actions pane, click **Start**.
7. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
8. Log on by using the following credentials:
    - User name: **TREYRESEARCH\Administrator**

MCT USE ONLY. STUDENT USE PROHIBITED

- Password: **Pa$$w0rd**

9. Repeat steps 6, 7, and 8 for **20342B-TREY-EX1**.

Exercise 1: Implementing Message Routing Coexistence

Scenario

The first step in implementing coexistence with Trey Research is to configure message routing between the two organizations. All messages sent between the two organizations should be sent across the wide area network (WAN) link between the company data centers and not over the Internet. Messages sent between the two organizations should be secured by using Transport Layer Security (TLS). You need to ensure that messages are routed between the two organizations. Also, you need to establish a forest trust to help ensure a smooth migration.

The main tasks for this exercise are as follows:

1. Implementing a cross-forest certification authority (CA) trust
2. Creating conditional forwarders and mail exchanger (MX) resource records
3. Establishing a forest trust
4. Create send and receive connectors
5. Test the domain security between adatum and Trey Research

▶ Task 1: Implementing a cross-forest certification authority (CA) trust

1. On LON-DC1, open the Certification Authority console.
2. Export the CA certificate in X.509 (.CER) format, and then save it to **C:\AdatumRoot.cer**.
3. Switch to TREY-DC1.
4. Open the Certification Authority console.
5. Export the CA certificate in X.509 (.CER) format, and then save it to **C:\TreyRoot.cer**.
6. From Trey-DC1, open File Explorer, navigate to **\\172.16.0.10\C$**, and then log on as **Adatum\administrator**.
7. Copy the file **AdatumRoot.cer** from **\\172.16.0.10\C$** to **C:** on TREY-DC1.
8. Copy the file **TreyRoot.cer** from **C:** to **\\172.16.0.10\C$\.**
9. On Trey-DC1, open Group Policy Management Console.
10. Edit the **Default Domain Policy**.
11. Navigate to **Computer Configuration**, expand **Policies**, expand **Windows Settings**, expand **Security Settings**, expand **Public Key Policies**, and then right-click **Trusted Root Certification Authorities**. Click **Import**.
12. Choose to import the **AdatumRoot.cer** file from **C:**.
13. Repeat steps 10, 11, and 12 on LON-DC1. Choose to import **TreyRoot.cer** from **C:**.
14. Close Group Policy Management Editor and Group Policy Management Console on both LON-DC1 and TREY-DC1.
15. Refresh Group Policy by executing **gpupdate /force** in Windows PowerShell on LON-CAS1, LON-CAS2, and TREY-EX1.

MCT USE ONLY. STUDENT USE PROHIBITED

▶ Task 2: Creating conditional forwarders and mail exchanger (MX) resource records

1. On LON-DC1, open the DNS Manager console.
2. In **Forward Lookup zone**, create an MX resource record to point to **LON-CAS1.adatum.com**.
3. Create a conditional forwarder for **treyresearch.net**. Type **172.16.0.100** for the DNS server for treyresearch.net.
4. On TREY-DC1, open the DNS Manager console.
5. In **Forward Lookup zone**, create an MX resource record to point to **TREY-EX1.treyresearch.net**.
6. Create a conditional forwarder for **adatum.com**. Type **172.16.0.10** for the DNS server for adatum.com.
7. On LON-CAS1, open Windows Internet Explorer®, type **https://trey-ex1.treyresearch.net/owa**, and then press Enter.
8. Ensure that you do not receive a certificate warning message and that Outlook Web App opens.
9. Close Internet Explorer.
10. Switch to TREY-EX1.
11. Open Internet Explorer, type **https://lon-cas1.adatum.com/owa**, and then press Enter.
12. Ensure that you do not receive a certificate warning message and that Outlook Web App opens.

▶ Task 3: Establishing a forest trust

1. On LON-DC1, open the **Active Directory Domains and Trusts** console.
2. Open **Properties** for the **Adatum.com** domain object.
3. Choose to create a new trust relationship with **treyresearch.net**.
4. Choose to implement a two-way forest trust with forest-wide authentication on both sides.
5. Confirm the trusts.

▶ Task 4: Create send and receive connectors

1. On LON-CAS1, open the Exchange admin center at **https://LON-CAS1.adatum.com/ecp**, and then log on as **Adatum\Administrator** with the password **Pa$$w0rd**.
2. Navigate to mail flow – send connectors.
3. Create a send connector dedicated to the treyresearch.net domain. Click **Partner** type of connector. Select LON-MBX1 as a source server, and select option to **proxy through client access server**.
4. Create a receive connector dedicated to treyresearch.net.
5. Click **Partner** type of connector, and then configure the connector to accept email only from **172.16.0.101**.
6. Assign Webmail.adatum.com certificate on LON-CAS1 to SMTP service.
7. On LON-CAS1, in Exchange Management Shell type: **Set-TransportConfig – TLSSendDomainSecureList adatum.com** and press **Enter.**
8. On LON-CAS1, in Exchange Management Shell type: **Set-TransportConfig – TLSReceiveDomainSecureList treyresearch.net** and press **Enter.**
9. Switch to TREY-EX1, and then open the Exchange Management Console.
10. Navigate to Organization Configuration, and then click **Hub Transport**.

MCT USE ONLY. STUDENT USE PROHIBITED

11. Create a send connector dedicated to the adatum.com domain. Click **Partner type of connector**.
12. In Exchange Management Console, under Server Configuration, assign **Trey Mail Certificate** SMTP service. Overwrite existing certificate for SMTP.
13. In the Exchange Management Console, expand Server Configuration, click **Hub Transport**, and then, in the Hub Transport pane, click **TREY-EX1**.
14. Create a receive connector dedicated to adatum.com.
15. Click **Partner** type of connector, and then configure it to accept email only from **172.16.0.20**.
16. On TREY-EX1, in Exchange Management Shell, type: **Set-TransportConfig –TLSSendDomainSecureList treyresearch.net** and press **Enter.**
17. On TREY-EX1, in Exchange Management Shell, type: **Set-TransportConfig –TLSReceiveDomainSecureList adatum.com** and press **Enter**.

▶ Task 5: Test the domain security between adatum and Trey Research

1. On LON-CL1, open Outlook 2013 and complete the profile creation using default settings.
2. Open a command prompt and type **gpupdate /force** to refresh Group Policy.
3. Send an email message to **Administrator@treyresearch.net**.
4. Open Internet Explorer, and then navigate to **https://trey-ex1.treyresearch.net/owa**.
5. Sign in as **Treyresearch\Administrator** with the password of **Pa$$w0rd**.
6. Ensure that you receive the message from the Adatum administrator.
7. Reply to the message.
8. Open Outlook 2013, and ensure that you receive the message from the Trey Research administrator. Also, ensure that the message has a green check mark. Click the green check mark. (Note: If you don't receive any messages, restart the MSExchangeTransport service on TREY-EX1, the MSExchangeFrontEndTransport service on LON-CAS1, and the MSExchangeSubmission, MSExchangeDelivery and MSExchangeTransport services on LON-MBX1 machine)
9. In Outlook 2013, send three or four messages to **cindy@treyresearch.net**.
10. In Internet Explorer, open https://trey-ex1.treyresearch.net/owa.
11. Sign in as **Treyresearch\cindy** with the password of **Pa$$w0rd**.
12. Ensure that you receive the messages from Adatum administrator.

Results: After completing this exercise, you will have successfully implemented message routing coexistence.

Exercise 2: Migrate User Mailboxes

Scenario

IT management is considering moving all user mailboxes from Trey Research to A. Datum mail servers. As an initial proof of concept, you need to test the process of moving mailboxes between the Exchange organizations. You will prepare servers for the mailbox move and will perform a test move for user Cindy White.

The main tasks for this exercise are as follows:

1. Prepare the source server for a cross-forest mailbox move

MCT USE ONLY. STUDENT USE PROHIBITED

2. Prepare the object for the move

3. Move the user object from Treyresearch to Adatum

4. Validate the move

5. To prepare for the next module

▶ Task 1: Prepare the source server for a cross-forest mailbox move

1. On TREY-EX1, open Exchange Management Shell.
2. Type **Get-WebServicesVirtualDirectory | FL**, and then press Enter.
3. Verify that the value of the **Name** attribute is **EWS (Default Web Site)** and that the value of the **MRSProxyEnabled** attribute is **false**.
4. Type **Set-WebServicesVirtualDirectory -Identity "EWS (Default Web Site)" -MRSProxyEnabled $true**, and then press Enter.
5. Click **Start**, navigate to **Administrative Tools**, and then click **Internet Information Services (IIS) Manager**.
6. Expand Trey-EX1, expand **Sites**, and then click **Default Web Site**.
7. Double-click **Handler Mappings**. Scroll through the list, and verify the presence of ***.svc** entries in the **Path** column. (Note: If you find entries for *.svc, proceed directly to step 15. If not, go to step 8.)
8. Open a Command Prompt window as an Administrator.
9. Change the path to **C:\Windows\Microsoft.Net\Framework\v3.0\Windows Communication Foundation**.
10. Type **ServiceModelReg.exe –r**, and then press Enter.
11. Type **Y**, and then press Enter.
12. Switch to IIS.
13. Click **Default Web Site**.
14. Double-click **Handler Mappings**. Scroll through the list, and search for the ***.svc** in the **Path** column. You should find entries for ***.svc**.
15. Switch to the Exchange Management Shell.
16. Type **Test-MRSHealth**, and then press Enter.
17. Ensure that all three tests have the value **True** in the **Passed** row.

▶ Task 2: Prepare the object for the move

1. On LON-CAS1, open the Exchange Management Shell.
2. Change the path to "**C:\Program Files\Microsoft\Exchange Server\v15\scripts**".
3. Type **$Local = Get-Credential**, and then press Enter.
4. In the Windows PowerShell Credential window, in **User name** type Adatum\Administrator and in Password type **Pa$$w0rd**. Click **OK**.
5. Type **$Remote= Get-Credential**, and then press Enter.
6. In the Windows PowerShell Credential window, for **User name** type **Treyresearch\Administrator** and for **Password** type **Pa$$w0rd**. Click **OK**.

MCT USE ONLY. STUDENT USE PROHIBITED

7. Type **.\Prepare-MoveRequest.Ps1 -Identity Cindy@treyresearch.net -RemoteForestDomainController trey-dc1.treyresearch.net -RemoteForestCredential $Remote -LocalForestDomainController lon-dc1.adatum.com -LocalForestCredential $Local -TargetMailUserOU "OU=IT,dc=adatum,dc=com"**, and then press Enter.
8. Ensure that you receive the message **1 mailbox(es) ready to move**.
9. Switch to LON-DC1. Open Active Directory Users and Computers.
10. Click the **IT** organizational unit.
11. Ensure that there is an object called **Cindy White** there and that it is disabled.
12. Close Active Directory Users and Computers.

▶ Task 3: Move the user object from Treyresearch to Adatum

1. On LON-CAS1, in the Exchange admin center, navigate to **recipients**, and then click the **migration** tab.
2. Click **Move to this forest**.
3. Click **Cindy White**.
4. On the Enter on-premises account credentials page, type Treyresearch\administrator for the Source forest administrator name(domain\administrator name) and Pa$$w0rd for the password. Click next.
5. On the **Confirm the migration endpoint** page, type **trey-ex1.treyresearch.net** in the **Remote MRS Proxy Server** text box, and then click **next**. (Note: If you get an error that connection to trey-ex1.treyresearch.net can not be made, restart TREY-EX1 machine and then try again).
6. For the New migration batch name, type **Cindy**.
7. Select target database to be Mailbox Database 1.
8. In **bad and large mail threshold**, type **10**.
9. Start the batch job.
10. After the job reaches the status **Synced**, click **Complete this migration batch**.
11. Verify that the migration finishes without errors.

▶ Task 4: Validate the move

1. On LON-DC1, open Active Directory Users and Computers.
2. Navigate to the **IT** organizational unit.
3. Set the password for Cindy White to be **Pa$$w0rd**.
4. Enable the account for Cindy White.
5. Open Internet Explorer, and then navigate to **https://lon-cas1.adatum.com/owa**.
6. Sign in as **Adatum\Cindy**.
7. Verify that you can access all email messages that user Cindy had in the treyresearch.net organization.
8. Switch to TREY-EX1. Open the Exchange Management Console.
9. Expand **Recipient Configuration**, and then click **Mailbox**.
10. Ensure that Cindy White is not there anymore.

MCT USE ONLY. STUDENT USE PROHIBITED

▶ Task 5: To prepare for the next module

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

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20342B-LON-CAS1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20342B-LON-MBX1**, **20342B-LON-DC1**, **20342B-LON-CL1**, **20342B-TREY-DC1**, and **20342B-TREY-EX1**.

Results: After completing this exercise, students will have moved a mailbox between Microsoft Exchange Server organizations.

Question: If you are using the internal public key infrastructure (PKI) to issue certificates in both Exchange organizations, why do you need to set up a certification authority (CA) cross-forest trust before you establish a relationship between the organizations?

Question: Why is the user object that is copied from the source domain in a disabled state?

MCT USE ONLY. STUDENT USE PROHIBITED

Module Review and Takeaways

Best Practices

- Always carefully design coexistence before you implement it.
- Use Sharing policies to precisely define the level of information sharing between organizations.
- Always implement secure message routing between partner Exchange organizations.
- Use public trusted certificates to establish federation.
- Use tools such as ADMT or FIM to synchronize user objects between organizations.
- Use organization relationships for a large number of users to share calendar information with an external organization, such as a partner or subsidiary.
- Specify a security distribution group in an organization relationship to limit the sharing of calendar data to specific users.

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
Can't establish connection to Microsoft Federation Gateway	
Secure SMTP messaging between organizations is not working	
Mailbox move between organizations doesn't work	

Review Question(s)

Question: How can FIM help synchronize GALs between two Exchange Server organizations?

Question: Which option for sharing calendar information can you use for both Exchange Server 2013 and Exchange Server 2010?

MCT USE ONLY. STUDENT USE PROHIBITED

Module 12

Designing and Implementing Microsoft Exchange Server Upgrades

Contents:

Module Overview **12-1**

Lesson 1: Planning the Upgrade from Previous Exchange Server Versions **12-2**

Lesson 2: Implementing the Upgrade from Previous Exchange Versions **12-16**

Lab: Upgrading from Exchange Server 2010 to Exchange Server 2013 **12-23**

Module Review and Takeaways **12-35**

Module Overview

Many organizations are already running Microsoft® Exchange Server in their IT infrastructure. In most cases, these organizations will choose to upgrade their current Exchange Server environment to Microsoft Exchange Server 2013 rather than creating a new Exchange Server 2013 organization.

If you already have a previous Exchange Server version installed in your organization, you must plan the upgrade to Exchange Server 2013 from your existing version of Exchange Server. This module provides an overview of the options you have when choosing to implement Exchange Server 2013, and provides details on how to upgrade an existing Exchange Server 2007 or Exchange Server 2010 organization to Exchange Server 2013.

Objectives

After completing this module, you will be able to:

- Plan the upgrade to Exchange Server 2013.
- Implement the upgrade to Exchange Server 2013.

MCT USE ONLY. STUDENT USE PROHIBITED

Lesson 1

Planning the Upgrade from Previous Exchange Server Versions

The first step in upgrading your existing Exchange Server organization to Exchange Server 2013 is to create a plan for the upgrade. During this phase, you need to choose your upgrade strategy and, if required, decide how you will enable coexistence with your current organization. You also need to plan how to migrate various components to the new Exchange servers.

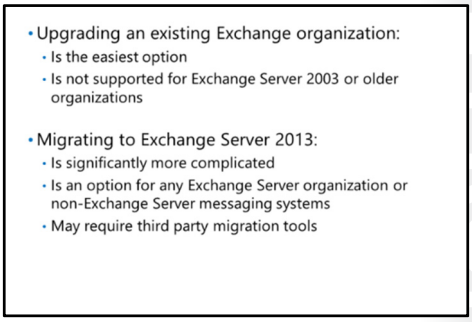
Lesson Objectives

After completing this lesson, you will be able to:

- Describe the options for implementing Exchange Server 2013.
- Describe the supported upgrade scenarios for Exchange Server 2013.
- Choose an upgrade strategy.
- Determine an upgrade scenario.
- Describe the Active Directory® Domain Services (AD DS) requirements for upgrading an existing Exchange Server organization to Exchange Server 2013.
- Describe how client access works during coexistence.
- Describe how message transport works during coexistence.
- Describe how public folders work in Exchange Server 2013.
- Plan the coexistence between messaging policy and compliance features.
- Plan for Exchange Server administration during coexistence.
- Plan multisite upgrades.

Options for Implementing Exchange Server 2013

Almost all organizations deploy some type of messaging system. It could be a cloud-based solution or an on-premises deployment of a previous version of Exchange Server or another messaging system. You have several options when implementing Exchange Server 2013 in an organization that is already running a messaging system.



Upgrading an Existing Exchange Organization

In this scenario, you upgrade an existing Exchange Server 2007 or Exchange Server 2010 organization to Exchange Server 2013. To perform the upgrade, you install Exchange Server 2013 servers into an existing Exchange Server 2007 server or Exchange Server 2010 organization, and then move data and functionality from the existing Exchange servers to new Exchange Server 2013 servers. This is the easiest and least disruptive scenario for integrating Exchange Server-based messaging systems, because the

MCT USE ONLY. STUDENT USE PROHIBITED

different Exchange Server versions share configuration and recipient information automatically. However, you can implement this option only if your organization is currently running the required versions of Exchange Server.

Migrating to Exchange Server 2013

In this scenario, you migrate from a non-Exchange Server messaging system to Exchange Server 2013 or from an existing Exchange Server organization to a new Exchange Server 2013 organization. In both migration scenarios, you install Exchange Server 2013 into a new Exchange organization, and then migrate the current messaging system's data and services to Exchange Server 2013. You can migrate from Exchange Server 2003, Exchange Server 2007, or Exchange Server 2010 to Exchange Server 2013.

You must manually configure all settings and features in the Exchange Server 2013 organization, because nothing is automatically migrated in this scenario. If you need to have both messaging systems coexist for a period of time, you must manually configure all connections between the systems.

When you perform a migration from one Exchange Server organization to another, you also need to deploy a second Active Directory forest, and then migrate all user accounts to the second forest. Each Exchange Server organization requires a unique Active Directory forest.

Exchange Server 2013 does not provide any migration tools or connectors to other messaging systems such as Novell GroupWise, IBM Domin, or cloud-based messaging systems. You can configure email transfer between Exchange Server 2013 and other messaging systems by using Simple Mail Transfer Protocol (SMTP) connectors. However, Exchange Server 2013 does not provide any tools for enabling coexistence or for migrating mailboxes to Exchange Server 2013. In most cases, organizations will use third-party migration tools to simplify the process.

Note: This module focuses on the process for upgrading Exchange organizations to Exchange Server 2013. The previous module covered many of the considerations for configuring coexistence between different Exchange Server organizations. These considerations also apply in a migration scenario when migrating to Exchange Server 2013 from a previous version of Exchange Server.

Supported Upgrade Scenarios

When you upgrade an existing Exchange Server organization to Exchange Server 2013, you should know which upgrade strategies are supported.

The following table identifies some common upgrade strategies.

Exchange Server version	Exchange Server organization upgrade
Exchange 2003 Server	Not supported
Exchange Server 2007 with SP3 and Rollup Update 10 or newer	Supported
Exchange Server 2010 with SP3 or newer	Supported
Mixed Exchange Server 2007 SP3 RU10 and Exchange Server 2010 SP3	Supported

- Upgrades are only supported with Exchange Server 2013 CU1

MCT USE ONLY. STUDENT USE PROHIBITED

Exchange Server version	Exchange organization upgrade	Comments
Exchange Server 2003	Not supported	Although an upgrade is not supported, you can use a migration strategy to transition to Exchange Server 2013. Alternately, you can upgrade the Exchange 2003 server organization completely to Exchange Server 2007 or Exchange Server 2010, and then perform an upgrade to Exchange Server 2013. If you still have Exchange 2003 servers deployed in an organization which also includes Exchange 2007 or Exchange 2010 servers, you must also remove all Exchange Server 2003 servers from the organization before starting the upgrade.
Exchange Server 2007 with SP3 and Update Rollup 10 or newer	Supported	Before upgrading from Exchange Server 2007, you must upgrade all of your organization's Exchange Server 2010 servers, including Edge Transport servers, to SP3 and Update Rollup 10 or newer.
Exchange Server 2010 with SP3 or newer	Supported	Before upgrading from Exchange Server 2010, you must upgrade all of your organization's Exchange Server 2010 servers, including Edge Transport servers, to SP3.
Mixed Exchange Server 2010 and Exchange Server 2007 organization	Supported	When you are ready to upgrade your mixed-mode environment, upgrade each Active Directory site individually. If you have Active Directory sites with only Exchange 2010 or Exchange 2007 in them, follow the instructions for upgrading from that version for that Active Directory site. For example, if you have Exchange Server 2010 in Active Directory site A, then follow the upgrade instructions for Exchange Server 2010. If you have Exchange Server 2007 in Active Directory site B, then follow the upgrade instructions for Exchange Server 2007.



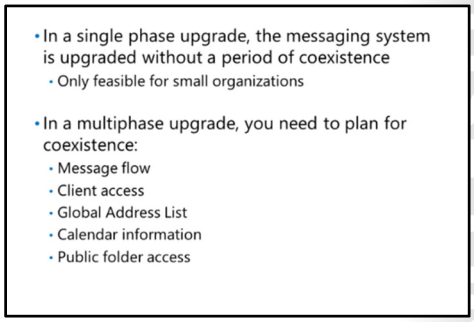
Note: Important: To upgrade an Exchange Server 2007 or Exchange Server 2010 organization to Exchange Server 2013, you must deploy Exchange Server 2013 Cumulative Update 1 (CU1). In this module, all references to Exchange Server 2013 include CU1.

After you deploy a new Exchange Server 2013 organization, you cannot add servers running earlier versions of Exchange Server to the organization. In other words, Exchange Server 2013 does not support the addition of earlier Exchange Server versions to an Exchange organization that includes only Exchange Server 2013 servers.

MCT USE ONLY. STUDENT USE PROHIBITED

Choosing an Upgrade Strategy

Once you have decided to perform an upgrade, you must select the appropriate upgrade strategy for your organization. You can choose between several options. The selection you make depends upon your current environment, your organization's requirements for data migration, and your project timeline.



Your first choice when planning the upgrade is to decide whether to use a single-phase upgrade or a multiphase upgrade.

Single-Phase Upgrade

In a single-phase upgrade, you replace your existing messaging system with Exchange Server 2013, and move all required data and services to the new system. You do not need to plan for an extended period of coexistence between the two systems.

Typically, you perform this type of upgrade over a short period, perhaps a weekend. This approach enables you to shut down the entire messaging system and replace it with Exchange Server 2013, so that when users return to work the new messaging system is operational. In this scenario, the period of coexistence or interoperability is quite short.

While this upgrade is the fastest option, it also introduces a significant risk if the upgrade fails. This scenario is feasible only for small organizations that must replace one or two servers, and that have only a small number of users to migrate.

Multiphase Upgrade with Coexistence

In a multiphase upgrade, you upgrade one server or site at a time to Exchange Server 2013. Because you spread this incremental upgrade over a longer period, you decrease your organization's risk. However, in this scenario, you also must plan for coexistence or interoperability. This is the best approach for medium and large organizations, because of their complex messaging requirements.

In most coexistence scenarios, you must ensure that there is no disruption for users. This requirement means that you need to deal with the following components during the upgrade:

- *Email message flow*. When you run two versions of Exchange Server, users must be able to send email to any other organizational users, and to and from users on the Internet. Message flow should be transparent to users. Users do not need to know—nor should it matter—which version of Exchange Server hosts their mailbox or their recipient's mailbox.
- *Client access*. When you run two versions of Exchange Server, the client access methods may not be the same for both versions. However, this needs to be transparent to users. They must be able to connect to either Exchange Server version without error. Users should experience very minimal, if any, disruption when their mailbox is moved from one server version to another.
- *Global Address List (GAL)*. The GAL must contain all messaging recipients, regardless of the Exchange Server version that hosts the user's mailbox. In addition, when users reply to messages received before or after their mailbox is moved, the message must be correctly delivered.
- *Calendar information*. To facilitate scheduling of meetings between the two Exchange Server versions, you must ensure that Free/Busy information is available from both systems as user mailboxes are moved.

MCT USE ONLY STUDENT USE PROHIBITED

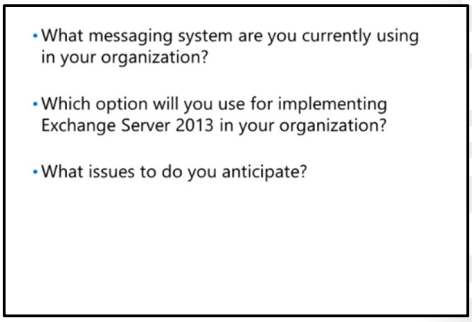
- *Public folder contents*. If the organization stores important information in public folders, you may need to ensure that public folder contents are accessible in both Exchange Server versions.

When you upgrade an existing Exchange Server organization to Exchange Server 2013, it is fairly easy to ensure full functionality during the period of coexistence. However, it is important to keep the user experience in mind during the upgrade project to ensure that users experience as little disruption in email services as possible.

Discussion: Choosing an Upgrade Scenario

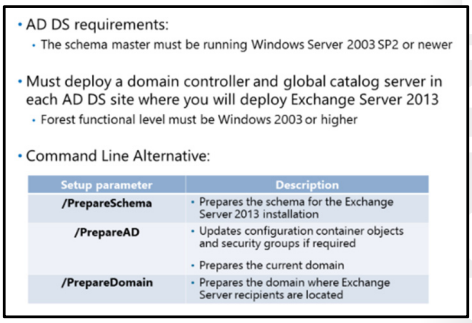
Discussion questions:

1. What messaging system are you currently using in your organization?
2. Which option will you use for implementing Exchange Server 2013 in your organization?
3. What issues do you anticipate?



AD DS Requirements for Installing Exchange Server 2013 in an Existing Exchange Server Organization

The first step in upgrading an Exchange Server organization to Exchange Server 2013 is to ensure that the AD DS environment in your Exchange organization meets the following requirements:



- Your schema master must be running the Windows Server® 2003 operating system with Service Pack 2 (SP2) or newer.
- You must have configured your AD DS forest to be at least at the Windows Server 2003 forest-functional level or higher.
- You must deploy at least one domain controller and one global catalog server with a writeable AD DS copy in each Active Directory site. The domain controller must be running the Windows Server 2003 operating system with SP2 or newer. Exchange Server 2013 cannot use read-only domain controllers (RODCs) or read-only global catalog servers running Windows Server 2008 or Windows Server 2012.

Preparing AD DS for the Exchange Server 2013 Installation

Before beginning the upgrade, you must prepare AD DS for Exchange Server 2013. You can prepare your AD DS by running the Exchange Server 2013 Setup Wizard with a user account that has the permissions required to prepare AD DS and the domain. To prepare the AD DS schema and configuration partition, you must use an account that is a member of the Schema Admins and Enterprise Admins groups. By using this type of account, the wizard automatically prepares AD DS and the current domain.

Note: You must install the AD DS Tools option from the Remote Server Administration Tools on the server from which you are upgrading the AD DS domain.

MCT USE ONLY. STUDENT USE PROHIBITED

Command-Line Alternatives

Instead of running the Exchange Server 2013 Setup Wizard to prepare AD DS for Exchange Server, you can alternatively run the Exchange Server 2013 setup utility from the command line. There are two different approaches you can use to prepare AD DS using the command line option.

Note: Whenever you run the Exchange Server 2013 setup command from the installation DVD, you must include the **/IAcceptExchangeServerLicenseTerms** parameter. This is the only way to confirm your acceptance of the license agreement from the command line.

If you are signed in with an account that is a member of the Enterprise Admins and Schema Admins groups, you can prepare AD DS by running the setup /PrepareAD command. In an upgrade scenario, this command performs the following tasks:

- Prepares the AD DS schema as described in the following list:
    - Verifies that the schema has been updated and that the organization is up to date by checking the **objectVersion** property in AD DS. The **objectVersion** property is in the CN=<your organization>,CN=Microsoft Exchange,CN=Services,CN=Configuration,DC=<domain> container. The **objectVersion** value for Exchange Server 2013 CU1 is 15449. If the schema has not been updated, it will be updated when you run this command.
    - Sets the **msExchProductId** of the Exchange organization object to 15.00.0516.032. The **msExchProductId** property is in the CN=<your organization>,CN=Microsoft Exchange,CN=Services,CN=Configuration,DC=<domain> container.
- If you are upgrading from Exchange Server 2007, this command makes a significant number of changes in AD DS. The command will create several new objects and containers needed for Exchange Server 2013 under CN=<Organization Name>,CN=Microsoft Exchange,CN=Services,CN=Configuration,DC=<root domain> and assign specific permissions throughout the configuration partition. It also creates the Microsoft Exchange Security Groups organizational unit (OU) in the root domain of the forest, assigns specific permissions to this OU, and creates the management role groups within the Microsoft Exchange Security Groups OU.
- If you are upgrading from Exchange Server 2010, the changes are less significant than an Exchange Server 2007 upgrade, but several of the role-based access control (RBAC) roles have been updated to enable administrators to run new cmdlets and configure new properties.
- Prepares the local domain for Exchange Server 2013.

Note: The **Setup /PrepareAD** command performs several additional tasks when run in an organization that does not already have a previous version of Exchange Server deployed.

To perform this command, you must be a member of the Enterprise Admins security group, and you must run this command on a computer that is in the same domain as the schema master domain controller. If you have more than one domain, you should wait after running this command, so that changes performed to AD DS are replicated to all other domains and domain controllers.

In some organizations, only specified users can ever be added to the Schema Admins group. If you do not have the option of running the **setup /PrepareAD** command as a member of the Schema Admins group, then you must run the **setup /PrepareSchema** command to upgrade the Exchange schema to the Exchange Server 2013 version. You must run this command before you run the **setup /PrepareAD**. To execute this command, you must also be a member of the Enterprise Admins or Schema Admins groups.

MCT USE ONLY. STUDENT USE PROHIBITED

This command performs the following tasks:

- Connects the Exchange Server to the schema master domain controller.
- Imports Lightweight Directory Access Protocol (LDAP) Data Interchange Format (DIF) files to update the schema with attributes specific to Exchange Server 2013.
- Sets the schema version (ms-Exch-Schema-Version-Pt) to 15137. This is the schema version for Exchange Server 2013 CU1.

If you have multiple domains in your organization, then you also need to run the **setup /PrepareDomain** command in each domain where Exchange recipients will be located. You do not need to run this command in a domain where you ran **setup /PrepareAD**. Alternatively, you can also run **setup /PrepareDomain:<FQDN of domain you want to prepare>** to prepare a specific domain, or you can run **setup /PrepareAllDomains** or **setup /pad** to prepare all domains in your organization.

This command performs the following tasks in an upgrade scenario:

- Sets the objectVersion property in the Microsoft Exchange System Objects container under DC=<root domain>. This objectVersion property contains the version of domain preparation. The version for Exchange Server 2013 is 13236.
- Modifies some permissions assigned within the domain partition.

Understanding Client Access Server Coexistence

Client Access option	Mailbox on Exchange 2007	Mailbox on Exchange 2010
Autodiscover	Proxy to Exchange 2013 Mailbox server	Proxy
Outlook Anywhere	Proxy	Proxy
Outlook Web App	Redirect to legacy URL	Proxy or redirect to ExternalURL
Exchange ActiveSync	Proxy to Exchange 2013 Mailbox server	Proxy
Exchange Web Services	Autodiscover provides URL	Proxy
POP3/IMAP4	Proxy	Proxy

As you plan your upgrade of an existing Exchange organization to Exchange Server 2013, it is important that you understand how client access will work during the period of coexistence and what impact this will have on your upgrade plan. When you deploy Exchange Server 2013 in an existing Exchange Server 2007 or Exchange Server 2010 organization, one of your first steps is to configure all client connections to connect to the Exchange 2013 Client Access servers. To do this, you normally change the Domain Name System (DNS) records for any generic client access names, such as mail.contoso.com, to refer to the Exchange 2013 Client Access servers. You may also need to change firewall rules to enable client access to the Client Access server from the Internet.

Client Access During Coexistence

When users with mailboxes on Exchange Server 2007 or Exchange Server 2010 connect to an Exchange 2013 Client Access server, the Exchange 2013 Client Access server forwards each client request to the correct Exchange servers using the following processes:

- During the Autodiscover process, the client needs to obtain information required to configure the client profile. If the user mailbox is located on an Exchange 2007 server, the client request is proxied to an Exchange 2013 Mailbox server to retrieve this information. If the user mailbox is located on an Exchange 2010 server, the user request is proxied to an Exchange 2010 Client Access server, which generates the Autodiscover response, and sends it back the Exchange 2013 Client Access server, which then forwards it to the client.
- When a Microsoft Outlook® Anywhere client connects to the Client Access server and the user mailbox is located on either an Exchange 2007 or Exchange 2010 mailbox server, the Exchange 2013

MCT USE ONLY. STUDENT USE PROHIBITED

Client Access server proxies the client request to the corresponding Exchange 2007 or Exchange 2010 Client Access server. The Exchange 2007 or Exchange 2010 Client Access server connects to the corresponding Mailbox server. The mailbox information is returned through the Exchange 2007 or Exchange 2010 Client Access server to the Exchange 2013 Client Access server, and then to the client.

- When a Microsoft Outlook Web App client connects to the Exchange 2013 Client Access server and the user mailbox is located on an Exchange Server 2007 Mailbox server, the client redirects to the Exchange Server 2003 URL configured on the Client Access server. For example, if the client connects to the Exchange Server 2010 Client Access server by using the URL https://Mail.contoso.com, the request might be redirected to https://legacy.contoso.com. The client then communicates with the Exchange Server 2003 front-end server to access the user mailbox.
- When an Outlook Web App client connects to the Client Access server and the user mailbox is located on an Exchange Server 2007 Mailbox server, the web client is redirected to a legacy name which references an Exchange Server 2007 Client Access server. The Outlook Web App client connection will use the legacy URL, and the request is not proxied through the Exchange 2013 Client Access server again.
- When an Outlook Web App client connects to the Client Access server and the user mailbox is located on an Exchange Server 2010 Mailbox server in the same Active Directory site, the Exchange 2013 Client Access server will proxy the request to an Exchange Server 2010 Client Access server.
- When an Outlook Web App client connects to the Client Access server and the user mailbox is located on an Exchange Server 2010 Mailbox server in a different Active Directory site, the process depends on whether the Exchange 2010 Client Access server in the target site has the ExternaURL configured. If the External URL is configured, the client will be redirected to that URL. If the Exchange 2010 Client Access server does not have the ExternalURL configured, the client request will be proxied to an Exchange 2010 Client Access server in the target site.
- When a Microsoft Exchange ActiveSync® client connects to the Exchange 2013 Client Access server and the user mailbox is located on an Exchange 2007 Mailbox server, the Exchange 2013 Client Access server proxies the client request to the Exchange 2013 Mailbox server. The Exchange 2013 Mailbox server proxies the request an Exchange 2007 Client Access server, which uses a remote procedure call (RPC) to connect to the Exchange 2007 Mailbox server.
- When an Exchange ActiveSync client connects to the Exchange 2013 Client Access server and the user mailbox is located on either an Exchange 2007 or Exchange 2010 mailbox server, the Exchange 2013 Client Access server proxies the client request to a corresponding Exchange 2007 or Exchange 2010 Client Access server.
- When a client tries to access the Exchange Web Services virtual directory and the user mailbox is located on an Exchange 2007 mailbox server, the Autodiscover information provided to the client is used to connect the client to an Exchange 2007 Client Access server for Exchange Web Services. If the user mailbox is located on Exchange Server 2010, the client request is proxied to an Exchange 2010 Client Access server by the Exchange 2013 Client Access server.
- When a client connects to an Exchange Server 2013 Client Access server using either Post Office Protocol version 3 (POP3) or Internet Message Access Protocol 4 (IMAP4) and the user mailbox is on either an Exchange 2007 or Exchange 2010 mailbox server, the Exchange 2013 Client Access server proxies the client request to the corresponding Exchange 2007 or Exchange 2010 Client Access server.

Considerations for Client Access During Coexistence

When implementing client access during coexistence, consider the following:

- Outlook Web App clients will always display the user interface for the Exchange Server version where the user mailbox is located. For example, if the user's mailbox is located on an Exchange Server 2007

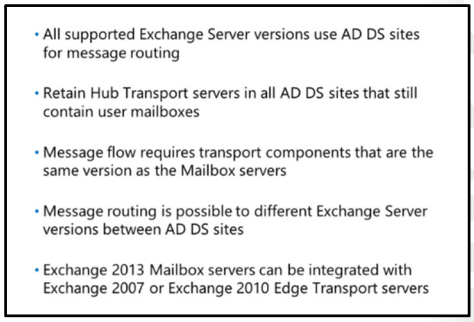
MCT USE ONLY. STUDENT USE PROHIBITED

Mailbox server, the user will see the Exchange Server 2007 version of Microsoft Office Outlook Web Access.

- Outlook Web App redirection does not support single sign-on (SSO). Users will be prompted for their credentials when they connect to the Exchange 2013 Client Access server. When users are redirected to an Exchange 2007 or Exchange 2010 Client Access server, they will be prompted again for their credentials.
- Users will not be able to connect to their mailbox on an Exchange 2013 Mailbox server if they first connect to an Exchange 2007 or Exchange 2010 Client Access server. Users can connect to the Exchange 2007 or Exchange 2010 Client Access server and gain access to their Exchange 2007 or Exchange 2010 mailboxes. Therefore, before you begin moving mailboxes to the Exchange 2013 Mailbox servers, you must first configure all client connections to use the Exchange 2013 Client Access server.
- You must maintain Exchange 2007 or Exchange 2010 Client Access servers as long as any user mailboxes remain on the corresponding Mailbox servers. The Exchange 2013 Client Access server always proxies or redirects client requests to previous Client Access servers, not to Mailbox servers.

Understanding Message Transport Coexistence

A second coexistence component between the two Exchange Server versions is message transport. Message transport coexistence is configured automatically, as long as you do not remove Exchange 2007 or Exchange 2010 Hub Transport servers before moving all required functionality and data to Exchange Server 2013.



Message Routing During Coexistence

As you deploy Exchange Server 2013 Mailbox servers in an Exchange 2007 or Exchange 2010 organization, message transport works as follows:

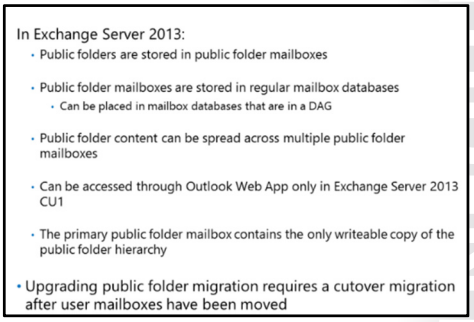
- All three Exchange server versions use AD DS sites for message routing. This means that the message routing topology will not change significantly during and after the upgrade.
- You must retain message routing functionality for each version of Exchange server in each AD DS site as long as there are mailboxes located on that Exchange server version. Since Exchange Server 2013 message routing is provided by the Mailbox server role, you will have message routing for Exchange Server 2013 when you deploy the first server. For Exchange Server 2007 and Exchange Server 2010, you must not remove the last Exchange 2007 Hub Transport server from an Active Directory site until you have removed all of the mailboxes from the Exchange 2007 or Exchange 2010 Mailbox servers in that site.
- If you have Exchange 2007 or Exchange 2010 servers deployed in a site, messages will flow from the Exchange 2013 Mailbox server to the Exchange 2007 or Exchange 2010 Hub Transport server, and then to the Exchange 2007 or Exchange 2010 Mailbox server. Messages sent from an Exchange 2007 or Exchange 2010 mailbox would follow the reverse route. Exchange 2013 Mailbox servers cannot communicate directly with Exchange 2007 or Exchange 2010 Mailbox servers.
- Message routing between AD DS sites can use a mixture of Exchange Server versions. Exchange 2013 Mailbox servers in one AD DS site can send mail to Exchange 2007 or Exchange 2010 Hub Transport servers in another site.

MCT USE ONLY. STUDENT USE PROHIBITED

- Message routing to and from the Internet can use either the Exchange 2013 infrastructure or the Exchange 2007 or Exchange 2010 infrastructure. If your current deployment uses Exchange 2007 or Exchange 2010 Edge Transport servers for inbound email, you can continue to have the Edge Transport servers forward all messages to the Exchange 2007 or Exchange 2010 Hub Transport server. As you deploy the Exchange 2013 Mailbox and Client Access servers, you can add Exchange 2013 Mailbox servers to the edge subscription so that the existing Edge Transport servers can forward messages to the Exchange 2013 Mailbox server. Alternatively, you can deploy Edge Transport servers in Exchange Server 2013, which now includes the Edge Transport role in Service Pack 1 (SP1). If you are using a third-party SMTP gateway server, the Default Frontend receive connector created on Exchange 2013 Client Access servers is automatically configured to accept anonymous connections from the SMTP gateway server.
  For outbound messages, you can add Exchange 2013 Mailbox servers to the SMTP Send connector that is responsible for sending messages to the Internet. This enables outbound messages to be sent through either the Exchange 2013 Mailbox servers or through the Exchange 2007 or Exchange 2010 Hub Transport servers.

Understanding Exchange Server 2013 Public Folders

In Exchange Server 2013, the underlying architecture for public folders has changed entirely without significantly affecting the user experience with public folders.



Public Folders in Exchange Server 2013

In Exchange Server 2013:

- Public folders are stored in a special type of mailbox called a public folder mailbox. In previous versions of Exchange Server, public folders were stored in a separate public folder database. In Exchange Server 2013, the public folder mailboxes are stored in regular mailbox databases. The public folder mailboxes store the public folder hierarchy as well as the public folder contents.
- Public folder mailboxes can be stored in mailbox databases that are part of a DAG. In previous versions of Exchange Server, public folders used a public folder replication process to enable redundancy. By storing the public folder mailboxes in a mailbox database that is part of a database availability group (DAG), you can provide high availability for the public folder deployment using the same mechanism as the one used for providing high availability for mailboxes.
- Public folders are spread across multiple public folder mailboxes. In previous versions of Exchange Server, you could replicate public folder contents to public folder databases located in different locations to enhance client access to public folder contents. In Exchange Server 2013, you can create public folders and store them in different mailboxes, which can be located on Mailbox servers in different locations.

Note: An important difference between public folder replication in previous versions of Exchange Server and distributing public folders across multiple mailboxes in Exchange Server 2013 is that, in Exchange Server 2013, you can have only a single copy of the data. In previous versions of Exchange Server, you can have multiple copies of the public folder contents, and public folder replication is a multi-master process. In Exchange Server 2013, you can only

MCT USE ONLY. STUDENT USE PROHIBITED

store the public folder contents in one mailbox, and all clients must access that mailbox to see the public folder contents. If you put the public folder mailbox in a database that is part of a DAG, the mailbox is highly available, but all clients still only access the mailbox in the active copy of the database.

- Public folders are accessed by Outlook 2007 or newer clients. In Exchange Server 2013 Release to Manufacture (RTM), Outlook Web App clients cannot access the public folders. In Exchange Server 2013 CU1, you can add public folders to the Favorites list in Outlook Web App and access individual public folders. You cannot browse the whole public folder tree in Outlook Web App.

To implement public folders in Exchange Server 2013, you first must create a primary public folder hierarchy mailbox. The primary public folder mailbox contains the only writeable copy of the public folder hierarchy. After creating the primary public folder mailbox, you can create additional public folder mailboxes as secondary public folder mailboxes. The secondary public folders will contain read-only versions of the public folder hierarchy.

After creating the primary public folder mailbox, you can begin creating public folders. By default, all public folders are created in the primary public folder mailbox. If you create a secondary public folder mailbox, you can create public folders in the secondary public folder mailbox only if you create the public folder using the **new-publicfolder** cmdlet with the **–mailbox** parameter.

Considerations for Migrating Public Folders

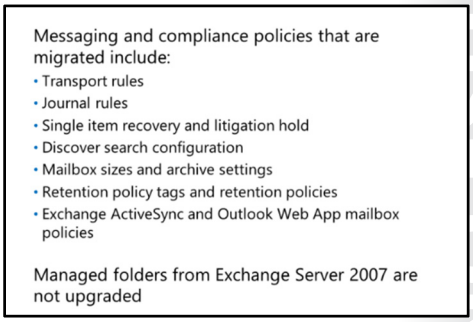
Because of the new public folder architecture, there are some special considerations for migrating public folders:

- During coexistence, users with mailboxes on Exchange 2013 can access public folders on Exchange 2007 or Exchange 2010 Mailbox servers. But users with mailboxes on Exchange 2007 or Exchange 2010 cannot access public folders on Exchange Server 2013. This restriction means that you will need to switch over access to public folders at some point in the upgrade.
- Because only Exchange 2013 mailbox users can access public folders in Exchange 2013, you should migrate all users who require public folder access to Exchange 2013 before migrating the public folders.
- In some companies, the public folders contain a very large amount of data. During the migration, you can copy this data to the Exchange 2013 public folders incrementally. After all the data is copied over, you can complete the cutover.
- You can switch the public folders back to the previous version of Exchange Server, but all changes made to the public folders while they were hosted on Exchange Server 2013 will be lost.

MCT USE ONLY. STUDENT USE PROHIBITED

Planning Messaging Policy and Compliance Coexistence and Migration

Exchange Server 2010 introduced several messaging and policy compliance solutions that could be used to enforce email related corporate policies. During coexistence, these policies are applied for mailboxes on both the Exchange 2010 and Exchange 2013 servers. When you remove the Exchange 2010 servers, the policies continue to be applied. Some examples of how policies are migrated to Exchange 2013 include:



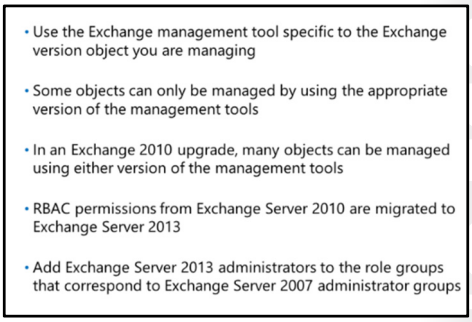
- Transport rules continue to be applied to messages sent and received on both messaging systems. When you complete the upgrade to Exchange Server 2013, you can modify the transport rules to take advantage of the new data loss protection features in Exchange Server 2013.
- Journal rules continue to be applied on both Exchange Server 2010 and Exchange Server 2013 during and after mailbox moves.
- Single item recovery and litigation hold policies are migrated to Exchange Server 2013 when the mailboxes are migrated. Any messages being held in the Recoverable Items folder while the mailbox is on Exchange Server 2010 are migrated with the mailbox to Exchange Server 2013.
- Discovery search configurations are migrated to Exchange Server 2013. In Exchange Server 2013, the Discovery Management group continues to have permission to search mailboxes and apply compliance policies. Any discovery search mailboxes, including the default mailbox created in Exchange Server 2010, are migrated to Exchange Server 2013, and can continue to be used to store eDiscovery results.
- Any policies related to mailbox sizes and mailbox archive configuration are migrated to Exchange Server 2013. When a mailbox is configured with a custom size, the size is retained during the mailbox move. If a mailbox is configured with an archive mailbox, you can move the archive to Exchange Server 2013 at the same time as the regular mailbox, or at a different time.
- Retention policy tags and retention policies are available on Exchange Server 2013 as soon the first Exchange 2013 server is migrated. If any policies are applied to Exchange 2010 mailboxes, they remain applied after the mailbox is moved to Exchange Server 2013.
- Exchange ActiveSync mailbox policies and Outlook Web App mailbox policies continue to be applied during and after the mailbox moves.

In Exchange Server 2007, you could use Managed Folders to manage the contents of user mailboxes. These settings are not upgraded to Exchange Server 2013 and cannot be converted to Retention Policies in Exchange Server 2013. If you are currently running a mixed environment with Exchange 2007 and Exchange 2010 servers, you can use the Exchange 2010 tool to migrate Managed Folder settings to Retention Policies before upgrading to Exchange Server 2013.

MCT USE ONLY. STUDENT USE PROHIBITED

Planning for Administration

During the upgrade to Exchange Server 2013, you also need to plan for administrative coexistence. In this scenario, you need to consider how you will use the Exchange Server management tools and how you will delegate permissions.



Management Console Coexistence

In Exchange Server 2013, the Exchange Administration Center (EAC) replaces the Exchange Management Console, which is used in Exchange Server 2007 and Exchange Server 2010. While most of the functionality is similar between the two management tools, you need to consider the following during an upgrade:

- In general, always use the management tool that matches the version of the Exchange objects that you are managing.
- Some objects can only be managed from the appropriate version of the Exchange management tools. For example, if you are creating a new mailbox on an Exchange 2007 or Exchange 2010 Mailbox server, you must use the Exchange Management Console that corresponds to the Exchange Server version. Mailbox databases, public folder databases, and Exchange server settings must be configured using the appropriate version of the management tools.
- When upgrading from Exchange Server 2010 to Exchange Server 2013, you can modify and manage many objects using either version of the management tools. For example, you can modify mailboxes, transport rules, and global message delivery settings using either version of the management tools.
- When upgrading from Exchange Server 2007, fewer settings can be modified using both management tools. In some cases, you can view objects, but not modify objects. In some cases, you cannot even view objects configured on the Exchange Server version that is different from the management tool you are using.

Delegating Administration During Coexistence

Exchange Server 2013 and Exchange Server 2010 both use RBAC to assign administrative permissions in Exchange. During an upgrade from Exchange Server 2010 to Exchange Server 2013, all permissions are migrated. The only time you might need to make a change is if you have granted users administrative rights to a particular Exchange server. If you want to replicate this configuration in Exchange Server 2013, you must grant permissions to an Exchange 2013 server that replaces an Exchange 2010 server.

However, the model for delegating administrative permissions is quite different between Exchange Server 2007 and Exchange Server 2013. Exchange 2007 Setup creates several Active Directory groups with designated permissions in AD DS and in the Exchange organization. To delegate permissions, you just add users to the appropriate Active Directory groups. RBAC replaces this model in Exchange Server 2013.

When you install Exchange Server 2013 servers in an Exchange Server 2007 organization, the Exchange Server 2010 role groups are added to AD DS, and the Exchange Server 2007 groups are retained. When assigning permissions on Exchange Server 2007 servers, use the Exchange Server 2007 groups. When assigning permissions on Exchange Server 2010 servers, use the Exchange Server 2010 role groups.

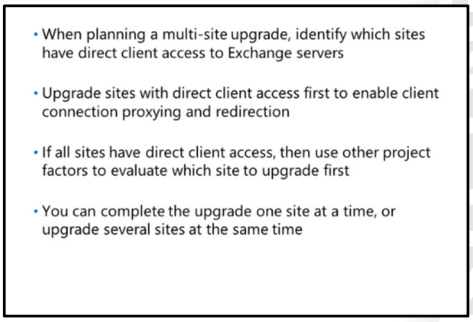
MCT USE ONLY. STUDENT USE PROHIBITED

You also can delegate permissions in an Exchange 2007 organization. The following table describes some options for creating an Exchange Server 2013 administrative design that emulates an Exchange Server 2007 design.

Exchange Server 2007 administrative option	Exchange Server 2010 equivalent
Assign users to the Exchange Organization Administrators group	Add users or groups to the Organization Management role group
Assign users to the Exchange View-Only Administrators group	Add users or groups to the View-Only Organization Management role group
Assign users to the Exchange Recipient Administrators group	Add users or groups to the Recipient Management role group
Assign users to the Exchange Public Folder Administrators group	Add users or groups to the Public Folder Management role group
Assign users as server administrators for a specific Exchange 2007 server	Create a custom role group that includes only server management roles and with a scope limited to a single server

Planning a Multisite Upgrade

If your organization has multiple AD DS sites that contain Exchange servers, it is important to consider the order in which you upgrade these sites. The critical factor in choosing the order in which to upgrade the Exchange servers in each site is whether the client can access the Exchange servers in each site directly from the Internet, or whether only some of the sites have direct Internet access.



Note: Important: As described earlier, when users connect to an Exchange 2013 Client Access server, the client connection can be proxied or redirected, depending on the location of the user mailbox and the Exchange server version that is hosting the user mailbox. If users connect to an Exchange 2007 or Exchange 2010 Client Access server, they will not be able to connect to a mailbox on an Exchange 2013 Mailbox server.

If only some of the Exchange Servers in the different AD DS sites are directly accessible from the Internet, then you must start the upgrade with one of the sites that is accessible from the Internet. By doing this, you can ensure that client connections will always connect first to an Exchange 2013 Client Access server, and then the connections can be proxied or redirected to other Exchange servers, including different versions of Exchange Server in different AD DS sites.

If the Exchange servers in all of the AD DS sites are directly accessible from the Internet, you will need to use other criteria for evaluating which Exchange servers to upgrade first. Some organizations may choose to upgrade the site with the most mailboxes first to get maximum benefit of the Exchange Server 2013 features. Other organizations may choose a smaller site as a pilot site before upgrading a larger site.

MCT USE ONLY. STUDENT USE PROHIBITED

When upgrading organizations with multiple sites, you also need to decide whether to complete the upgrade in one site first, and then move on to additional sites, or to upgrade multiple sites at a time. Either approach is possible as long as you ensure that all required Exchange server roles are deployed and retained in each site as long as they are required. Usually this decision is based on other project factors such as personnel availability or budget rather than technical requirements. Some organizations may have a team of administrators who move from one site to another deploying the servers, while another team manages the mailbox migration in multiple sites at a time. Other organizations may want to dedicate the team to completing the migration in one location before moving on to the next one.

MCT USE ONLY. STUDENT USE PROHIBITED

Lesson 2

Implementing the Upgrade from Previous Exchange Versions

Now that you understand how coexistence will work during the upgrade and you have completed the upgrade plan, you are ready to implement the upgrade. This lesson describes the steps that you must complete to upgrade from previous versions of Exchange Server to Exchange Server 2013.

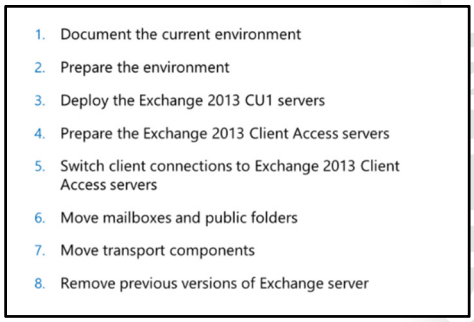
Lesson Objectives

After completing this module, you will be able to:

- Describe the steps required to upgrade to Exchange Server 2013.
- Upgrade client access to Exchange Server 2013.
- Upgrade message transport to Exchange Server 2013.
- Move user mailboxes to Exchange Server 2013.
- Migrate public folders to Exchange Server 2013.
- Remove previous versions of Exchange Server.

Exchange Server 2013 Upgrade Overview

Most Exchange Server upgrades will follow these high level steps.



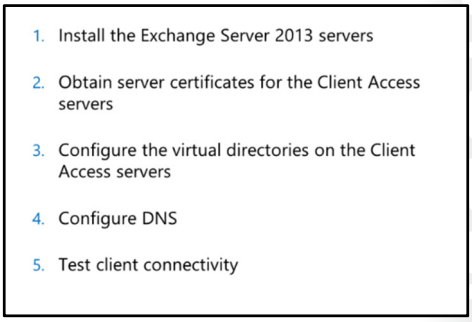
1. *Document the current environment*. It is important that you have a good understanding of the current environment before starting the upgrade. By documenting the current environment, you can make sure that you do not forget to upgrade all components to Exchange Server 2013.
2. *Prepare the environment*. Before starting the upgrade, you need to make sure that your organization meets all prerequisites for the upgrade. This requirement may mean that you need to install the required updates to Exchange Server 2007 or Exchange Server 2010. You will also need to prepare AD DS for the upgrade to Exchange Server 2013 CU1.
3. *Deploy the Exchange Server 2013 CU1 servers*. You can start deploying the Exchange Server 2013 servers without impacting the current environment. If you are deploying separate Client Access and Mailbox servers, it is not critical whether you deploy Mailbox servers or Client Access servers first. You need to deploy at least one Exchange server with each server role before you can move on.
4. *Prepare the Exchange 2013 Client Access servers*. To prepare the Exchange Client Access servers, you need to install the required server certificates, configure load balancing, and configure the virtual directories.

MCT USE ONLY. STUDENT USE PROHIBITED

5. *Switch client connections to the Exchange 2013 Client Access servers*. After you are confident that the Client Access servers are configured correctly, change DNS records and any appropriate firewall rules to configure all clients to connect to the Exchange 2013 Client Access servers.
6. *Move mailboxes and public folders*. Once you are sure that client connections are working correctly, you can start moving content to the Exchange 2013 Mailbox servers. The first step in this process is often creating the DAGs required for high availability. Then you can start moving mailboxes and public folder contents to the new Mailbox servers.
7. *Move transport components*. While you are moving mailbox servers, you can also start moving the message transport components from the previous version of Exchange Server to Exchange Server 2013.
8. *Remove previous versions of Exchange Server*. When you have removed all data and functionality from the Exchange 2007 or Exchange 2010 servers, you can remove these servers.

Process for Upgrading Client Access

During coexistence, you need to ensure that users with mailboxes on both Exchange Server 2007 Mailbox servers and Exchange Server 2010 Mailbox servers can access their mailboxes. At a high level, you will do this by preparing the Exchange 2013 Client Access servers and then configuring all client connections to start using these Client Access servers. The following steps describe how to enable this in more detail:



1. *Install the Exchange Server 2013 Client Access and Mailbox servers*. You can install the servers without disrupting any functionality in the existing environment.
2. *Obtain the required server certificates for the Exchange 2013 Client Access servers. You should request a certificate that supports at least the following subject alternative names:*
    - The primary URL for accessing the previous versions of Exchange Server for client access. For example, users may be accustomed to using a name such as mail.adatum.com whenever they need to configure any email clients. You should continue to use this name in Exchange Server 2013, so include this name in the certificate request.
    - The AutoDiscover server name. Normally, you would use a name such as autodiscover.adatum.com.
    - If you are upgrading from Exchange Server 2007, you should include an alternate name for the URL that will be used to redirect Outlook Web App clients to the Exchange 2007 Client Access server. For example, you might use a name such as legacy.contoso.com.
    - You can also include the Exchange 2013 Client Access server name in the certificate, but this is not recommended. In most cases, you will reconfigure both the internal and external URLs of the virtual directory to use a single name such as mail.adatum.com. This DNS name will resolve to the shared virtual IP address on a hardware load balancer or on a Network Load Balanced cluster. If you are not using a load balancing mechanism and you want to be able to connect to the servers using the server name, then include the server name in the certificate request.

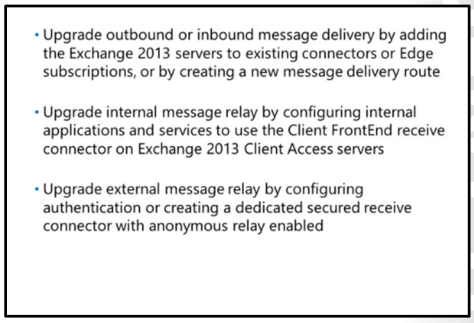
MCT USE ONLY. STUDENT USE PROHIBITED

Note: The Exchange Server 2013 Client Access server requires this certificate, but you also might install the same certificate on the Exchange 2007 Client Access server. Since the Exchange 2007 Client Access server will now need to accept connections to the legacy name, you need to ensure that this name is included in the certificate on the server.

3. *Configure the virtual directories on the Exchange 2013 Client Access servers. At a minimum, you should modify the external URLs for all virtual directories to use the shared client access name. If you did not include the server name in the certificate request, then change the internal URL and the AutoDiscoverServiceInternalUri to use the shared client access name as well.*
4. *Configure DNS. To configure DNS, you should:*
    - Create the legacy host record, which is legacy.contoso.com, in your external DNS infrastructure, and configure it to reference the Exchange Server 2007 Client Access server. Create or modify the host record for Autodiscover, which is Autodiscover.contoso.com, and configure it to reference the Exchange 2013 Client Access server. Create or modify the host record for the primary client access name and configure it to reference the Exchange Server 2013 Client Access server.
5. *Test all client scenarios, and ensure they function correctly. Use the Exchange Remote Connectivity Analyzer to test external connectivity.*

Process for Upgrading Message Transport

You may need to include several components when upgrading message transport from previous versions of Exchange Server to Exchange Server 2013.



- *Upgrading outbound Internet message delivery*. When upgrading outbound message delivery, you can either add Exchange Server 2013 Mailbox servers to the connectors currently used to deliver messages to the Internet, or you can configure a separate message routing mechanism. If you are currently using Exchange 2007 or Exchange 2010 Edge Transport servers, you can recreate the edge subscription so that the Exchange 2013 mailbox servers will route outbound mail to the Edge Transport servers. In addition, since Exchange Server 2013 SP1 now includes the Edge Transport role, you can upgrade your existing Edge server to Exchange Server 2013 SP1. If you have an SMTP send connector that routes email to a third-party SMTP gateway or directly to the Internet, you can add the Exchange 2013 Mailbox server as a source server on the send connector. Alternately, you can create a new send connector that uses only the Exchange 2013 Mailbox servers as source servers. While it is easier to add the Exchange 2013 servers to the existing connectors, creating new connectors for Exchange Server 2013 provides more options for managing outbound mail flow during testing and migration.
- *Upgrading inbound Internet message delivery*. If you have deployed Exchange 2007, Exchange 2010, or Exchange 2013 SP1 Edge Transport servers and you can recreate the edge subscription so that the Exchange 2013 mailbox servers are part of the edge subscription, inbound messages are routed to the Exchange 2013 Client Access servers. If you have a third-party SMTP gateway, you can configure the SMTP gateway to distribute the inbound email between the Exchange 2007 or Exchange 2010 Hub Transport servers and the Exchange 2013 Client Access servers. The Client FrontEnd receive

MCT USE ONLY. STUDENT USE PROHIBITED

connector on the Exchange 2013 Client Access servers is configured to accept anonymous SMTP connections for inbound messages.

- *Upgrading internal SMTP message relay*. Many organizations deploy internal applications and services that need to send SMTP messages to internal recipients. During a migration, you need to document which applications and services perform this function and modify these applications to start using the Exchange 2013 Client Access servers rather than the Exchange 2007 or Exchange 2010 Hub Transport servers for routing inbound email. If the applications and services are configured to use a DNS name as the SMTP server, you can complete the upgrade by changing the IP address for the server in DNS.
- *Upgrading external SMTP message relay*. In some cases, the internal applications or services might need to send email to Internet recipients. For example, an organization may have a website that needs to send SMTP mail to customers on the Internet. By default, this functionality is blocked on Exchange 2013 servers because it requires an open relay. The best way to configure this type of functionality is to configure the services and applications to use authentication when they try to send SMTP email. If the applications can authenticate to the receive connector on an Exchange 2013 Client Access or Mailbox server, the message can be delivered to the Internet.

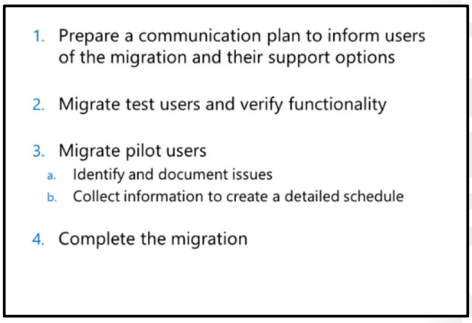
  In some cases, the applications cannot be configured to use authentication and you need to enable anonymous relay on a receive connector. Because this enables open relay, you should create a dedicated receive connector on a Client Access server and configure the connector to accept SMTP connections only from specified internal SMTP addresses. To enable the receive connector to allow anonymous users to relay to external email addresses, you need to run the following command in the Exchange Management Shell.

```
Get-ReceiveConnector "Receive Connector Name" | Add-ADPermission -User "NT
AUTHORITY\ANONYMOUS LOGON" -ExtendedRights "Ms-Exch-SMTP-Accept-Any-Recipient"
```

Note: Under no circumstances should SMTP servers from the Internet be able to relay anonymously through your Exchange servers.

Process for Moving the User Mailboxes

In most upgrades, the process of moving mailboxes will take up the majority of the migration time. After the Exchange Server 2013 servers have been deployed and the infrastructure prepared, organizations may still need to spend weeks or even months moving mailboxes from the previous version of Exchange Server to Exchange Server 2013. Use the following process when implementing the mailbox moves:

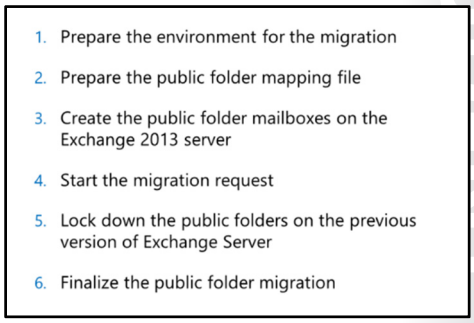


1. Prepare a communication plan for informing users that their mailboxes will be moved to a new environment. Most users should not experience any disruption in client functionality as a result of the migration, but it is important to let users know who they should contact if they do experience issues when their mailbox is moved. Normally, organizations provide a general communication to all users shortly before any mailboxes are moved, and then more specific communication to users a day or two before their mailbox is moved.

MCT USE ONLY. STUDENT USE PROHIBITED

2. After you have deployed the Exchange 2013 servers and verified functionality by using test accounts created on the Exchange 2013 servers, migrate a few test mailboxes to the Exchange 2013 servers. Use the test migration to verify the user experience during the mailbox move, and then verify message delivery and client access functionality after the mailboxes have been moved.
3. Migrate several groups of pilot users. Many organizations move the users in the IT departments first. In addition to technical users such as the IT users, it is also important to include a wide spectrum of users in the pilot migrations. This mix of users will provide more useful information about any issues that might be encountered during the move. Use the pilot migrations to:
    a. Identify and document any issues that users experience during the migration. Provide this information to the front-line service support team so that they can resolve as many issues as possible.
    b. During the pilot migration, you can identify the speed with which you can move mailboxes, and also the number of users that are likely to experience issues when their mailbox is moved. Use this information to create a detailed schedule for completing the mailbox moves.
4. Complete the mailbox moves.

Process for Migrating Public Folders



Because of the entirely new architecture for Exchange Server 2013 public folders, it is more complicated to migrate public folders from previous versions of Exchange Server than it was in older versions. To complete the migration, you must copy the contents of public folders from Exchange Server 2007 Service Pack 3 (SP3) Update Rollup 10 (RU10) or Exchange Server 2010 SP3 to the Exchange Server 2013 public folder mailboxes, and then switch all access to public folders to the new environment. Exchange Server 2013 provides several new ***PublicFolderMigrationRequest** cmdlets, in addition to two Windows® PowerShell® scripts, to help you complete the migration. These cmdlets use the Microsoft Exchange Mailbox Replication Service to perform the migration.

You can use the following high-level steps to complete the public folder migration from both Exchange Server 2007 and Exchange Server 2010.

1. Prepare the environment for the migration. To prepare the environment, perform the following steps:
    a. On the Exchange Server 2010 SP3 server, take a snapshot of the current public folder deployment. This snapshot is used to verify that the migration includes all the same folders, items, and permissions at the end of the migration. Use the **Get-PublicFolder**, **Get-PublicFolderStatistics**, and **Get-PublicFolderClientPermission** cmdlets to take this snapshot.
    b. On the Exchange Server 2010 SP3 server, verify that there is no previous record of a successful or ongoing migration.
    c. On the Exchange Server 2013 server, verify that there are no existing public folder migration requests. If any exist, clear them.
    d. Ensure that there are no existing public folders on the Exchange Server 2013 servers.
2. Prepare the public folder mapping file by performing the following steps:

MCT USE ONLY. STUDENT USE PROHIBITED

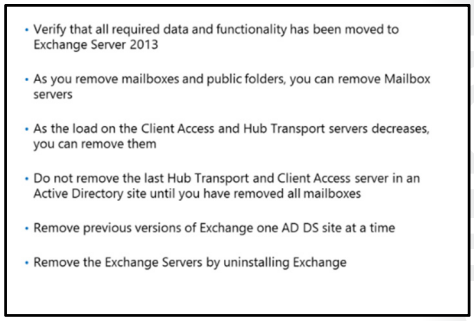
a. On the Exchange Server 2010 or Exchange Server 2007 server, generate the comma-separated values (CSV) file that lists all of the public folders on the previous Exchange Server versions. To do this, run the **Export-PublicFolderStatistics.ps1** script to create the mapping file that maps the folder name to the folder size. The file will have two columns: **FolderName** and **FolderSize**.

b. Create the Folder-to-Mailbox mapping file. This file will be used to create the correct number of public folder mailboxes on the Exchange 2013 Mailbox server. Run the **PublicFolderToMailboxMapGenerator.ps1** script to create the public folder-to-mailbox mapping file. You can edit the names of the public folder mailboxes that are listed in this file.

3. Create the public folder mailboxes on the Exchange 2013 server. Verify that the public folder mailboxes that you create match the name of the **TargetMailbox** in the mapping file. When you create the first public folder mailbox, use the **HoldForMigration** parameter.
4. Start the migration request. On an Exchange Server 2013 Mailbox server, run the **New-PublicFolderMigrationRequest** cmdlet to start the migration. This command can take a long time to complete if you have several gigabytes (GBs) or more of data in the public folders.
5. Lock down the public folders on the previous versions of Exchange Server for final migration. During the public folder migration, users are able to access public folders. To finish the migration, you must log users off of the public folders and lock them for a final synchronization. Run the **Set-OrganizationConfig -PublicFoldersLockedForMigration:$true** command on an Exchange Server 2010 SP3 server. If you have multiple public folder databases, wait until the public folder replication has completed to make sure that all public folder databases are locked.
6. Finalize the public folder migration. In the final step, run the **Set-PublicFolderMigration** cmdlet and set the **PreventCompletion** flag to **false**. Then resume the public folder migration. Exchange will now complete a final synchronization of the public folder contents and set the public folder databases on the Exchange Server2013 servers as active. After you complete the migration, all clients will need to access the public folders on the Exchange Server 2013 servers. If you experience issues with the migration, you can roll back to the previous version of Exchange Server by unlocking the public folders and setting the migration as **not completed**.

Note: For more detailed information on migrating public folders from a previous version of Exchange Server, see http://go.microsoft.com/fwlink/?LinkID=290962.

Removing Previous Versions of Exchange Server



Before removing the previous Exchange Server versions, you need to ensure that all functionality and data has been moved from the Exchange Servers to Exchange Server 2013. Verify that the following components have been upgraded to Exchange Server 2013:

- *Client access*. Verify that all clients are connecting to the Exchange 2013 Client Access servers. This may include POP3 or IMAP4 clients.
- *Transport connectors*. If you have added the Exchange 2013 servers to the existing SMTP connectors, remove the Exchange 2007 or Exchange 2010 Hub Transport servers. If you created a

MCT USE ONLY. STUDENT USE PROHIBITED

separate routing path using Exchange Server 2013, remove the previous connectors. Check message headers for inbound and outbound messages to verify that all messages are flowing through the Exchange 2013 servers.

- *Mailboxes*. Verify that there are no mailboxes left on the Exchange 2007 or Exchange 2010 Mailbox servers. In particular, ensure that all arbitration mailboxes have been moved. When you try to delete the mailbox databases on these servers, you will be blocked if any mailboxes remain on the server.
- *Public folders*. If you require public folders after the upgrade, complete the migration of the public folders before removing the Exchange servers. Verify that all data has been moved to the Exchange 2013 public folder mailboxes and that users can access the public folders on Exchange 2013.

Note: One way to validate that all required functionality has been removed from the previous versions of Exchange Server is to turn off the servers before removing them from the organization. If you can shut down the servers for several days without anyone noticing, it is likely safe to remove the server.

As you move data and functionality to the Exchange 2013 servers, you can begin removing the previous Exchange Server versions. You do not need to wait until all functionality has been moved to Exchange Server 2013 before you start removing the previous servers. Consider the following as you start removing Exchange servers from the organization:

- As you move mailboxes and public folders to Exchange Server 2013 Mailbox servers, you can start decommissioning the Exchange 2007 or Exchange 2010 Mailbox servers. If you want to reuse the hardware from the existing servers, you can move all mailboxes or public folders from one of the previous Mailbox servers and decommission that server.
- As the number of mailboxes on Exchange 2007 or Exchange 2010 servers decreases, you can also consider removing Hub Transport or Client Access servers. Since these server roles are only required for users with mailboxes on the previous Exchange servers, the load on these servers will decrease steadily as you move mailboxes.
- Do not remove the last Client Access server or Hub Transport server in an Active Directory site until you have moved all mailboxes and public folders to Exchange Server 2013. If you remove these server roles, users with mailboxes on Exchange 2007 or Exchange 2010 will not be able to access their mailboxes or send email.
- You can remove the previous Exchange Server versions one Active Directory site at a time. As you complete the migration in one Active Directory site, you can remove all previous versions of Exchange in that site before moving on to the next site.
- To remove the Exchange servers, uninstall Exchange. Do not just remove the Exchange Servers from the network as this will leave the objects related to the previous version of Exchange in AD DS.

.

MCT USE ONLY. STUDENT USE PROHIBITED

Lab: Upgrading from Exchange Server 2010 to Exchange Server 2013

Scenario

The A. Datum management team has decided to upgrade the Trey Research Exchange 2010 organization to Exchange Server 2013. The existing infrastructure consists of one server with Mailbox, Hub, and Transport server roles installed.

During the migration, all of the functionality currently deployed on the Exchange 2010 server needs to be moved to an Exchange 2013 server. Trey Research is currently using public folders, and it needs to migrate the public folder content to Exchange Server 2013. Because moving all of the mailboxes to Exchange Server 2013 will take several weeks, the two Exchange Server versions will need to co-exist during this time. Trey Research has deployed several retention policies and transport rules on the Exchange 2010 server that they want to migrate to the new Exchange Server.

Lab Setup

Estimated Time: 90 minutes

Virtual machines	**20342B-TREY-DC1** **20342B-TREY-EX1** **20342B-TREY-EX13**
User name	TreyResearch\Administrator
Password	Pa$$w0rd

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

1. On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In Microsoft Hyper-V® Manager, click **20342B-TREY-DC1**, and then, in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Log on by using the following credentials:
    a. User name: **TreyResearch\Administrator**
    b. Password: **Pa$$w0rd**
5. Repeat steps 2, 3, and 4 for **20342B-TREY-EX1** and **20342B-TREY-EX13**.

Exercise 1: Documenting the Exchange Server 2010 Organization

Scenario

The first step in the Exchange Server upgrade is to document the Exchange 2010 organization. This process will help in identifying any potential issues that may be encountered during the migration. The documentation can also be used as a check list to ensure that all current functionality has been moved to Exchange Server 2013 before removing the Exchange 2010 server.

The main tasks for this exercise are as follows:

1. Document the Exchange Organization configuration

MCT USE ONLY. STUDENT USE PROHIBITED

2. Document the Exchange Server configuration
3. Document the public folder configuration
4. Document the Exchange recipient configuration

▶ Task 1: Document the Exchange Organization configuration

1. On TREY-EX1, open the Exchange Management Console.
2. Document the following settings for Mailbox Database 1:
   a. Limits: Issue warning at (MB):
   b. Limits: Prohibit Send at (MB):
   c. Limits: Prohibit send and receive at (MB):
3. Document the TreyResearch Retention Policy Tags:
   - TreyResearch – Business Critical
     - Tag type:
     - Age limit:
     - Action:
   - TreyResearch – Default Delete
     - Tag type
     - Age limit:
     - Action:
   - TreyResearch – DefaultMovetoArchive
     - Tag type:
     - Age limit:
     - Action:
   - TreyResearch – Deleted Items
     - Tag type:
     - Age limit:
     - Action:
4. Document the TreyResearch Retention Policies:
   - Default Policy
     - Retention Policy Tags:
     - Mailbox Organizational Units:
   - Default Policy
     - Retention Policy Tags:
     - Organizational Units:
5. Document the Generation Server and Distribution Mechanism for the Offline Address Book.
6. Document the Password Change setting for Default and Executives Policy under Outlook Web App Mailbox Policies.
7. Document the password settings for the Executives EAS Policy under Exchange ActiveSync Mailbox Policies.
8. Under Organization Configuration, click **Hub Transport**.

MCT USE ONLY. STUDENT USE PROHIBITED

9. Document the settings for the E-Mail Disclaimer transport rule.
10. Document the settings for the Research Department Message Journaling rule.
11. Document the settings for the Internet Send Connector.

▶ Task 2: Document the Exchange Server configuration

1. Access the Exchange Server Configuration node.
2. Document the External URL and Authentication settings for the owa (Default Web Site) virtual directory.
3. Under Server Configuration, click **Hub Transport**.
4. Document the permission group configuration for the Default TREY-EX1 Receive Connector.

▶ Task 3: Document the public folder configuration

1. On TREY-EX1, open the Exchange Management Shell.
2. Run the following commands:
   - **Get-PublicFolder –recurse** and document the public folder structure.
   - **Get-PublicFolder –recurse | Get-PublicFolderClientPermission** and document the public folder client permissions for the IT, Research, and Sales public folders.
   - **Get-PublicFolderStatistics** and document the item count in the IT, Research, and Sales public folders.

▶ Task 4: Document the Exchange recipient configuration

1. In the Exchange Management Console, access the Recipient Configuration node.
2. Document the total number of mailboxes.
3. Document the following settings for Anders Riis, who is a member of the Executive team, Aaron Nicholls, who is a member of the Production team, and April Stewart, who is a member of the Research team:
   - Archive mailbox:
   - Retention policy: Default
   - Outlook Web App mailbox policy:
   - Exchange ActiveSync mailbox policy:
4. Document which user mailboxes have full access to the Research Journal Mailbox.
5. Document the groups to which the Mailbox Auditor belongs.
6. Document the litigation hold settings for the Kai Axford mailbox.

Results: After completing this exercise, you will have documented the Microsoft® Exchange Server 2010 organization.

MCT USE ONLY. STUDENT USE PROHIBITED

Exercise 2: Deploying Exchange Server 2013

Scenario

The second step in the upgrade is deploying the Exchange 2013 server in the existing organization. To do this, you need to prepare the AD DS environment, and then you need to install the Exchange 2013 server. Both the Client Access server role and the Mailbox server role will be deployed on the new server.

The main tasks for this exercise are as follows:

1. Preparing AD DS for the Exchange Server 2013 deployment
2. Install Exchange Server 2013
3. Verify a successful installation

▶ Task 1: Preparing AD DS for the Exchange Server 2013 deployment

1. On **TREY-EX13**, attach **D:\Program Files\Microsoft Learning\20342\Drives\ExchangeServer2013CU1.iso** to the virtual machine.
2. On TREY-EX13, from the desktop, open a Windows PowerShell window, and then use the **Install-WindowsFeature RSAT-ADDS** command to install the AD DS management tools.
3. Switch to D:\.
4. Execute the following command to prepare AD DS for your Exchange Server installation.

```
.\Setup.exe /PrepareAD /IAcceptExchangeServerLicenseTerms
```

▶ Task 2: Install Exchange Server 2013

1. On **TREY-EX13**, start Exchange Server setup. Choose the following options for the installation:
   - Do not check for updates.
   - Select the options to install both **Client Access** and **Mailbox Server** roles.
   - Do not disable malware protection.
   - Ensure that prerequisites are met.
   - Install the Exchange Server. Wait until the installation completes. It can take 30 to 40 minutes to finish.
   - On the **Setup Completed** page, click **finish.**
2. Restart TREY-EX13 and sign in as **TreyResearch\Administrator** with the password **Pa$$w0rd.**

▶ Task 3: Verify a successful installation

1. On TREY-EX13, from the Start screen, open the Exchange Management Shell.
2. Use the Exchange Management Shell to rename the default database to **EX13MDB1**.
3. Create a new mailbox with the following settings:
   a. Name EX13Test
   b. Password: Pa$$w0rd
   c. Organizational Unit: default
   d. Mailbox Database: EX13MDB1.

MCT USE ONLY. STUDENT USE PROHIBITED

4. Open Windows Internet Explorer®, in the Address bar, type **https://TREY-EX13.TreyResearch.net/owa**, and then press Enter.
5. Sign in as **TreyResearch\EX13Test** with the password **Pa$$w0rd**.
6. At the **Outlook Web App** page, click **save**. Verify that Outlook Web App opens.
7. Send a new message to Aaron Nicholls with a subject of "Test from Exchange 2013."
8. On TREY-EX1, open Internet Explorer and connect to **https://TREY-EX1.treyresearch.net/owa**.
9. Sign in as **Aaron** using the password **Pa$$w0rd**. Verify that the email from the EX13Test account is received in the inbox. Reply to the message.
10. On Trey-EX13, verify that EX13Test receives the reply from Aaron.
11. Close Internet Explorer.

Results: After completing this exercise, you will have deployed an Exchange 2013 server in the Trey Research Exchange organization.

Exercise 3: Upgrading from Exchange Server 2010 to Exchange Server 2013

Scenario

Now that you have deployed the Exchange 2013 server and verified basic functionality, you are ready to start moving all data and functionality from Exchange Server 2010 to Exchange Server 2013. First, you will move the administrator account to Exchange Server 2013 so that you can start using the Exchange Administration Console to manage Exchange. You will then install a certificate on the Exchange Server 2013 server and configure client connections to start using the Exchange 2013 Client Access server. Next, you will move a pilot group of mailboxes to the new environment and verify functionality. Then you will migrate all mailboxes to Exchange Server 2013. When you complete the mailbox migration, you will complete a public folder migration and move other functionality to the Exchange Server 2013 server.

The main tasks for this exercise are as follows:

1. Move the Administrator mailbox to Exchange Server 2013
2. Obtain a certificate for the Exchange 2013 server
3. Change the Client Access configuration to use Exchange Server 2013
4. Move the IT department mailboxes to Exchange Server 2013
5. Configure and verify message routing
6. Move the remaining mailboxes to Exchange Server 2013
7. Migrate public folders to Exchange Server 2013
8. Verifying the upgrade of compliance features
9. Verifying additional upgrade components (optional)

▶ Task 1: Move the Administrator mailbox to Exchange Server 2013

1. On Trey-EX13, in the Exchange Management Shell, use the **New-MoveRequest** command to move the Administrator mailbox to the EX13MDB1 mailbox. Use the **Get-MoveRequest** command to verify when the move is complete.

MCT USE ONLY. STUDENT USE PROHIBITED

2. After the move completes, connect to the EAC and sign in as **TreyResearch\Administrator** using the password **Pa$$w0rd**.
3. Verify that the Administrator can now access the EAC.
4. Connect to Outlook Web App and send a message to Aaron.
5. On TREY-EX1, connect to Outlook Web App, sign in as **TreyResearch\Aaron** using the password **Pa$$w0rd**, and verify that Aaron receives the message.

▶ Task 2: Obtain a certificate for the Exchange 2013 server

1. On TREY-EX13, connect to **https://TREY-EX13.TreyResearch.net/ecp**.
2. Sign in as **TreyResearch\administrator** with the password **Pa$$w0rd**.
3. Click the **servers** node, click on **Certificates**, and start the wizard for creating a new certificate request.
4. Provide **mail.TreyResearch.net** as a friendly name for a certificate.
5. Do not use wildcard certificates.
6. Provide the name **mail.TreyResearch.net** for all values that are not defined.
7. Ensure that the certificate request contains the following domain names: **mail.TreyResearch.net**, **TREY-EX13.TreyResearch.net**, **AutoDiscover.TreyResearch.net, TREY-EX13**, and **TreyResearch.net**.
8. Fill in additional data as follows:
    a. Organization name: **A.Datum**
    b. Department name: **IT**
    c. City/Locality: **London**
    d. State/Province: **England**
    e. Country/Region name: **United Kingdom**
9. Save certificate request to **\\TREY-EX13\C$\users\administrator.treyresearch\downloads\certreq.req**.
10. Open the certificate request file with Notepad, and copy all content to the clipboard.
11. Connect to **http://TREY-DC1.TreyResearch.net/certsrv** as **TreyResearch\Administrator** with the password **Pa$$w0rd**.
12. The browser displays a message that it does not support the generation of certificate requests. Press F12.
13. In the **Browser Mode** drop down list, click **Internet Explorer 10 Compatibility View**. Close the bottom tab.
14. Choose to perform an advanced certificate request.
15. Paste the certificate request content in to the appropriate field, and select Trey Web template.
16. Save the certificate.
17. On TREY-EX13, open the EAC.
18. Complete the **Mail.TreyResearch.net** Exchange certificate request using the **\\TREY-EX13\C$\users\administrator.treyresearch\downloads\certnew.cer** file.

MCT USE ONLY. STUDENT USE PROHIBITED

19. Assign the certificate to IIS service.

▶ Task 3: Change the Client Access configuration to use Exchange Server 2013

1. On Trey-EX13, in the EAC, configure the external URL for the following virtual directories to use the mail.treyresearch.net server name.
   a. owa (Default Web Site)
   b. ecp (Default Web Site)
   c. EWS (Default Web Site)
   d. Microsoft-Server-ActiveSync (Default Web Site)
   e. OAB (Default Web Site)
2. From the Exchange Management Shell, reset IIS.
3. On TREY-DC1, in DNS, change the IP address for mail.treyresearch.net to use the IP address 172.16.0.102.
4. Clear the cache in DNS.
5. On TREY-EX13, use the nslookup command to verify that the host name mail.treyresearch.net is assigned the IP address 172.16.0.102.
6. Clear the local DNS resolver cache.
7. Open Internet Explorer and connect to https://mail.treyresearch.net/owa. Sign in as TreyResearch\Administrator using the password **Pa$$w0rd**. Verify that the Administrator can access the Exchange 2013 mailbox. Close Internet Explorer.
8. Open Internet Explorer and connect to **https://mail.treyresearch.net/owa**.
9. Sign in as **TreyResearch\Aaron** using the password **Pa$$w0rd**. Verify that the Aaron can access his Exchange 2010 mailbox. Close Internet Explorer.

▶ Task 4: Move the IT department mailboxes to Exchange Server 2013

1. Use the **get-mailbox –organizationalunit IT | New-MoveRequest –TargetDatabase EX13MDB1** command to move all IT mailboxes to the Exchange 2013 server.

▶ Task 5: Configure and verify message routing

1. Sign into Outlook Web App as **TreyResearch\Administrator** using the password **Pa$$w0rd**.
2. Create a new mail for Kai Axford with the subject **Message before migration**.
3. Connect to the EAC. In the left pane, click on **mail flow**, and then click **delivery reports**.
4. Track the message that you just sent to Kai's mailbox by tracking the message from the Administrator's mailbox.
5. Track the message that you just sent to Kai's mailbox by tracking the message from Kai's mailbox.
6. Modify the Internet Send Connector to add TREY-EX13 as a source server.
7. Verify that the Default Frontend Trey-EX13 receive connector is configured to accept SMTP connections from anonymous users.

▶ Task 6: Move the remaining mailboxes to Exchange Server 2013

1. Use the **Get-MoveRequest** cmdlet to verify that all mailboxes have been successfully moved. If the accounts show Queued, wait for 5 minutes and rerun the command until it shows Completed.

MCT USE ONLY. STUDENT USE PROHIBITED

2. Run the **Get-Mailbox –Arbitration -Database "Mailbox Database 1" | New-MoveRequest –TargetDatabase EX13MDB1** command to move the arbitration mailboxes to Exchange Server 2013.
3. In the EAC, on the **migration** tab, create a new local move request. Include all mailboxes that are still on Mailbox Database 1.
4. Name the migration batch **CompleteMigration**, and use the **EX13MDB1** mailbox database as the destination for both the mailbox and archive mailboxes.
5. Start the migration. The migration will take some time to complete, so continue with the following tasks.

▶ Task 7: Migrate public folders to Exchange Server 2013

1. On TREY-EX13, in the Exchange Management Shell, use the **Get-PublicFolderMigrationRequest** to verify that there are no outstanding public folder migration requests.
2. Use the **Get-Mailbox -PublicFolder** to verify that there are no public folder mailboxes on the Exchange 2013 Mailbox server.
3. On TREY-EX1, create a folder named Migration on the C drive. Share the folder with the default permissions.
4. Copy the following files from \\TREY-EX13\C$\Program Files\Microsoft\Exchange Server\v15\Scripts to the C:\Migration folder on TREY-EX1:
   - Export-PublicFolderStatistics.ps1
   - Export-PublicFolderStatistics.strings.psd1
   - PublicFolderToMailboxMapGenerator.ps1
   - PublicFolderToMailboxMapGenerator.strings.psd1
5. Open the Exchange Management Shell, change to the Migration folder path, and then type **.\Export-PublicFolderStatistics.ps1 PFStats.csv TREY-EX1**, and press Enter. This command exports the public folder statistics to a .CSV file.
6. In the **C:\Migration** folder, open the **PFStats.csv** file in Notepad. Review the information and close the file.
7. In the Exchange Management Shell, type **.\PublicFolderToMailboxMapGenerator.ps1 2000 C:\Migration\PFStats.csv PFtoMBXMap.csv**, and then press Enter.



Note: The value "2000" in the previous command specifies the maximum public folder mailbox size in bytes planned for the Exchange Server 2013 environment. This number does not set a limit on the mailbox size; it is only a value used by the script to determine how many public folder mailboxes will be required. In a production environment, this value would be much larger. The smaller number is used here so that the script will require more than one public folder mailbox on Exchange Server 2013.

8. Open the PFtoMBXMap.csv file in Notepad.
9. Edit the target mailbox names by adding a PF to the mailbox name. For example, Mailbox1 should be changed to PFMailbox1. After changing all three mailbox names, save and close the file.
10. On TREY-EX13, in the Exchange Management Shell, type **New-Mailbox –PublicFolder PFMailbox1 –HoldForMigration** and press Enter.
11. Type **New-Mailbox –PublicFolder PFMailbox2** and press Enter.

MCT USE ONLY. STUDENT USE PROHIBITED

12. Type **New-Mailbox –PublicFolder PFMailbox3** and press Enter.
13. Type **New-PublicFolderMigrationRequest –SourceDatabase (Get-PublicFolderDatabase –Server TREY-EX1) –CSVData (Get-Content \\TREY-EX1\Migration\PFToMBXMap.csv -Encoding Byte)** and press Enter.
14. Wait a minute, and then type **Get-PublicFolderMigrationRequest | Get-PublicFolderMigrationRequestStatistics** and press Enter. Verify that the StatusDetail is displayed as AutoSuspended, and PercentComplete is set to 95. If these values are not displayed, wait another minute and then run the command again.
15. Type **Set-OrganizationConfig -PublicFoldersLockedForMigration:$true** and press Enter.
16. Type **Set-PublicFolderMigrationRequest -Identity \PublicFolderMigration -PreventCompletion:$false** and press Enter.
17. Type **Resume-PublicFolderMigrationRequest -Identity \PublicFolderMigration** and press Enter.
18. This request can take several minutes to finish. You can continue with the next steps while the migration finishes.

▶ Task 8: Verifying the upgrade of compliance features

1. On TREY-EX13, in the EAC, verify that all mailboxes have been migrated to Exchange Server 2013.
2. Connect to Outlook Web App and sign in as **TreyResearch\Administrator**.
3. Create a new mail for Kai Axford with the subject heading **Message after migration**.
4. On TREY-EX1, connect to **https://mail.treyresearch.net/owa** and sign in as **Treyresearch\Kai** using the password **Pa$$w0rd**.
5. Verify that the message from the Administrator arrived and that it includes the email disclaimer configured by the transport rule configured in Exchange Server 2010.
6. Delete both messages that you sent from the Administrator.
7. Empty the deleted items folder.
8. Purge the messages from the deleted items folder.
9. On TREY-EX13 use the **Get-mailbox Discover* | FL Hidden*** command to verify that the DiscoverySearchMailbox is hidden from the address lists.
10. Use the **Set-mailbox Discover* -HiddenFromAddressListsEnabled $false** command to unhide the mailbox. This step is required so that the Mailbox Auditor can open the DiscoverySearchMailbox from Outlook Web App.
11. On TREY-EX1, open Internet Explorer and connect to **https://mail.treyresearch.net/owa**.
12. Sign in as Treyresearch\MailboxAuditor using the password **Pa$$w0rd**.
13. Open the Research Journal Mailbox, and verify that the two messages sent to Kai Axford are listed. Kai is a member of the Research group, and the messages sent to any member of the Research group are journaled to this mailbox.
14. On TREY-EX1, connect to the EAC as Mailbox Auditor.
15. Click **compliance management**.
16. Create a new eDiscovery search named **Search Kai's mailbox** that will return all messages in Kai's mailbox.
17. Copy the search results to the **Discovery Search Mailbox**.

MCT USE ONLY. STUDENT USE PROHIBITED

18. Open the Discover Search Mailbox and verify that the two messages purged by Kai are in the **Purges** folder. Kai's mailbox was placed on Litigation Hold in Exchange Server 2010, and the hold and all saved messages were retained during the migration.

▶ Task 9: Verifying additional upgrade components (optional)

1. Review the configuration options that you documented in Exercise 1. If time permits, verify that all of the options have been upgraded to Exchange Server 2013.

Results: After completing this exercise, you will have completed the upgrade of all data and functionality to the Exchange 2013 server.

Exercise 4: Removing Exchange Server 2010

Scenario

Now that you have completed the upgrade to Exchange Server 2013, the final step is to remove the Exchange Server 2010 server from the organization. Before you uninstall the server, you will verify that all functionality has been migrated to Exchange Server 2013. Once you have verified the successful migration, you will uninstall Exchange Server 2010.

The main tasks for this exercise are as follows:

1. Removing Exchange Server components

2. Remove the Exchange 2010 server

3. Shut down the virtual machines

▶ Task 1: Removing Exchange Server components

1. On TREY-EX1, in the Exchange Management Console, remove the **Offline Address Book** located on Exchange Server 2010. When Exchange Server 2013 is installed, a new Offline Address Book is created that is only distributed through the web-based mechanism.
2. On TREY-EX13, in the Exchange Management Shell, use the **Get-PublicFolderMigrationRequest | Get-PublicFolderMigrationRequestStatistics** command to verify that the public folder migration is complete.
3. If the public folder migration is not complete use the **Suspend-PublicFolderMigration** and **Resume-PublicFolderMigrationRequest** commands to pause and restart the migration request. Wait a few minutes, and then verify that the public folder migration is complete.
4. On TREY-EX1, open the Exchange Management Shell, use the **Get-Mailbox** command to verify that no regular mailboxes and no arbitration mailboxes are listed on the server.
5. Use the **Get-PublicFolder** to verify that the public folders on TREY-EX1 were no longer available.
6. On TREY-EX1, in the Exchange Management Console, remove TREY-EX1 as a source server on the **Internet Send Connector**.
7. In the Exchange Management Console, dismount and delete the Mailbox Database 1 and Public Folder Database 1.
8. If you get an error message that the public folder still contains public folders, use Active Directory Services Interfaces Editor (ADSI Edit) to delete the public folder database from **CN=Configuration [TREY-DC1.TreyResearch.Net]**, **CN=Configuration,DC=TreyResearch,DC=net**, **CN=Services**,

MCT USE ONLY. STUDENT USE PROHIBITED

CN=Microsoft Exchange, **CN=TreyResearchOrg**, **CN=Administrative Groups, CN=Exchange Administrative Group (FYDIBOHF23SPDLT)**, and **CN=Databases.**

▶ Task 2: Remove the Exchange 2010 server

1. On TREY-EX1, close the Exchange Management Console and the Exchange Management Shell.
2. In Control Panel, uninstall Microsoft Exchange Server 2010.

▶ Task 3: Shut down the virtual machines

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

1. On the host computer, start Microsoft Hyper-V® Manager.
2. In the Virtual Machines list, right-click **20342B-TREY-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20342B-TREY-EX1** and **20342B-TREY-EX13**.

Results: After completing this exercise, you will have removed Exchange Server 2010 from the Exchange organization.

Question: When you changed the Domain Name System (DNS) settings for Mail.TreyResearch.net to point to TREY-EX13, how could users access their mailboxes on TREY-EX1 while using Microsoft Outlook® Web App?

Question: What would happen if you did not migrate all mailboxes to Exchange Server 2013 before migrating the public folders?

MCT USE ONLY. STUDENT USE PROHIBITED

Module Review and Takeaways

Best Practices

- Before attempting an Exchange Server upgrade in a production environment, verify the migration path in a test lab. This is particularly important in a complex Exchange environment that contains multiple sites and domains.

- Ensure that all email clients have been upgraded in your organization before moving users' mailboxes to Exchange Server 2013. Only Outlook 2007 SP3 with the Outlook 2007 November 2012 update (12.0.6665.5000) or newer clients are supported. Apple Macintosh clients must be upgraded to Microsoft Entourage® 2008 for Macintosh, Web Services Edition, or Outlook for Mac 2011.

- Always upgrade the Exchange Servers in Internet-facing sites before upgrading internal sites. This will enable the Exchange 2013 Client Access servers in the Internet-facing sites to proxy client requests to previous versions of Exchange Server in the internal sites.

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
You are upgrading your Exchange Server 2007 organization to Exchange Server 2013, and you have configured Client Access servers for Internet access. Users with mailboxes on Exchange Server 2013 Mailbox servers can access their mailbox using Outlook Web App from the Internet, but users with mailboxes on the Exchange Server 2007 Mailbox servers cannot.	
When you try to remove the public folder database in Exchange Server 2007 or Exchange Server 2010, you are prevented from doing so. The error states that the database still contains public folders.	

Review Question(s)

Question: Why do you need to use a legacy name for Exchange Server 2007 Client Access servers when you deploy Exchange Server 2013 Client Access servers?

Question: Your organization includes two locations and Active Directory sites. You have deployed Exchange Server 2010 servers in both sites. You now are deploying Exchange Server 2013 servers in one of the sites and removing the Exchange Server 2010 servers. When can you remove the last Exchange 2010 Hub Transport server in the site?

MCT USE ONLY. STUDENT USE PROHIBITED

Course Evaluation

Include this slide only in the last module of the Course.

<insert slide here>

Keep this evaluation topic page if this is the final module in this course. Insert the Product_Evaluation.ppt on this page.

If this is not the final module in the course, delete this page

Your evaluation of this course will help Microsoft understand the quality of your learning experience.

Please work with your training provider to access the course evaluation form.

Microsoft will keep your answers to this survey private and confidential and will use your responses to improve your future learning experience. Your open and honest feedback is valuable and appreciated.

MCT USE ONLY. STUDENT USE PROHIBITED

Module 1: Designing and Implementing Site Resilience

Lab: Designing and Implementing Site Resiliency

Exercise 1: Implementing Site Resilience

▶ Task 1: Add DNS Entries for LON-CAS1 and LON-CAS2 to the webmail.adatum.com and autodiscover.adatum.com DNS A Records

1. On LON-DC1, open **Server Manager**, click **Tools**, and then click **DNS**.
2. In Domain Name System (DNS) Manager, in the left navigation pane, expand **Forward Lookup Zones**, select and then right-click **Adatum.com**, and then click **New Host (A or AAAA)**.
3. In the **New Host** dialog box, in **Name** field type **webmail**, in the **IP address** field type **172.16.0.221**, and then click **Add Host**. Click **OK**.
4. In the **New Host** dialog box, in the **Name** field, type **autodiscover**, in the **IP address** field type **172.16.0.221**, and then click **Add Host**. Click **OK**.
5. In the **New Host** dialog box, in the **Name** field, type **webmail**, in the **IP address** field, type **172.16.0.20**, and then click **Add Host**. Click **OK**. Click **Done**.
6. In the left navigation pane, expand **Forward Lookup Zones**, and then select **Adatum.com**.
7. Verify that the new records are listed in the results pane.
8. Close the DNS Manager.

▶ Task 2: Configure the Client Access virtual directories

1. Switch to the **LON-CAS1** virtual machine. Open **Windows® Internet Explorer®**, type **https://webmail.adatum.com/ecp**, and then press Enter.
2. In the Exchange Admin Center, sign in as **Adatum\administrator** with the password **Pa$$w0rd**.
3. In the Exchange admin center, in the Feature pane, click **servers**.
4. In the virtual directories tab, click **configure external access domain**.
5. In the configure external access domain window, click **Add,** and then select **LON-CAS1** and **LON-CAS2**. Click **add,** and then click **ok**.
6. Type **webmail.adatum.com** in the **Enter the domain name you will use with your external Client Access servers** field. Click **save**.
7. Click **close**.
8. Click to the **Start** screen, and then click **Exchange Management Shell**.
9. In the Exchange Management Shell, type **Set-ClientAccessServer LON-CAS1 –AutoDiscoverServiceInternalUri https://autodiscover.adatum.com/Autodiscover/Autodiscover.xml**, and then press Enter.
10. In the Exchange Management Shell, type **Set-ClientAccessServer LON-CAS2 –AutoDiscoverServiceInternalUri https://autodiscover.adatum.com/Autodiscover/Autodiscover.xml**, and then press Enter.

MCT USE ONLY. STUDENT USE PROHIBITED

11. Type **Get-OutlookAnywhere | Set-OutlookAnywhere –ExternalHostname webmail.adatum.com –ExternalClientAuthenticationMethod negotiate –ExternalClientsRequireSsl $true –InternalHostname webmail.adatum.com –InternalClientsRequireSsl $true**, and then press Enter.

▶ Task 3: Prepare the cluster network object for a database availability group (DAG)

1. On LON-DC1, open **Server Manager**, click **Tools**, and then click **Active Directory Users and Computers**.
2. In **Active Directory Users and Computers**, on the menu bar, click **View**, and then click **Advanced Features**.
3. In the navigation pane on the left, expand **Adatum.com**, click **Computers**, right-click **Computers**, point to **New**, and then click **Computer**.
4. In the **New Object – Computer** dialog box, in the **Computer name** field, type **DAG1**, and then click **OK**.
5. In the details pane on the right, right-click **DAG1**, and then click **Properties**.
6. In the **DAG1 Properties** dialog box, click the **Security** tab.
7. On the **Security** tab, click **Add**, and then, in **Enter the object names to select,** type **Exchange Trusted Subsystem**. Click **Check Names**, and then click **OK**.
8. On the **Security** tab, click **Add**, and then click **Object Types**.
9. In the **Object Types** dialog box, select the **Computers** check box, and then click **OK**.
10. In the Select Users, Computers, Service Accounts, or Groups window, in **Enter the object names to select** field, type **LON-MBX1;LON-MBX2**, click **Check Names**, and then click **OK**.
11. On the **Security** tab, select **LON-MBX1 (ADATUM\LON-MBX1$)**, and then, in the **Allow** column, in the **Permissions for LON-MBX1** list, select the **Full control** check box.
12. On the **Security** tab, select **LON-MBX2 (ADATUM\LON-MBX2$)**, and then, in the **Allow** column, in the **Permissions for LON-MBX2** list, select the **Full control** check box.
13. On the **Security** tab, select **Exchange Trusted Subsystem (ADATUM\Exchange Trusted Subsystem)**, in the **Allow** column, in the **Permissions for Exchange Trusted Subsystem** list, click **Full control**, and then click **OK**.
14. In the Active Directory Users and Computers window, in the details pane on the right, right-click **DAG1**, and then click **Disable Account**.
15. In the warning window, click **Yes**, and then, on the next information window, click **OK**.

▶ Task 4: Create a DAG and add Mailbox servers to it

1. Switch to **LON-CAS1**.
2. If necessary, open **Internet Explorer**, and type **https://webmail.adatum.com/ecp**, and then press Enter.
3. Sign in as **Adatum\administrator** with the password **Pa$$w0rd**.
4. In the Exchange admin center, in the **Feature** pane, click **servers**.
5. On tabs, click **database availability groups**, and then, on the toolbar, click **New**.
6. In the new database availability group window, in the **Database availability group name** field, type **DAG1**, click **Witness server**, and then type **LON-CAS1** in the **Witness server** field.
7. Click **Witness directory**, in the **Witness directory** field, type **C:\FSWDAG1**, click **Enter an IP address**, in **Database availability group IP addresses** field, and then type **172.16.0.33**.

MCT USE ONLY. STUDENT USE PROHIBITED

8. Click **Add**. In the **Database availability group IP addresses** field, type **172.16.0.225**. Click **Add**, and then click **save**.
9. In the list view, click **DAG1**, and on the toolbar click **Manage DAG Membership**.
10. Click **Add**, and then select **LON-MBX1** and **LON-MBX2**. Click **add**, and then click **ok**.
11. Click **save**. Wait until the task completes successfully.
12. Click **close**. Wait for 2-3 minutes before continuing with the next task.

▶ Task 5: Add a copy of the Mailbox database on LON-MBX2

1. In Exchange admin center, in the navigation pane, click **servers**. In the tabs click **databases**, click **Mailbox Database 1**, click **More**, and then click **Add database copy**.
2. In the Add Mailbox Database Copy window, click **Browse**.
3. In the Select Server window, click **LON-MBX2**, and then click **ok**.
4. In the Add Mailbox Database Copy window, click **save**.
5. Wait until the saving finishes, and then click **close**.

▶ Task 6: Verify the successful copying of a database

1. In tabs, click **Refresh**, and wait until the details pane shows **Mailbox Database 1\LON-MBX2** as **Passive Healthy**. This might take several minutes and up to several hours depending on the size of the database.
2. In the details pane, under **Mailbox Database 1\LON-MBX2**, click **View details**.
3. Make sure that the **Status** displays **Healthy** and the **Content index state** also displays **Healthy**. Then click **cancel**.
4. Close Internet Explorer.

▶ Task 7: Configure an alternate file share witness and configure Datacenter Activation Mode

1. On the **Start** screen, click **Exchange Management Shell**.
2. In Exchange Management Shell, type **Set-DatabaseAvailabilityGroup DAG1 –AlternateWitnessServer LON-CAS2 –AlternateWitnessDirectory C:\FSWDAG1**, and then press Enter.
3. In Exchange Management Shell, type **Set-DatabaseAvailabilityGroup DAG1 –DatacenterActivationMode DAGOnly**, and then press Enter.
4. Type **Restart-Computer –ComputerName LON-CAS2 –Force**, and then press Enter.

MCT USE ONLY. STUDENT USE PROHIBITED

5. Type **Restart-Computer –ComputerName LON-CAS1 –Force**, and then press Enter.

Results: After completing this exercise, you will have successfully:

Added DNS entries for LON-CAS2 to the webmail.adatum.com and autodiscover.adatum.com host (A) resource records.

Configured virtual directories on LON-CAS2.

Prestaged the cluster network object for a DAG.

Created a DAG and added mailbox servers to it.

Added a copy of the mailbox database on LON-MBX2.

Configured LON-CAS2 as an alternate witness server.

MCT USE ONLY. STUDENT USE PROHIBITED

Exercise 2: Validating Site Resilience

▶ Task 1: Verify the location of the active Mailbox copy

1. Log on to **LON-CAS1** by using the following credentials:
   - User name: **Adatum\Administrator**
   - Password: **Pa$$w0rd**
2. On LON-CAS1, on the **Start** screen, click **Exchange Management Shell**.
3. In the Exchange Management Shell, type **Get-MailboxDatabase -Status | Format-Table Name, MountedOnServer**, and then press Enter.
4. If **Mailbox Database 1** is mounted on **LON-MBX1.adatum.com**, continue to the next task. If the database is mounted on **LON-MBX2.adatum.com**, type **Move-ActiveMailboxDatabase "Mailbox Database 1" –ActivateOnServer LON-MBX1**, and then press Enter. Press **Y** to complete the move.

▶ Task 2: Verify the normal functionality of Outlook

1. Switch to **LON-CL1**.
2. From the **Start** screen, right-click the screen, and then click **All apps**. Click **Outlook 2013**.
3. Create an Outlook profile. On the **Welcome to Microsoft® Outlook® 2013** screen, click **Next**.
4. Click **Next**.
5. Click **Next**.
6. Click **Finish** to create the profile.
7. On the **First things first** screen, select **Ask me later,** and then click **Accept**.
8. Click **New Email**.
9. In the **To** field, type **amr@adatum.com**.
10. In the **Subject** field, type **Test**.
11. In the message body, type **This is a test message**, and then click **Send**.
12. Verify that the message is not in the Outbox folder.
13. Close Outlook.

▶ Task 3: Initiate a failure of the active Mailbox copy on LON-MBX1 and verify Outlook functionality

1. On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In Microsoft Hyper-V® Manager, click **20342B-LON-MBX1**, and then, in the Actions pane, click **Shut Down**. Click **Shut Down** and then wait until the server powers down.
3. Switch to the **LON-CL1** virtual machine.
4. From the **Start** screen, right-click the screen, and then click **All apps**. Click **Outlook 2013**.
5. If you are prompted to **Allow this website to configure Kim@adatum.com server settings?** click **Allow**.
6. After a minute or so, verify that you are connected to Microsoft Exchange and then click **New Email**.
7. In the **To** field, type **amr@adatum.com**.
8. In the **Subject** field, type **Test After LON-MBX1 Powered Down**.

MCT USE ONLY. STUDENT USE PROHIBITED

9. In the message body, type **This is a test message**, and then click **Send**.
10. Verify that the message is not in the Outbox folder.
11. Close Outlook.

▶ Task 4: Initiate a failure of the Client Access services on LON-CAS1

1. On LON-CAS1, on the **Start** screen, click **Exchange Management Shell**.
2. In the Exchange Management Shell, type **Stop-Service MSExchangeHM**, and then press Enter.
3. In the Exchange Management Shell, type **Stop-Service W3SVC**, and then press Enter.
4. Switch to the **LON-CL1** virtual machine.
5. From the **Start** screen, right-click the screen, and then click **All apps**. Click **Outlook 2013**.
6. If you are prompted to **Allow this website to configure Kim@adatum.com server settings?** click **Allow**.
7. Click **New Email**.
8. In the **To** field, type **amr@adatum.com**.
9. In the **Subject** field, type **Test After LON-CAS1 IIS Stopped**.
10. In the message body, type **This is a test message**, and then click **Send**.
11. Verify that the message is not in the Outbox folder.
12. Close Outlook.

▶ Task 5: Initiate a failure of the witness server, and test Outlook functionality

1. On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In Hyper-V Manager, click **20342B-LON-CAS1**, and then, in the Actions pane, click **Shut Down**. Click **Shut Down** and then wait until the server powers down.
3. Switch to the **LON-CL1** virtual machine.
4. From the **Start** screen, right-click the screen, and then click **All apps**. Click **Outlook 2013**.
5. Verify that Outlook is disconnected. Note that it may take a couple of minutes for the disconnected state to show.
6. Close Outlook.

▶ Task 6: Recover the DAG in the secondary site, and verify Outlook functionality

1. On the LON-MBX2 virtual machine, on the **Start** screen, click **Exchange Management Shell**.
2. In the Exchange Management Shell, type **Stop-DatabaseAvailabilityGroup DAG1 –ActiveDirectorySite London**, and then press Enter. Type **Y**, and then press Enter to confirm.
3. Type **Stop-Service clussvc** and then press Enter.
4. Type **Restore-DatabaseAvailabilityGroup DAG1 –ActiveDirectorySite Swindon**, and then press Enter. Type **Y**, and then press Enter to confirm.
5. Type **Start-DatabaseAvailabilityGroup DAG1 –ActiveDirectorySite Swindon**, and then press Enter.
6. Switch to the **LON-CL1** virtual machine.
7. From the **Start** screen, right-click the screen, and then click **All apps**. Click **Outlook 2013**.
8. Click **New Email**.

MCT USE ONLY. STUDENT USE PROHIBITED

9. In the **To** field, type **amr@adatum.com**.
10. In the **Subject** field, type **Test After DAC**.
11. In the message body, type **This is a test message**, and then click **Send**.
12. Verify that the message is not in the Outbox folder.

▶ Task 7: To prepare for the next module

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

1. On the host computer, start **Hyper-V Manager**.
2. In the **Virtual Machines** list, right-click **20342B-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20342B-LON-CAS1**, **20342B-LON-CAS2**, **20342B-LON-MBX1**, **20342B-LON-MBX2**, and **20342B-LON-CL1**.
5. In Hyper-V Manager, click **20342B-LON-DC1**, and in the Actions pane, click **Start**.
6. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
7. Sign in using the following credentials:
   - User name: **Adatum\Administrator**
   - Password: **Pa$$w0rd**
8. Repeat steps 5 to 7 for **20342B-LON-CL1**.

Results: After completing this exercise, you will have successfully:

Verified the location of the active mailbox copy.

Verified the normal functionality of Outlook.

Initiated a failure of the active mailbox copy on LON-MBX1 and verified continued Outlook functionality.

Initiated a failure of the Client Access services on LON-CAS1.

Initiated a failure of the witness server and tested Outlook functionality.

Recovered the DAG in the secondary site and verified continued Outlook functionality.

MCT USE ONLY. STUDENT USE PROHIBITED

MCT USE ONLY. STUDENT USE PROHIBITED

Module 2: Planning Virtualization for Microsoft Exchange Server 2013

Lab: Planning the Virtualization of Exchange Server Roles

Exercise 1: Designing a Microsoft Exchange Server 2013 Deployment for a Large Organization

▶ Task 1: Use the Exchange Mailbox Calculator to create a configuration

1. Sign in to the **LON-CL1** virtual machine with username of **Adatum\Administrator** and password **Pa$$w0rd**.
2. On the Start screen, click **Excel 2013**. On the **First things first** dialog box, click **Ask me later** and then click **Accept**. Close the Welcome screen.
3. In Microsoft® Excel®, click **Open Other Workbooks**, click **Computer**, and then click **Browse**.
4. Select **C:\Files\E2013Calc.xlsm**, and then click **Open**.
5. Verify you are viewing the **Input** tab.
6. Select **No** in the **Value** column of **Server Multi-Role Configuration (MBX+CAS)**.
7. Select **Yes** in the **Value** column of **Server Role Virtualization**.
8. Type **16** in the **Value** column of **Number of Mailbox Servers Hosting Active Mailboxes / DAG**.
9. Select **No** in the **Value** column of **Consider Storage Designs Utilizing JBOD (if applicable)** value.
10. For Mailbox Server Guest Machines, select **8** for the **Processor Cores / Server** column, and then type **43** for the **SPECint2006 Rate Value** column.
11. Verify that 10% is configured in the **Value** column of **Hypervisor CPU Adjustment Factor**.
12. Type **25000** in the **Value** column of **Total Number of Tier-1 User Mailboxes / Environment**.
13. Type **5120** in the **Value** column of **Mailbox Size Limit (MB)**.

▶ Task 2: Verify the processor configuration generated by the Mailbox Role Calculator

1. Click the **Role Requirements** tab in Excel 2013.
2. Document the following information.

Recommended Minimum Number of Mailbox server Cores	**86**
Recommended Minimum Number of Client Access server Cores	54
Recommended Minimum Number of Global catalog Cores	9
Total cores needed	149

MCT USE ONLY. STUDENT USE PROHIBITED

3. The number of servers is calculated by taking the total number of cores and dividing it by the number of cores assigned to each Client Access server. The calculator recommended 54 cores, and this scenario calls for 8 cores for each server.

```
54 cores / 8 cores/server = 6.75 servers or 7 servers.
```

4. The number of servers is calculated by taking the total number of cores and dividing it by the number of cores assigned to each global catalog server. The calculator recommends 9 cores, and the scenario calls for 4 cores for each server.

```
9 cores / 4 cores/server = 2.25 servers or 3 servers.
```

▶ Task 3: Verify the memory configuration recommended by the Mailbox Configuration Calculator

1. Document the value in the **Recommended RAM Configuration** row in the **/ Primary Datacenter Server (Double Failure)** column. This value should be 24 gigabytes (GB).
2. Calculate the total amount of memory needed for all Mailbox servers. This is calculated by multiplying the number of servers by the amount of memory assigned to each.

```
16 servers * 24 GB/server = 384 GB total.
```

3. The amount of memory required for Client Access servers is obtained by multiplying the number of Client Access servers by the amount of memory assigned to each.

```
7 servers * 8 GB/server = 56 GB total.
```

4. The amount of memory required for Client Access servers is obtained by multiplying the number of Client Access servers by the amount of memory assigned to each.

```
3 servers * 8 GB/server = 24 GB total.
```

5. Document the total recommended amount of memory needed to support the configuration.

```
384 GB + 56 GB + 24 GB = 464 GB
```

▶ Task 4: Decide whether to deploy Exchange using either virtual or physical servers

1. Are there enough processor cores to virtualize this Exchange configuration? Review the information you have gathered: There are 128 hypervisor cores available, but the configuration requires 149 cores. Either more hypervisor resources need to be added or additional physical servers should be allocated.
2. Is there enough memory to virtualize this Exchange configuration? Review the information you have gathered above: 512 GB of memory is available, and the configuration requires 464 GB of memory. There is enough memory available to virtualize the configuration.
3. Will you deploy physical or virtual servers? Enough memory resources are available to virtualize the configuration; however enough processor resources are not available. This configuration should not be virtualized.
4. Close Excel 2013 and click **Don't Save** when prompted.

Results: After completing this exercise, you should have designed a Microsoft® Exchange Server 2013 deployment for a large organization.

MCT USE ONLY. STUDENT USE PROHIBITED

Exercise 2: Designing an Exchange Server 2013 Deployment for a Medium-Sized Organization

▶ Task 1: Input information into the Exchange 2013 Mailbox Server Role Requirements Calculator

1. Sign in to the **LON-CL1** virtual machine with username of **Adatum\Administrator** and password **Pa$$w0rd**.
2. On the **Start** screen, click **Excel 2013**.
3. In Excel 2013, click **Open Other Workbooks**, click **Computer**, and then click **Browse**.
4. Select **C:\Files\E2013Calc.xlsm**, and then click **Open**.
5. Verify you are viewing the **Input** tab.
6. Select **No** in the **Value** column of **Server Multi-Role Configuration (MBX+CAS)**.
7. Select **Yes** in the **Value** column of **Server Role Virtualization**.
8. Select **No** in the **Value** column of **High Availability Deployment**.
9. Type **4** in the **Value** column of **Number of Mailbox Servers** value.
10. For Mailbox Server Guest Machines select **4** for the **Processor Cores / Server** column, and then type **20** for the **SPECint2006 Rate Value** column.
11. Verify the Hypervisor CPU Adjustment Factor in the **Value** column is set to **10%**.
12. Type **2500** in the **Value** column of **Total Number of Tier-1 User Mailboxes / Environment**.
13. Type **5120** in the **Value** column of **Mailbox Size Limit (MB)**.

▶ Task 2: Verify the Processor Configuration generated by the Mailbox Role Calculator

1. Click the **Role Requirements** tab in Excel 2013.
2. Document the following information.

Recommended Minimum Number of Mailbox server Cores	**9**
Recommended Minimum Number of Client Access server Cores	6
Recommended Minimum Number of Global catalog Cores	1
Total cores needed	16

3. The minimum number of highly available Client Access servers is two.
4. The minimum number of highly available global catalog servers is two.

MCT USE ONLY. STUDENT USE PROHIBITED

5. The minimum number of cores needed for high availability is as follows.

Recommended Minimum Number of Mailbox server Cores	32
Recommended Minimum Number of Client Access server Cores	16
Recommended Minimum Number of Global catalog Cores	8
Total cores needed	56

▶ Task 3: Verify the memory configuration recommended by the Mailbox Configuration Calculator

1. Document the value in the **Recommended RAM Configuration** row in the **/ Primary Datacenter Server** column. This value should be 12 GB.
2. Calculate the total amount of memory needed for all Mailbox servers. This is calculated by multiplying the number of servers by the amount of memory assigned to each.

```
4 servers * 12 GB/server = 48 GB total.
```

3. The amount of memory required for Client Access servers is obtained by multiplying the number of Client Access servers by the amount of memory assigned to each.

```
2 servers * 8 GB/server = 16 GB total.
```

4. The amount of memory required for Client Access servers is obtained by multiplying the number of Client Access servers by the amount of memory assigned to each.

```
2 servers * 8 GB/server = 16 GB total.
```

5. Document the total recommended amount of memory needed to support the configuration.

```
48 GB + 16 GB + 16 GB = 80 GB
```

▶ Task 4: Decide whether to deploy Exchange using either virtual or physical servers

1. Are there enough processor cores to virtualize this Exchange configuration? Review the information you have gathered: There are 128 hypervisor cores available, and the configuration requires 56 cores.
2. Is there enough memory to virtualize this Exchange configuration? Review the information you have gathered above: 512 GB of memory is available, and the configuration requires 80 GB of memory. There is enough memory available to virtualize the configuration.
3. Will you deploy physical or virtual servers? Enough memory and processor resources are available to virtualize the configuration.
4. Close Excel 2013 and click **Don't Save** when prompted.

Results: After completing this exercise, you should have designed an Exchange Server 2013 deployment for a medium-sized organization.

MCT USE ONLY. STUDENT USE PROHIBITED

Exercise 3: Designing an Exchange Server 2013 Deployment for a Medium Complexity Organization

▶ Task 1: Input information into the Exchange 2013 Mailbox Server Role Requirements Calculator

1. Sign in to the **LON-CL1** virtual machine with username of **Adatum\Administrator** and password **Pa$$w0rd**.
2. On the **Start** screen, click **Excel**.
3. In Excel 2013, click **Open Other Workbooks**, click **Computer**, and then click **Browse**.
4. Select **C:\Files\E2013Calc.xlsm** and click **Open**.
5. Verify you are viewing the **Input** tab.
6. Select **No** for the **Value** column of **Server Multi-Role Configuration (MBX+CAS)**.
7. Select **Yes** for the **Value** column of **Server Role Virtualization**.
8. Verify that **Yes** is selected in the **Value** column of the **High Availability Deployment**.
9. Type **5** in the **Value** column of **Number of Mailbox Servers Hosting Active Mailboxes / DAG** value.
10. Select **3** in the **Value** column of **Total Number of HA Database Copy Instances (Includes Active Copy) within DAG**.
11. Select **No** in the **Value** column of **Consider Storage Designs Utilizing JBOD (if applicable)** value.
12. In the Mailbox Server Guest Machines section, select **8** for the **Processor Cores / Server** column, and type **46** for the **SPECint2006 Rate Value** column.
13. Verify the **Hypervisor CPU Adjustment Factor Value** column is set to **10%**.
14. Type **5000** in the **Value** column of **Total Number of Tier-1 User Mailboxes / Environment**.
15. Type **1024** in the **Value** column of **Mailbox Size Limit (MB)**.

▶ Task 2: Verify the processor configuration generated by the Mailbox Role Calculator

1. Click the **Role Requirements** tab in Excel 2013.
2. Document the following information.

Recommended Minimum Number of Mailbox server Cores	**20**
Recommended Minimum Number of Client Access server Cores	12
Recommended Minimum Number of Global catalog Cores	2
Total cores needed	34

3. The minimum number of highly available Client Access servers is two.
4. The minimum number of highly available global catalog servers is two.

MCT USE ONLY. STUDENT USE PROHIBITED

5. The minimum number of cores needed for high availability is as follows.

Recommended Minimum Number of Mailbox server Cores	**40**
Recommended Minimum Number of Client Access server Cores	16
Recommended Minimum Number of Global catalog Cores	8
Total cores needed	64

▶ Task 3: Verify the memory configuration recommended by the Mailbox Configuration Calculator

1. Document the value in the **Recommended RAM Configuration** row in the **/ Primary Datacenter Server (Double Failure)** column. This value should be 24 GB.
2. Calculate the total amount of memory needed for all Mailbox servers. This is calculated by multiplying the number of servers by the amount of memory assigned to each.

```
5 servers * 24 GB/server = 120 GB total.
```

3. The amount of memory required for Client Access servers is obtained by multiplying the number of Client Access servers by the amount of memory assigned to each.

```
2 servers * 8 GB/server = 16 GB total.
```

4. The amount of memory required for Client Access servers is obtained by multiplying the number of Client Access servers by the amount of memory assigned to each.

```
2 servers * 8 GB/server = 16 GB total.
```

5. Document the total recommended amount of memory needed to support the configuration.

```
120 GB + 16 GB + 16 GB = 152 GB
```

▶ Task 4: Decide whether to deploy Exchange using either virtual or physical servers

1. Are there enough processor cores to virtualize this Exchange configuration? Review the information you have gathered: 72 hypervisor cores are available, and the configuration requires 64 cores.
2. Is there enough memory to virtualize this Exchange configuration? Review the information you have gathered above: 192 GB of memory is available, and the configuration requires 152 GB of memory.
3. Will you deploy physical or virtual servers? Enough memory and processor resources are available to virtualize the configuration.
4. Close Excel 2013, and click **Don't Save** when prompted.

Results: After completing this exercise, you should have designed an Exchange Server 2013 deployment for a medium-sized organization.

MCT USE ONLY. STUDENT USE PROHIBITED

▶ Task 5: To prepare for the next module

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

1. On the host computer, start **Microsoft Hyper-V® Manager**.
2. In the **Virtual Machines** list, right-click **20342B-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 to 3 for **20342B-LON-CL1**.
5. In Hyper-V Manager, click **20342B-LON-DC1**, and in the Actions pane, click **Start**.
6. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
7. Sign in using the following credentials:
   a. User name: **Adatum\Administrator**
   b. Password: **Pa$$w0rd**
8. Repeat steps 5 to 7 for **20342B-LON-MBX1**. When you have successfully signed in to LON-MBX1, repeat steps 5 to 7 for **20342B-LON-CAS1**, **20342B-LON-CL1**, **20342B-LON-CL2**, and **20342B-LON-LY1**.

MCT USE ONLY. STUDENT USE PROHIBITED

MCT USE ONLY. STUDENT USE PROHIBITED

Module 3: Overview of Exchange Server 2013 Unified Messaging

Lab: Unified Messaging Overview

Exercise 1: Identifying Unified Messaging Components

▶ Task 1: Matching the Unified Messaging terminology

- Review the following list of terms, and then match them to the correct definition or description.

VoIP gateway	A device that would be deployed between an analog PBX and an Exchange 2013 server
Dial plan	Defines user extensions
Auto Attendant	Provides a searchable phone list
Call answering	Can be configured by each user
Hunt group	Identifies a group of telephone users
Unified Messaging IP Gateway	Identifies the first hop when Exchange 2013 servers need to communicate with the PSTN
IP-PBX	A device that can be used with VoIP phones
UM Mailbox Policy	Defines some of the user experience with Unified Messaging
Pilot number	One call to this number could reach many phones
Analog PBX	A device that cannot communicate directly with Exchange 2013 servers

Results: After completing this exercise, you should be able to identify the main Unified Messaging components.

MCT USE ONLY. STUDENT USE PROHIBITED

MCT USE ONLY. STUDENT USE PROHIBITED

Module 4: Designing and Implementing Exchange Server 2013 Unified Messaging

Lab: Designing and Implementing Exchange Server 2013 Unified Messaging

Exercise 1: Designing the Unified Messaging Implementation

▶ Task 1: Create a design for the Exchange Server 2013 Unified Messaging components

Review the information in the Exercise Scenario and answer the following questions:

1. What Exchange Server 2013 dial plans will you need to configure? How will you configure the dial plans?

   Answer: You will need to configure three dial plans: one for each office. You will need to configure the dial plans to use a five-digit extension. You also must configure the dial plan to use SIP URI, and to use mutual TLS.

2. How will you meet the requirement to enable external and internal users to reach the organization's telephone directory by dialing local or internal numbers?

   Answer: You will need to configure an auto attendant for each dial plan, and configure the access numbers to match the local phone number for the auto attendant, as well as the local extension number. When you run the OcsUMUtil.exe, you will need to configure the auto attendant contact for each local auto-attendant access number.

3. How will you meet the requirement that users should be to get service in English or French?

   Answer: On each auto attendant, provide key mappings that allow the user to request the service language. Install the English and French language packs on the Exchange 2013 mailbox servers.

4. How will you meet the requirement that users should be able to search for recipients in each office?

   Answer: On each auto attendant, provide key mappings that allow the user to search for recipients in each office.

5. How will you meet the requirement that enables employees to access their email and voice mail, by phone, by dialing a local number or internal extension?

   Answer: You will need to configure the Outlook Voice Access settings for each dial plan, and configure the access numbers to match the local Outlook Voice Access phone number and the local extension number. When you run the OcsUMUtil.exe, you will need to configure the Subscriber contact for each local auto attendant access number.

6. How will you meet the requirement for the different PIN settings for different groups of employees?

   Answer: You will need to configure at least two UM mailbox policies, and then assign users to the correct policy when you UM-enable them.

▶ Task 2: Discuss your design with the rest of the class

- As a group, discuss the proposed solutions from the students, to find the solution that best fits the requirements of A. Datum.

Results: After completing this exercise, you will have designed an Exchange Unified Messaging deployment.

MCT USE ONLY. STUDENT USE PROHIBITED

Exercise 2: Configuring Unified Messaging Features

▶ Task 1: Configure Unified Messaging dial plans

1. On LON-CAS1, open **Internet Explorer**, and then connect to **https://LON-CAS1.adatum.com/ecp**.
2. Sign in as **Adatum\administrator** using the password **Pa$$w0rd**.
3. In the Features pane, click **unified messaging**.
4. On the **UM dial plans** tab, click **New**.
5. In the new UM dial plan window, type **Lync-Dialplan** in the **Name** field.
6. Verify that the **Extension length (digits)** field is set to **5**.
7. Under **Dial plan type**, click **SIP URI**.
8. Under **VoIP security mode**, click **Secured**.
9. Under **Audio language**, accept the default. Only the English (United States) is available because it is the only language pack that is installed.
10. In the **Country/Region Code** field, type **44** for the UK country code.
11. Click **save**.
12. Click **Lync-Dialplan**, and then click **Edit**.
13. In the Lync-Dialplan window, click **configure**.
14. In the second Lync-Dialplan window, click **dial codes**.
15. In the **Outside line access code** field, type **9**.
16. In the **International access code** field, type **00**.
17. In the **Outlook Voice Access** section, in the **Outlook Voice Access numbers** field, type +**4417144442000**, and then click **Add**.
18. On the **settings** tab, review the default settings.
19. On the **dialing authorization** tab, select the **Allow calls to any extension** check box.
20. On the **transfer & search** tab, click **In the entire organization**.
21. Click **save**, and then click **close**.

▶ Task 2: Configure the UM IP Gateway

1. On LON-CAS1, in the EAC, click the **UM IP gateways** tab.
2. Click **New**.
3. In the new UM IP gateway window, type **LON-UM-Gateway** in the **Name** field.
4. In the **Address** field, type **172.16.0.40**.
5. Under **UM dial plan**, click **browse**.
6. Click **Lync-Dialplan**, and then click **ok**.
7. Click **save**.

MCT USE ONLY. STUDENT USE PROHIBITED

▶ Task 3: Review the default Unified Messaging hunt group

1. On LON-CAS1, on the **Start** screen, click **Exchange Management Shell**.
2. At the command prompt, type **Get-UMHuntGroup**, and then press Enter. Verify that a default hunt group has been created.
3. Type **Get-UMHuntGroup | FL**, and then press Enter. Review the hunt group configuration.

▶ Task 4: Configure Unified Messaging mailbox policies

1. On LON-CAS1, in the Exchange Management Shell, type **Get-UMMailboxPolicy**, and then press Enter. Verify that a default mailbox policy has been created and associated with the Lync-Dialplan dial plan.
2. Type **Get-UMMailboxPolicy | FL**, and then press Enter to see the detailed information about the dial plan.
3. On LON-CAS1, in the EAC, in the unified messaging pane, click UM dial plans, and then click **Lync-Dialplan**.
4. Click **edit**.
5. Under **UM Mailbox Policies**, click. **New**.
6. In the new **UM mailbox policy** page, in the **Name** field, type **Managers-UMMailboxPolicy**, and then click **save**.
7. Click **Managers-UMMailboxPolicy**, and then click **Edit**.
8. On the **message text** tab, in the **When a user is enabled for Unified Messaging** field, type **Your mailbox has been enabled for Unified Messaging**.
9. On the **PIN policies** tab, change the **Minimum PIN length** to **8**, the **PIN recycle count** to **8**, and the **Enforce PIN lifetime** setting to **30**.
10. Click **save**, and then click **close**.

▶ Task 5: Configure a Unified Messaging Auto Attendant

1. On LON-CAS1, in the EAC, in the unified messaging pane, click **Lync-Dialplan**.
2. Click **edit**.
3. In the Lync-Dialplan window, under **UM Auto Attendants**, click **New**.
4. In the new UM auto attendant window, in the **Name** field, type **Adatum-AutoAttendant**.
5. Select the **Create this auto attendant as enabled** check box.
6. Select the **Set this auto attendant to respond to voice commands** check box.
7. Under **Access numbers**, type **+4417144449999**, and then click **Add**.
8. Click **save**, and then click **close**.

▶ Task 6: Configure Unified Messaging users

1. On LON-CAS1, in the EAC, click **recipients** in the Features pane.
2. On the **mailboxes** tab, click **Search**.
3. In the **search mailboxes** box, type **Kelly**, and then press Enter.
4. Click **Kelly Rollin**, and then click **edit**.
5. Click **mailbox features**.

MCT USE ONLY. STUDENT USE PROHIBITED

6. Under **Phone and Voice Features**, click **Enable**.
7. In the Enable UM mailbox window, under **UM mailbox policy**, click **browse**.
8. Click **Lync-Dialplan Default Policy**, and then click **ok**.
9. Click **next**.
10. In the **SIP address** field, type **Kelly@adatum.com**.
11. In the **Extension number** field, type **11006**.
12. Click **finish**, and then click **save**.
13. On LON-CAS1, in the EAC, on the **mailboxes** tab, click **Search**.
14. In the **search mailboxes** box, type **Benno**, and then press Enter.
15. Click **Benno Kurmann**, and then click **edit**.
16. Click **mailbox features**.
17. Under **Phone and Voice Features**, click **Enable**.
18. In the **Enable UM mailbox** window, under **UM mailbox policy**, click **browse**.
19. Click **Managers-UMMailboxPolicy**, and then click **ok**.
20. Click **next**.
21. In the **SIP address** field, type **Benno@adatum.com**.
22. In the **Extension number** field, type **11005**.
23. Click **finish**, and then click **save**.

MCT USE ONLY. STUDENT USE PROHIBITED

Exercise 3: Configuring Unified Messaging Integration with Lync Server 2013

▶ Task 1: Install a Certificate on the Mailbox Server

1. On LON-CAS1, open **File Explorer**.
2. Click **Computer**, and then double-click **Labfiles (E)**.
3. On the E: drive, create a folder named **Certs**.
4. Right-click **Certs**, point to **Share with**, and then click **Specific people**.
5. Click **Share**, and then click **Done**.
6. If required, open Internet Explorer, and then connect to **https://lon-CAS1.adatum.com/ecp**.
7. Sign in as **Adatum\Administrator** using the password **Pa$$w0rd**.
8. In the Features pane, click **servers**, and then click the **certificates** tab.
9. In the Select server list, click **LON-MBX1.Adatum.com**.
10. Click **New**, and on the new **Exchange certificate** page, click **next**.
11. Type **LON-MBX1.adatum.com** as the friendly name of the certificate, and then click **next** twice.
12. In the **Store certificate request on this server** field, click **browse**, click **LON-MBX1**, and then click **ok**.
13. Click **next** twice.
14. On the wizard page that lists the names that will appear in the certificate, click **Add**.
15. In the **New Domain** dialog box, type **LON-MBX1.adatum.com**, and then click **ok**.
16. Repeat the previous two steps twice to add the names **LON-MBX2** and **LON-MBX2.adatum.com** to the certificate request.
17. Click **next**.
18. Fill in the following information, and then click **next**.
    - Organization Name: **A. Datum**
    - Department Name: **Messaging**
    - Country/Region name: **United Kingdom**
    - City/Locality: **London**
    - State/Province: **EN**
19. On the new **Exchange certificate** page, type **\\lon-CAS1\certs\MBXcertrequest.req**, and then click **finish**. Close Internet Explorer.
20. In File Explorer, browse to the **E:\Certs** folder, and then double-click **MBXCertrequest.req**.
21. Click **Notepad**.
22. In the Notepad Window, press Ctrl+A, and then press Ctrl+C. Close Notepad.
23. In the Internet Explorer, connect to **http://lon-dc1.adatum.com/certsrv**. Sign in as **administrator** using the password **Pa$$w0rd**.
24. In the Welcome window, click **Request a certificate**.
25. In the Request a Certificate windows, click **advanced certificate request**.

MCT USE ONLY. STUDENT USE PROHIBITED

26. In the Advanced Certificate Request window, click **Submit a certificate request by using a base-64-encoded CMC or PKCS #10 file, or submit a renewal request by using a base-64-encoded PKCS #7 file**.
27. In the Submit a Certificate Request or Renewal Request window, click in the **Saved Request** section, and then press Ctrl+V.
28. Under Certificate Template, click **Adatum Web**, and then click **Submit**. If you see a prompt that Internet Explorer has blocked an ActiveX control, close the prompt.

Note: If you receive an error message that the certificate request was denied, restart the Active Directory Certificate Services service on LON-DC1, and then try the request again.

29. In the **Certificate Issued** page, click **Download Certificate**, click **Save as**. Browse to **E:\Certs**, and then click **Save**.
30. In Internet Explorer, connect to **https://lon-cas1.adatum.com/ecp**. If prompted, sign in as **Adatum\Administrator** using the password **Pa$$w0rd**.
31. On the servers pane, click **certificates**.
32. Click **LON-MBX1.adatum.com**, and then click **Complete**.
33. In the complete pending request windows, type **\\lon-cas1\certs\certnew.cer**, and then click **ok**.

▶ Task 2: Configure the Unified Messaging Services to Use TLS

1. On LON-CAS1, open the **Exchange Management Shell**, type **Get-UMService**, and then press Enter.
2. Verify that both Mailbox Servers are listed but that no dialplan is associated with the servers.
3. Type the following command, and then press Enter.

```
Get-MailboxServer | Set-UMService –DialPlans Lync-Dialplan –UMStartupMode TLS
```

4. Review the warnings. Type the following command, and then press Enter.

```
Get-ClientAccessServer | Set-UMCallRouterSettings –DialPlans Lync-Dialplan –
UMStartupMode TLS
```

5. Review the warnings. To view the default UM call router settings, type **Get-UMCallRouterSettings –Server lon-cas1.adatum.com**, and then press Enter.
6. On LON-CAS1, in the EAC, click the **servers** pane, and then click **certificates**.
7. Verify that **LON-MBX1.Adatum.com** is listed in the **Select server** list, and then click the **LON-MBX1.adatum.com** certificate.
8. Click **Edit**.
9. In the LON-MBX1.adatum.com window, on the **services** tab, select the **UM** check box, and then click **save**.
10. Click **yes**, and then click **ok**.
11. In the **Select server** list, click **LON-CAS1.adatum.com**.
12. Click the **Webmail.adatum.com** certificate.
13. Click **Edit**.

MCT USE ONLY. STUDENT USE PROHIBITED

14. In the Webmail.adatum.com window, on the **services** tab, select the **UM call router** check box, and then click **save**.
15. Click **yes**, and then click **ok**.
16. On LON-CAS1, in the Exchange Management Shell, type **restart-service msexchangeumcr**, and then press Enter.
17. On LON-MBX1, open **Windows PowerShell**®.
18. In the Windows PowerShell window, type **restart-service msexchangeUM**, and then press Enter.

Note: If you get an error message indicating that the service cannot be started, ignore this error for now.

▶ Task 3: Configure the Autodiscover Service URI

1. On LON-CAS1, in the Exchange Management Shell, type the following command, and then press Enter.

```
Get-ClientAccessServer | Set-ClientAccessServer -AutoDiscoverServiceInternalUri
https://autodiscover.adatum.com/autodiscover/autodiscover.xml
```

2. Type **nslookup autodiscover.adatum.com**, and then press Enter.
3. Verify that the autodiscover.adatum.com alias references 172.16.0.20.

▶ Task 4: Configure Exchange Server for Lync Server integration

1. On LON-CAS1, in the Exchange Management Shell, type **cd "\Program Files\Microsoft \Exchange Server\V15\Scripts"**, and then press Enter.
2. Type **.\ExchUCUtil.ps1**, and then press Enter.
3. In the Exchange Management Shell, type **Get-UMDialPlan | FL**, and then press Enter.
4. Verify that a new UM IP Gateway has been created, named LON-LY1, and associated with the dial plan Lync-Dialplan.

Note: If the Microsoft Exchange Unified Messaging service did not start previously, on LON-MBX1, in the Exchange Management Shell, type **Get-service msexchangeUM**, and then press Enter. If the service still shows as stopped, type **Start-service msexchangeUM**, and then press Enter. If the service still does not start, wait a few minutes, and then try starting the service again. It can take several minutes for the service to start.

▶ Task 5: Prepare Lync Server for Exchange Integration

1. On LON-LY1, on the **Start** screen, type **CMD**, and then press Enter.
2. At the command prompt, type **cd "C:\Program Files\Common Files \Microsoft Lync Server 2013\Support"**, and then press Enter.
3. Type **ocsumutil.exe**, and then press enter
4. In the Exchange UM Integration Utility window, click **Load Data**.
5. Verify that **Adatum.com** is displayed in the **Exchange UM Dial Plan Forest** field.
6. Under **SIP Dial Plans**, click **Lync-Dialplan.Adatum.com**, and then click **Add**.
7. In the **Contact** dialog box, under **Organizational Unit**, click **Browse**.

MCT USE ONLY. STUDENT USE PROHIBITED

8. Click **Make New OU**, type **UMIntegration**, and then click **OK**.
9. In the **Name** field, type **Lync-Subscriber-Access**.
10. In **SIP Address** field, type **sip:Lync-SA**.
11. Click **Use this pilot number from Exchange UM**, and verify that +**4417144442000** is the number listed. If the phone number is not listed, click **Enter phone number**, and type the phone number.
12. Under **Contact Type**, verify that **Subscriber Access** is selected, and then click **OK**.
13. Click **Add**.
14. In the **Name** field, type **Lync-Autoattendant**.
15. In the **SIP Address** field, enter **sip:Adatum-AA**.
16. Under **Contact Type**, click **Auto-Attendant**.
17. Click **Use this pilot number from Exchange UM**, and verify that +**4417144449999** is the number listed, and then click **OK**.
18. Close the Exchange UM Integration Utility window.

Note: The previous two tasks create two contact items in the organizational unit (OU) that you specified. The first contact routes messages to Outlook Voice Access, and the second contact routes messages to the auto-attendant.

Results: After you have configured the Exchange 2013 Unified Messaging integration with Lync 2013, you will be able to leave voice messages for UM- enabled Exchange users and use the AutoAttendant via Lync 2013 to connect a SIP call to Lync Enterprise Voice-enabled users.

MCT USE ONLY. STUDENT USE PROHIBITED

Exercise 4: Verify Unified Messaging Functionality

▶ Task 1: Enable Enterprise Voice for Lync users

1. On LON-LY1, on the **Start** screen, click **Lync Server Control Panel**.
2. In the Security window, sign in as **Administrator** using the password **Pa$$w0rd**.
3. In the Lync 2013 Control Panel, click **Users**.
4. Click **Enable users**.
5. In the New Lync Server User window, click **Add**.
6. In the **Search** field, type **Benno**, and then click **Find**.
7. Click **Benno Kurmann**, and then click **OK**.
8. Under **Assign users to a pool**, click **LON-LY1.ADATUM.COM**.
9. Under **Generate user's SIP URI** verify that **Use user's email address** is selected.
10. Under **Telephony**, click **Enterprise Voice**.
11. In the **Line URI** field, type **tel:+4417144441005;ext=11005**.
12. Accept the other defaults, and then click **Enable** in the taskbar.
13. Click **Enable users**.
14. In the New Lync Server User window, click **Add**.
15. In the **Search** field, type **Kelly**, and then click **Find**.
16. Click **Kelly Rollin**, and then click **OK**.
17. Under **Assign users to a pool**, click **LON-LY1.ADATUM.COM**.
18. Under **Generate user's SIP URI** verify that **Use user's email address** is selected.
19. Under **Telephony**, click **Enterprise Voice**.
20. In the **Line URI** field, type **tel:+4417144441006;ext=11006**.
21. Accept the other defaults, and then click **Enable** in the taskbar.

▶ Task 2: Verify Enterprise Voice functionality

1. On LON-LY1, on the **Start** screen, click **Lync Server Management Shell**.
2. At the command prompt, type **Test-CsExUMConnectivity –TargetFqdn lon-ly1.adatum.com –UserSipAddress kelly@adatum.com**, and then press Enter.
3. Verify that the test returns a result of **Success**.
4. At the command prompt, type **Test-CsExUMConnectivity –TargetFqdn lon-ly1.adatum.com –UserSipAddress benno@adatum.com**, and then press Enter.
5. Verify that the test returns a result of **Success**.

 **Note:** If you get an error message when you run the **Test-CsExUMConnectivity** command, type **Update-CsAddressBook** at the command prompt, and then press Enter. Wait a few minutes, and then run the **Test-CsExUMConnectivity** commands again.

6. On the host machine, open **Windows Explorer**, and then browse to **D:\Program Files \Microsoft Learning\20342\Drives**. Double-click **LON-CL1.rdp**. Click **Connect**.

MCT USE ONLY. STUDENT USE PROHIBITED

7. In the **User name** field, type **Adatum\Benno**, in the **Password** field, type **Pa$$w0rd**, and then click **OK**. Click **Yes**.
8. Right-click the **Start** screen, and then click **All apps**.
9. Click **Outlook 2013**.
10. On the **Welcome to Outlook 2013** page, click **Next**.
11. On the **Add an Email Account** page, click **Next**.
12. On the **Auto Account Setup** page, click **Next**.
13. Click **Finish**.
14. In the **First things first** dialog box, click **Ask me later**, and then click **Accept**.
15. Verify that Benno received an email welcoming him to Exchange Unified Messaging.

Note: If the message is in the Drafts folder rather than the Inbox, run the following commands in the Exchange Management Shell on LON-MBX, pressing Enter after each command:

- **Restart-Service msexchangesubmission**
- **Restart-Service msexchangedelivery**
- **Restart-Service msexchangetransport**

16. Right-click the **Start** screen, and then click **All apps**.
17. Click **Lync 2013**. In the **Windows Security Alert** dialog box, click **Allow Access**.
18. Sign in as **Administrator** using the password **Pa$$w0rd**, and click **Yes**.
19. Verify that Benno is automatically connected to Lync.
20. On the host machine, in **D:\Program Files\Microsoft Learning\20342\Drives**, double-click **LON-CL2.rdp**. Click **Connect**.
21. In the **User name** field, type **Adatum\Kelly**, and in the **Password** field, type **Pa$$w0rd**, and then click **OK**. Click **Yes**.
22. Right-click the **Start** screen, and then click **All apps**.
23. Click **Outlook 2013**.
24. On the **Welcome to Outlook 2013** page, click **Next**.
25. On the **Add an Email Account** page, click **Next**.
26. On the **Auto Account Setup** page, click **Next**.
27. If a **Security Alert** screen appears, click **Yes**.
28. Click **Finish**.
29. In the **First things first** dialog box, click **Ask me later**, and then click **Accept**.
30. Verify that Kelly received an email welcoming her to Exchange Unified Messaging.
31. Right-click the **Start** screen, and then click **All apps**.
32. Click **Lync 2013**. In the **Windows Security Alert** dialog box, click **Allow Access**.
33. Sign in as **Administrator** using the password **Pa$$w0rd**, and click **Yes**.

MCT USE ONLY. STUDENT USE PROHIBITED

34. Verify that Kelly is connected to Lync automatically
35. In the LON-CL2 Remote Desktop Connection window, in the Lync client, type **Benno@adatum.com** in the **Find someone or dial a number field**.
36. Double-click **Benno@adatum.com**.
37. In the Instant Messenging window that opens, type a short message to Benno, and then press Enter.
38. In the LON-CL1 Remote Desktop Connection window, click the Lync pop up window, type a short message for Kelly, and then press Enter.
39. In the LON-CL1 Remote Desktop Connection window, in the Lync client, click the **Phone** icon.
40. In the **Find someone or dial a number** field, type the number **+4417144441006**, and then click **Call**.
41. In the LON-CL2 Remote Desktop Connection window, click the Lync pop up window, and then answer the call from Benno.
42. Verify that you can communicate between the two clients, and then click **Hang up**.
43. In the LON-CL1 Remote Desktop Connection window, in the Lync client, click the **Phone** icon.
44. In the **Find someone or dial a number** field, type the number **11006**, and then click **Call**.
45. In the LON-CL2 Remote Desktop Connection window, click the Lync pop up window, and then answer the call from Benno.
46. Verify that you can communicate between the two clients, and then click **Hang up**.

▶ Task 3: Verify Unified Messaging integration

1. In the LON-CL1 Remote Desktop Connection window, in the Lync client, click the **Phone** icon.
2. In the **Find someone or dial a number** field, type the number **11006**, and then click **Call**.
3. Do not answer the call on LON-CL2. On LON-CL1, wait for the call to go to voice mail.
4. Leave a message for Kelly.
5. In the LON-CL2 Remote Desktop Connection window, in Outlook, wait for the voice-mail message to appear.

 **Note:** If the message is not delivered within a minute, run the following commands in the Exchange Management Shell on LON-MBX, pressing Enter after each command:

- **Restart-Service msexchangesubmission**
- **Restart-Service msexchangedelivery**
- **Restart-Service msexchangetransport**

6. Verify the accuracy of the message transcription.
7. Click **Play** to replay the message.
8. In the LON-CL1 Remote Desktop Connection window, in the Lync client, click the **Phone** icon.
9. In the **Find someone or dial a number** field, type the number **19999**, and then click **Call**.
10. When the auto-attendant answers, listen for the greeting.

MCT USE ONLY. STUDENT USE PROHIBITED

Note: If you receive a response to call back later, verify that the time in the virtual machines is between 8 am and 6 pm and the date is Monday to Friday.

11. When requested to provide the name of the person who you want to call, say **Kelly Rollin**.
12. When asked to verify the name, say **Yes**.
13. Verify the phone rings for Kelly, and then hang up.
14. In the LON-CL1 Remote Desktop Connection window, in Outlook, create a new meeting request with Kelly for later today.
15. Open the message welcoming Benno to Exchange Unified Messaging.
16. In the Lync client, click the **Phone** icon.
17. In the **Find someone or dial a number** field, type the number **12000**, and then click **Call**.
18. When prompted, enter the PIN that was provided in the Welcome email.
19. Follow the voice prompts to change the PIN.
20. Follow the voice prompts to record your name and to record a personal greeting.
21. When the mailbox is prepared, say **calendar**, and then listen to the Calendar options.
22. Follow the voice prompts to return to the main menu.
23. Access the mailbox by saying **email**.
24. Listen to the reading of the email message, and then hang up.

▶ Task 4: To prepare for the next module

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

1. On the host computer, start **Hyper-V Manager**.
2. In the **Virtual Machines** list, right-click **20342B-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20342B-LON-CAS1**, **20342B-LON-MBX1**, **20342B-LON-LY1**, **20342B-LON-CL1**, and **20342B-LON-CL2**.
5. In Hyper-V Manager, click **20342B-LON-DC1**, and in the Actions pane, click **Start**.
6. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
7. Sign in using the following credentials:
   - User name: **Adatum\Administrator**
   - Password: **Pa$$w0rd**
8. Repeat steps 5 to 7 for **20342B-LON-MBX1**, **20342B-LON-CAS1**, and **20342B-LON-CL1**.

Results: After completing this exercise, you will have configured two users for Enterprise Voice in Lync 2013, verified the Enterprise Voice functionality, and verified the integration between Exchange 2013 Unified Messaging and Lync 2013.

MCT USE ONLY. STUDENT USE PROHIBITED

Module 5: Designing and Implementing Message Transport Security

Lab: Designing and Implementing Message Transport Security

Exercise 1: Planning a Message Transport Implementation

▶ Task 1: Plan a message transport plan

Review the information in the Exercise Scenario and answer the following questions:

1. Do you need transport rules in order to meet the requirements? If so, how many transport rules do you need and how will you implement them?

   Answer: Yes, transport rules are needed to meet the requirements. We need to create four transport rules to fulfill the requirements, and all of them are created in the Exchange Administration Center (EAC).

2. Do you need journaling? If so, how will you implement it?

   Answer: Journaling is required to journal all messages sent from the distribution group Research to internal and external recipients. Because this type of journaling is a Premium feature, it requires an Enterprise client access license.

3. Do you need recipient moderation? If so, how will you implement it?

   Answer: Recipient moderation is required through the management team to approve all messages that are sent to the AllCompany distribution group.

4. How can you protect messages during the message delivery? Is Information Rights Management (IRM) an option? If so, which features can you use to meet the requirements?

   Answer: Transport protection rules allow you to use transport rules to protect messages with IRM by applying AD RMS rights policy templates.

▶ Task 2: Discuss your design with the class

- Discuss the proposed solutions from the students all together to find the solution that best fits the requirements of A. Datum.

Results: After this exercise, you should have created a message transport plan.

MCT USE ONLY. STUDENT USE PROHIBITED

Exercise 2: Implementing Message Transport Security

▶ Task 1: Configure the required transport rules

1. On LON-CAS1, open **Windows® Internet Explorer®**, in the address bar type **https://lon-cas1.adatum.com/ecp**, and then press Enter to open the Exchange admin center.
2. Sign in as **Adatum\administrator** with the password **Pa$$w0rd**.
3. In the Exchange Admin Center, navigate to **mail flow** – **rules**.
4. Click **New**, and then click **Apply disclaimers**.
5. Type **ADatum External Disclaimer** as the **Name** of the rule.
6. From the **Apply this rule if** dropdown menu, select the condition **The recipient is located**.
7. For the **select recipient location** dialog box, click **Outside the organization**, and then click **ok**.
8. Under **Do the following** click **Append the disclaimer**.
9. Next to **Append the disclaimer**, click **Enter text**, and then type **This message contains confidential information and is intended only for the individual named. If you are not the named addressee, you should not disseminate, distribute, or copy this email. Please notify the sender immediately by email if you have received this email by mistake and delete this email from your system.** Click **ok**.
10. Next to **Append the disclaimer**, click **Select one**, and then, in failback action, click **Reject**. Click **ok**.
11. Click **save**.
12. Click **New** to create another rule, based again on the built-in rule template **Apply disclaimers**.
13. Type **ADatum Info disclaimer** as the **name** of the rule.
14. From the **Apply this rule if** dropdown list, click **The sender is**, and then click the mailbox **info@adatum.com**. Click **Add**, and then click **ok**.
15. Under **Do the following** click **Append the disclaimer**.
16. Next to **Append the disclaimer**, click **Enter text**, and then type **This message is sent on behalf of the Information Department of A. Datum and is intended for internal recipients of A. Datum only. If you are not the intended recipient, you are notified that disclosing, copying, distributing, or taking any action in reliance on the contents of this information is strictly prohibited.** Click **ok**.
17. Next to **Append the disclaimer**, click **Select one**, and then, in failback action, click **Reject**. Click **ok**.
18. Click **save**.
19. Open the **Exchange Management Shell** as an Administrator, and then type the following command to create the transport rule **ADatum Customer Approval**:

```
New-TransportRule –Name “ADatum Customer Approval” –SubjectOrBodyMatchesPatterns
“\d\d\d\d(-|.)\d\d\d” –ModerateMessageByUser benno@adatum.com
```

20. In the Exchange Management Shell, type the following command to create the transport rule **ADatum Internal Confidential**:

```
New-TransportRule –Name “ADatum Internal Confidential” –SubjectOrBodyContainsWords
“Confidential” –SentToScope NotInOrganization -RejectMessageReasonText “You are not
allowed to send confidential messages outside the organization”
```

MCT USE ONLY. STUDENT USE PROHIBITED

▶ Task 2: Configure the required message moderation settings

1. In the Exchange admin center, navigate to **recipients – groups**.
2. In the result pane, click the **AllCompany** distribution group, and then click **Edit**.
3. On the **properties** page, click **message approval**, and then complete the following:
   a. Select the **Messages sent to this group have to be approved by a moderator** check box.
   b. In the **group moderators** list, click **Add**.
   c. In the Select group moderators window, find and select **Aidan**, click **Add**, and then click **ok**.
   d. In **Select moderation notifications**, select **Notify senders in your organization when their messages aren't approved**.
4. Click **Save**.

▶ Task 3: Configure the required journal rules

1. Navigate to **compliance management**, click **journal rules**, and then click **New**.
2. In the field **Send journal reports to**, type **Journal**.
3. In **journal rule**, type **Research Journal Rule** as the name of the journal rule.
4. In the field **If the message is sent to or received from**, click **A specific user or group**. Search for **Research**, click **add**, and then click **ok**.
5. In the field **Journal the following messages**, click **All messages**.
6. Click **save**.
7. Navigate to **recipients** – **mailboxes**. Search for the **Journal**, and then double-click to open the properties.
8. On **mailbox delegation**, navigate to **Full Access**, and then click **Add**. Search for the **Managers** distribution group, click **add**, and then click **ok**.
9. Click **save** and **yes** in the warning window.
10. Log on to **LON-DC1** as **Adatum\administrator** with the password **Pa$$w0rd**.
11. From Server Manager, open **Active Directory Users and Computers**, and then navigate to the **Users** container.
12. Right-click **Journal**, and then click **Disable Account**. Click **OK**.
13. Sign out from LON-DC1.

▶ Task 4: Validate the message transport security configuration

1. On LON-CAS1, close Internet Explorer.
2. Open **Internet Explorer**, and connect to **https://lon-cas1.adatum.com/owa**. Sign in to Outlook Web App as **ADatum\Info** with the password **Pa$$w0rd**.
3. On the **language and time zone** page, click **save**.
4. Create an email message and send it to **ed@adatum.com**.
5. Sign out from Outlook Web App.
6. Sign in to **Outlook Web App** as **ADatum\Ed**.
7. Check that the message from info@adatum.com contains the correct disclaimer.

MCT USE ONLY. STUDENT USE PROHIBITED

8. In Outlook Web App, create and send an email message to **Adam@adatum.com** that has the subject **Customer Information** and the number **2012-199** in the body.
9. Sign in to **LON-CL1** as user **ADatum\Benno** with the password **Pa$$w0rd**.
10. Open **Outlook**. Click **Next** three times and then **Finish** to create the Outlook profile for Benno.
11. In the **First things first** dialog box, click **Ask me later**, and click **Accept**.
12. Check that the message from Ed is in your Inbox and waiting for approval.
13. Click **Reject** and **OK** to reject the message from Ed. Close Outlook. Sign out Benno from LON-CL1.
14. On LON-CAS1, check the Inbox in Outlook Web App of Ed for the rejected message.
15. Read the information.
16. Validate the transport rule **ADatum Internal Confidential**, which rejects messages that have the word Confidential in the subject or body, if the recipients are outside the organization. Signed in to Outlook Web App as **Adatum\Ed**, create an email message that has the word **Confidential** in the subject or body, and then send it to **Troy@treyresearch.net**.
17. In Outlook Web App, create and send an email message to the distribution group **AllCompany**.
18. Check that a MailTip informs you that the group is moderated.
19. Sign out Ed from Outlook Web App.
20. Sign in to **LON-CL1** as user **ADatum\Aidan** with the password **Pa$$w0rd**.
21. Open **Outlook 2013** and create the user profile as requested.
22. Check that the message from Ed is in your Inbox and waiting for approval.
23. Click **Approve**.
24. Close Outlook. Sign out Aidan from LON-CL1.
25. Sign in to **LON-CL1** as user **ADatum\Benno** with the password **Pa$$w0rd**.
26. Open **Outlook**.
27. Create and send an email message to **chloe@adatum.com**.
28. Close Outlook. Sign out Benno from LON-CL1.
29. Sign in to **LON-CL1** as user **ADatum\Aidan** with the password **Pa$$w0rd**.
30. Open **Outlook**.
31. Click **File**, and then click **Add Account**.
32. On the **Auto Account Setup** page, type **Journal** in the **Your Name** field.
33. Type **Journal@adatum.com** in the **E-mail Address** field.
34. Type **Pa$$w0rd** in the **Password** and **Retype Password** fields, and then click **Next**.
35. Click **Finish**. Close and the reopen **Outlook**.
36. Verify that the Journal mailbox is listed in Outlook.
37. Check for the journaled message sent from Benno to Chloe.

Results: After this exercise, you should have implemented message transport security.

MCT USE ONLY. STUDENT USE PROHIBITED

Exercise 3: Implementing AD RMS and Exchange Server Integration

▶ Task 1: Configure AD RMS integration

1. If required, sign in to **LON-CAS1** as **Adatum\administrator** with the password **Pa$$w0rd**.
2. From the **Start** screen, open the **Exchange Management Shell**.
3. Type the following command to enable IRM for internal messages, and then press Enter:

```
Set-IRMConfiguration –InternalLicensingEnabled $True
```

4. In the Exchange Management Shell, type the following commands to create the distribution group **ADRMSSuperUser**, and to add the **FederatedEmail** system mailbox as a member. Press Enter after each command.

```
New-DistributionGroup –Name ADRMSSuperUser
Add-DistributionGroupMember ADRMSSuperUser –Member FederatedEmail.4c1f4d8b-8179-4148-
93bf-00a95fa1e042
```

5. Sign in to **LON-DC1** as **ADatum\administrator** with the password **Pa$$w0rd**.
6. Open the **Active Directory Rights Management Services** console, and then expand **lon-dc1.adatum.com (local)**.
7. In the console tree, expand **Security Policies**, and then click **Super Users**.
8. In the Actions pane, click **Enable Super Users**.
9. In the result pane, click **Change Super User Group** to open the Super Users property sheet.
10. In the **Super User group** box, type **ADRMSSuperUser@adatum.com**, and then click **OK**.
11. Close the Active Directory® Rights Management Services (AD RMS) console.
12. On LON-CAS1, in the Exchange Management Shell, type the following command to enable transport encryption, and then press Enter:

```
Set-IRMConfiguration –TransportDecryptionSetting Mandatory
```

13. In the Exchange Management Shell, type the following command to enable IRM on the Client Access server:

```
Set-IRMConfiguration –ClientAccessServerEnabled $True
```

14. On LON-DC1, open **File Manager** and browse to **c:\Inetpub\wwwroot_wmcs\certification**.
15. Right-click **ServerCertification.asmx**, and then click **Properties**.
16. In the **ServerCertification.asmx Properties** dialog box, on the **Security** tab, click **Edit**.
17. In the **Permissions for ServerCertification.asmx** dialog box, click **Add**.
18. In the **Select User, Computer, Service Account, or Group** dialog box, type **Exchange Servers**. Click **Check Names**, and then click **OK**.
19. Under **Allow**, make sure that the **Read & execute** and the **Read** check boxes are selected.
20. Click **OK**.
21. Repeat the previous six steps to add the **AD RMS Service Group** to the permissions.
22. Click **OK** and close all open windows.

MCT USE ONLY. STUDENT USE PROHIBITED

▶ Task 2: Configure the required transport protection rules

1. On LON-CAS1, open **Internet Explorer**, in the address bar type **https://lon-cas1.adatum.com/ecp**, and then press Enter to open the Exchange admin center.
2. Sign in as **Adatum\administrator** with the password **Pa$$w0rd**.
3. In the Exchange Admin center, navigate to **mail flow** – **rules**.
4. In the list view, click the **New** drop-down arrow and then click **Apply rights protection to messages**.
5. In **new rule**, complete the following fields:
   a. In **Name**, type **Manager group Transport Protection Rule**.
   b. In **Apply this rule if**, click **The sender**, and then click **is a member of this group**. Search for the **Managers** group, and then click **add**. Click **ok**.
   c. In **Do the following**, next to **Apply rights protection to the message with**, click **Select one**, and then click the RMS template **Do Not Forward**. Click **ok**.
6. Click **Save**.

▶ Task 3: Configure journal decryption

- In the Exchange Management Shell, type the following command to enable journal report decryption:

```
Set-IRMConfiguration –JournalReportDecryptionEnabled $True
```

▶ Task 4: Validate the AD RMS integration

1. On LON-CL1, verify that you are logged on as Aidan.
2. In Outlook, create a new message.
3. In the **To** field, type **Ed**, and then press Enter.
4. In **Subject**, type **Message from Manager**.
5. Click **Send**.
6. On LON-CAS1, open **Internet Explorer**, in the address bar, type **https://lon-cas1.adatum.com/owa**, and then press Enter to open Outlook Web App.
7. Sign in to **Outlook Web App** as **Adatum\Ed**.
8. Check the email from Aidan, and then try to forward it.

Results: After this exercise, you should have implemented AD RMS integration in Microsoft® Exchange Server 2013.

MCT USE ONLY. STUDENT USE PROHIBITED

▶ To prepare for the next module

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

1. On the host computer, start the **Microsoft Hyper-V® Manager**.
2. In the **Virtual Machines** list, right-click **20342B-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 to 3 for **20342B-LON-CAS1**, **20342B-LON-MBX1** and **20342B-LON-CL1**.
5. In Hyper-V Manager, click **20342B-LON-DC1**, and in the Actions pane, click **Start**.
6. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
7. Sign in using the following credentials:
   - User name: **Adatum\Administrator**
   - Password: **Pa$$w0rd**
8. Repeat steps 5 to 7 for **20342B-LON-MBX1**, **20342B-LON-MBX2**, and **20342B-LON-CAS1**.

MCT USE ONLY. STUDENT USE PROHIBITED

MCT USE ONLY. STUDENT USE PROHIBITED

Module 6: Designing and Implementing Message Retention

Lab: Designing and Implementing Message Retention

Exercise 1: Designing Message Retention and Archiving

▶ Task 1: Design the Mailbox database configuration required for this deployment

1. To save space, you need to make the databases for the archive mailboxes less redundant than the databases for the regular mailboxes.
2. Roughly estimate how much space you need:
   - Sales group:

     77 users x 20 GB = 1540 GB
   - Manager group:

     42 user x 20 GB = 840 GB
3. Next, calculate how many databases you need. The size of a database should not exceed 200 GB.
   - Sales group:

     1540 GB / 200 GB = 8 databases
   - Manager group:

     840 GB / 200 GB = 5 databases
4. In summary, you need 13 databases to store the archive mailboxes of both user groups.
5. Next, make these databases redundant, but less redundant than the databases for the regular mailboxes.
6. To do this, create one additional copy of every database on the second mailbox server. Now you have a total of 26 databases for the archive mailboxes.

▶ Task 2: Design the retention tags required for this deployment

1. You need to archive messages for the Sales and Managers groups, but with different retention periods, as follows:
   - Sales: Move to archive after one year
   - Managers: Move to archive after three years
2. To do this, create one default policy tag for the Sales group that has the following settings:
   - Name: Sales User 1 year move to archive
   - Retention Action: Move to Archive
   - Retention Period: 365 days

MCT USE ONLY. STUDENT USE PROHIBITED

3. The next requirement is to purge any deleted items after 30 days. This applies to both the Sales and Manager groups. The retention action is applied to a default folder.

   To do this, create one retention policy tag for both groups with the following settings:

   - Name: Purge Deleted Items 30 days
   - Retention Action: Permanently Delete
   - Retention Period: 30 days

4. The next requirement is to move any messages from the default folders of Outlook to the Deleted Items folder after two years. This requirement is only for the Sales group.

   To do this, create a default policy tag for the Sales group with the following settings:

   - Name: Default 2 year move to Deleted Items
   - Retention Action: Delete and Allow Recovery
   - Retention Period: 730 days

5. The next requirement is to give both user groups the possibility to tag their messages themselves. Users must be able to tag messages in the following ways:

   - Delete after a specified retention period.
   - Archive after a specified retention period.
   - Do not archive or delete.

6. To do this, create personal tags for all groups of users. Because personal tags are already available for most of the requirements, you need to create only the following personal tags:

   - Name: 2 Year Delete
   - Retention Action: Delete and Allow Recovery
   - Retention Period: 730 days
   - Name: Never archive
   - Retention Action: Move to Archive
   - Retention Period: Never

▶ Task 3: Design the retention policies required for this deployment

1. Create a new retention policy named **Sales MRM Policy**.
2. Create a new retention policy named **Manager MRM Policy**.

Results: After completing this exercise, you will have successfully:

Designed the required mailbox database configuration.

Designed the required retention tags.

Designed the required retention policies.

MCT USE ONLY. STUDENT USE PROHIBITED

Exercise 2: Implementing Message Retention and Archiving

▶ Task 1: Enable in-place archiving on mailboxes

1. On LON-CAS1, open the **Windows® Internet Explorer®** browser.
2. Type **https://lon-cas1.adatum.com/ecp**, and then press Enter.
3. Sign in as **adatum\administrator** with the password **Pa$$w0rd**.
4. In the left navigation pane, click **recipients**.
5. In the tabs pane, click **mailboxes**.
6. In the toolbar, click **More**. Select **Add / Remove columns**.
7. Scroll down the list, and then select **Archive Database** and **Department**. Click **ok**.
8. Click **Department** to sort the users. Change **Items per page** to **500**.
9. In the list view, select all users within the **Sales** department.
10. In the Bulk Edit pane, scroll down, click **More options**, and then click **Enable** under **Archive**.
11. In the Bulk Enable Archive window, click **browse**, and then select **Research** as the destination for the archive mailboxes. Click **ok**, click **save**, and then click **close**.
12. Repeat steps 9 to 11 for all users of the **Managers** department.
13. Refresh the list view, all users of the Sales and Managers departments now have an archive mailbox enabled.
14. Sign out as the administrator from the Exchange admin center.
15. In Internet Explorer, type **https://LON-CAS1.adatum.com/owa**, and then press Enter.
16. Type **adatum\Dan** and the password **Pa$$w0rd** to log on to Dan's mailbox with Microsoft® Outlook® Web App.
17. On the **language and time zone** page, click **save**.
18. On the left side navigation pane, check that **In-Place Archive – Dan Park** is visible. Expand the folder structure.
19. Create and send a new email with Dan and Bill as recipients and with **Message before setting new retention policy** as the subject.
20. Sign out the user Dan from Outlook Web App, and then close Internet Explorer.

▶ Task 2: Create the required retention tags

1. On LON-CAS1, open **Internet Explorer**.
2. Type **https://lon-cas1.adatum.com/ecp**, and then press Enter.
3. Sign in as **adatum\administrator** with the password **Pa$$w0rd**.
4. In the left navigation pane, select **compliance management**, and in the tabs pane, select **retention tags**.
5. In the toolbar, click **New tag (+)**, and then select **applied automatically to entire mailbox (default)**.
6. Type **Sales User 1 year move to archive** as the name.
7. Select **Move to Archive** as the **Retention action**.
8. Type **365** as the **Retention period**.

MCT USE ONLY. STUDENT USE PROHIBITED

9. Click **save**.
10. In the toolbar, click **New tag (+)**, and then select **applied automatically to entire mailbox (default)**.
11. Type **Default 2 years move to Deleted Items** as the name.
12. Select **Delete and Allow Recovery** as the **Retention action**.
13. Type **730** as the **Retention period**.
14. Click **save**.
15. In the toolbar, click **New tag (+)**, and then select **applied automatically to entire mailbox (default)**.
16. Type **Manager 3 years move to archive** as the name.
17. Select **Move to Archive** as the **Retention action**.
18. Type **1095** as the **Retention period**.
19. Click **save**.
20. In the toolbar, click **New tag (+)**, and then select **applied automatically to a default folder**.
21. Type **Purge Deleted Items 30 days** as the name.
22. Under **Apply this tag to the following default folder**, select **Deleted Items**.
23. Select **Permanently Delete** as the **Retention action**.
24. Type **30** as the **Retention period**.
25. Click **save**.
26. In the toolbar, click **New tag (+)**, and then select **applied by users to items and folders (personal)**.
27. Type **2 Year Delete** as the name.
28. Select **Delete and Allow Recovery** as the **Retention action**.
29. Type **730** as the **Retention period**.
30. Click **save**.
31. In the toolbar, click **New tag (+)**, and then select **applied by users to items and folders (personal)**.
32. Type **Never archive** as the name.
33. Select **Move to Archive** as the **Retention action**.
34. Select **Never** as the **Retention period**.
35. Click **save**.

▶ Task 3: Create and apply the retention policies

1. In the tabs pane, select **retention policies**.
2. In the toolbar, click **New** (+).
3. Type **Sales MRM Policy** as the name.
4. Click **Add** (+) below **Retention tags**.

MCT USE ONLY. STUDENT USE PROHIBITED

5. While holding the Ctrl key, select the following retention tags:
   - 6 Month Delete
   - 1 Year Delete
   - 2 Year Delete
   - Never Delete
   - Sales user 1 year move to archive
   - Default 2 years move to Deleted Items
   - Purge Deleted Items 30 days
   - Personal 1 year move to archive
   - Never archive
6. Click **add**, and then click **ok**. Click **save**.
7. In the toolbar, click **New** (+).
8. Type **Manager MRM Policy** as the name.
9. Click **Add** (+) below **Retention Tags**.
10. While holding the CTRL key, select the following retention tags:
    - 6 Month Delete
    - 1 Year Delete
    - 2 Year Delete
    - Never Delete
    - Manager 3 year move to archive
    - Purge Deleted Items 30 days
    - Personal 1 year move to archive
    - Never archive
11. Click **add**, and then click **ok**. Click **save**.
12. In the left navigation pane, select **recipients**, and then in the tabs pane, select **mailboxes**.
13. Change **Items per page** to **500**.
14. Select all of the **Sales** mailboxes, click **more options** and then under **Retention Policy**, click **Update**.
15. Select **Sales MRM Policy** and then click **save**.
16. Click **close**.
17. Repeat steps 14 to 15 for the **Managers**. Apply the **Manager MRM Policy**.
18. Close Internet Explorer.
19. Open the **Exchange Management Shell** as an administrator.
20. Type the following command to apply the retention polices to the mailboxes immediately:

```
Get-Mailbox | Where {$_.RetentionPolicy –ne $NULL} | Start-ManagedFolderAssistant
```

MCT USE ONLY. STUDENT USE PROHIBITED

▶ Task 4: Verify the configuration

1. On LON-CAS1, open **Internet Explorer**, and then type **https://lon-cas1.adatum.com/owa** in the address bar.
2. Log on to **Outlook Web App** as user Dan from the Sales group.
3. Select a message in the Inbox. Right-click the message, and then expand **assign policy**.
4. Verify that the retention tags that are linked to the retention policy are available.
5. Sign out from Outlook Web App.
6. Sign in to **Outlook Web App** as user Bill from the Managers group.
7. On the **language and time zone** page, click **save**.
8. Select a message in the Inbox, right-click the message, and then expand **assign policy**.
9. Verify that the retention tags that are linked to the retention policy are available.
10. Sign out from Outlook Web App.

▶ Task 5: To prepare for the next module

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

1. On the host computer, start **Hyper-V® Manager**.
2. In the **Virtual Machines** list, right-click **20342B-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 to 3 for **20342B-LON-CAS1**, **20342B-LON-MBX1**, and **20342B-LON-MBX2**.
5. In Hyper-V Manager, click **20342B-LON-DC1**, and in the Actions pane, click **Start**.
6. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
7. Sign in using the following credentials:
   - User name: **Adatum\Administrator**
   - Password: **Pa$$w0rd**
8. Repeat steps 4 to 6 for **20342B-LON-MBX1**, **20342B-LON-MBX2**, **20342B-LON-CAS1**, and **20342B-LON-CAS2**.
9. Repeat steps 4 and 5 for **20342B-LON-CL1**. Do not log on until directed to do so.

Results: After completing this exercise, you will have successfully:

Enabled In-Place Archiving on mailboxes.

Created the required retention tags.

Created and applied the retention policies.

Verified the configuration.

MCT USE ONLY. STUDENT USE PROHIBITED

Module 7: Designing and Implementing Messaging Compliance

Lab: Designing and Implementing Messaging Compliance

Exercise 1: Designing Messaging Compliance

▶ Task 1: Design DLP rules that meet the security requirements

1. The exercise scenario defines the following requirements for data loss prevention:
   - No messages with financial information can leave the organization. Specifically, credit card numbers must be blocked.
   - Users can never send information about IP addresses in email.
2. To address these requirements, you need to use Data Loss Protection (DLP) policies in Microsoft® Exchange Server 2013. For the first requirement, you can use a DLP policy template that helps prevent financial information, including credit card numbers, from leaving the organization. You can modify this template to notify an administrator's mailbox if personally identifiable information appears in an email that is directed outside of the organization.

▶ Task 2: Design hold policies that meet the security requirements

1. The exercise scenario defines the following requirements for data preservation:
   - Email about ProjectX must be preserved in mailboxes of users Amr Zaki, Brad Sutton, and Ed Meadows, unaltered, for at least two years.
2. To meet this requirement, define a mailbox search that identifies all email messages related to ProjectX. Search mailboxes for users Amr Zaki, Brad Sutton, and Ed Meadows. In the search options, specify that resulting email messages must be preserved for 720 days.

▶ Task 3: Design the eDiscovery configuration

1. The exercise scenario defines the following requirements for a mailbox search:
   - Members of the Auditing department must be able to search the contents of all mailboxes.
   - Only members of the Auditing department can put mailboxes on a legal hold.
2. To meet these requirements, first identify users who belong to the Auditing department. Then, assign those users to the Discovery Management role.

Results: After completing this lab, you will have designed and implemented a DLP strategy.

MCT USE ONLY. STUDENT USE PROHIBITED

Exercise 2: Implementing Data Loss Prevention

▶ Task 1: Configure a DLP policy by using a template

1. On LON-CAS1, open **Internet Explorer**, and then type **https://lon-cas1.adatum.com/ecp** to open the Exchange admin center.
2. Sign in as **adatum\administrator** with the password of **Pa$$w0rd**.
3. In the Exchange admin center, in the feature pane, click **compliance management**.
4. Click the **data loss prevention** tab.
5. Click the arrow next to **New** (the plus sign (+)), and then click **New DLP policy from template**.
6. In the DLP policy from template window, in the **Name** field, type **Prevent financial data flow**.
7. In **Choose a template**, scroll down and select **U.S. Financial Data**.
8. Click **save**.
9. In the Exchange admin center, click **Prevent financial data flow**, and then click **Edit**.
10. In the edit DLP policy window, on the **general** tab, click **Enforce**.
11. Click **rules**.
12. Click **U.S. Financial: Scan email sent outside – high count**, and then click **Edit**.
13. In the **Do the following** section, in the drop-down list, select **Block the message**, and then click **reject the message and include an explanation**.
14. In the specify rejection reason window, type **This message contains financial information and can't be sent outside the organization**. Click **ok**.
15. Click **add action**.
16. From the drop-down list, click **Generate incident report and send it to**.
17. Click on the **Select one** link to the right of the drop-down list.
18. In the Select Mailbox window, click **Administrator**, and then click **ok**.
19. Click on the **Include message properties** link to the right of the drop-down list.
20. In the Include message properties window, select **original mail**, and then click **ok**.
21. Scroll down and select **Activate this rule on the following date**. Click an arrow next to a date, and then click **today**.
22. In the **Choose a mode for this rule** list, select **Enforce**.
23. Click **save** twice.
24. In the Exchange admin center, click **Customize Policy Tips**.
25. In the Policy Tips window, click **New**.
26. In the New Policy Tip setting window, under **Policy Tip**, click **notify the sender**.
27. In **Locale**, select **English**.
28. In **Text**, type **This message contains information that you are not allowed to send.**
29. Click **save**, and then click **close**.

MCT USE ONLY. STUDENT USE PROHIBITED

▶ Task 2: Configure a DLP by using a custom rule

1. In the Exchange admin center, click the arrow next to **New** (the plus sign (+)), and then click **New custom DLP policy**.
2. In the New custom DLP policy window, in **Name**, type **Prevent IP addresses**.
3. Click **Enforce**.
4. Click **save**.
5. In the Exchange admin center, click **Prevent IP addresses**, and then click **Edit**.
6. In the edit DLP policy window, click **rules**.
7. Click the arrow next to **New** (the plus sign (+)).
8. Click **Block messages with sensitive information unless the sender overrides with a business justification**.
9. In the new rule window, in the **Name** text box, type **Block IP**.
10. Click **Select sensitive information types**.
11. In the sensitive information types window, click **Add**.
12. In the select sensitive information types window, from the list, select **IP Address**, click **add->**, and then click **ok** twice.
13. Click **Select one** in the **Do the following** section.
14. In the Select Mailbox window, click **Administrator**, and then click **ok**.
15. Click on the **Include message properties** link to the right of the drop-down list.
16. In the Include message properties window, select **original mail**, and then click **ok**.
17. Click **Block the message, but allow the sender to override with a business justification and send**.
18. In the notify the sender with a Policy Tip window, in **Enter the message for the NDR that users will receive**, type **You are not allowed to send an IP address in email.** Click **ok**.
19. In **Audit this rule with severity level**, click **Medium**.
20. Click **Activate this rule on the following date**, and then in the drop-down box click **today**.
21. Click **save** twice.

▶ Task 3: Validate the DLP deployment

1. On LON-CL1, sign in as **Adatum\Aidan** with the password **Pa$$w0rd**.
2. From the Desktop, open **File Explorer**, and then browse to **C:\Files**. Open the file **Northwind Customer Data**. If the Microsoft Office welcome window appears, click the Close button in the upper right corner. Examine the content of the file. Close the Microsoft Excel® spreadsheet software.
3. Click the **Start** screen, open **Microsoft Outlook® 2013** by typing **Outlook** and pressing Enter.
4. In the Welcome to Microsoft Outlook 2013 window, click **Next**.
5. On the **Add an Email Account** page, click **Next**.
6. On the **Auto Account Setup**, click **Next**.
7. On the **Configuring** page, click **Finish**.
8. In the First things first window, click **Ask me later** and click **Accept**.

MCT USE ONLY. STUDENT USE PROHIBITED

9. In the Outlook window, click **New Email**.
10. In the new email window, in the **To** field, type **Ben@contoso.com**.
11. In the **Subject** field, type **Northwind data**.
12. Click **Attach file** on the toolbar.
13. Browse to **C:\Files** folder and select **Northwind Customer Data**. Click **Insert**.
14. In the message body, type **Find attached data**. If a policy tip appears, examine the content.
15. Click **Send**.
16. Open **Internet Explorer**, and then type **https://lon-cas1.adatum.com/owa**.
17. Sign in as **Adatum\Administrator** with the password of **Pa$$w0rd**.
18. Ensure that you receive a message with **Rule Detected** words in subject. Examine the message body and ensure that message with financial data is attached.
19. Close Outlook Web App.
20. Switch to **Outlook 2013**.
21. Click **New Email**.
22. In the **To** field, type **Ben@contoso.com**.
23. In the **Subject** field, type **My IP**.
24. In the message body, type **This is my IP address: 172.16.0.100**. Wait for a few moments until policy tip is displayed in title bar of the new mail window.
25. Click **Send**.
26. Verify that you received a message from Microsoft Outlook that says that message can't be send.
27. Click **OK** and close message window. Click **No**.
28. Leave Outlook 2013 open, and stay logged on as Aidan on LON-CL1.

Results: After completing this exercise, you will have implemented DLP.

MCT USE ONLY. STUDENT USE PROHIBITED

Exercise 3: Implementing In-Place eDiscovery

▶ Task 1: Send emails between users

1. On LON-CL1, make sure that you are logged on as **Adatum\Aidan**.
2. Open **Outlook 2013**.
3. Click **New Email**.
4. In the **To** field, type **Bill**.
5. In the **Subject** field, type **just for you**.
6. In the message body, type **It seems like the company won the project for delivering tools to Contoso. We must make sure that we take advantage of this information before authorized personnel do. Let me know what you think.**
7. Click **Send**.
8. Open **Internet Explorer**. Type **https://lon-cas1.adatum.com/owa**.
9. Sign in as **Adatum\Bill** with password **Pa$$w0rd**.
10. On the page with language and time zone settings, click **save**.
11. In Outlook Web App, read the message from Aidan, and then click **Reply**.
12. Type **We must meet with Contoso people as soon as possible. Can you keep this confidential?** Click **Send**.
13. Close Outlook Web App.

▶ Task 2: Configure permissions required for In-Place eDiscovery

1. On LON-DC1, open **Server Manager**, click **Tools**, and then click **Active Directory Users and Computers**.
2. In Active Directory® Users and Computers, expand **Adatum.com**, and then click **Microsoft Exchange Security Groups**.
3. In the details pane on the right, double-click **Discovery Management**.
4. In the Discovery Management Properties window, click the **Members** tab.
5. Click **Add**.
6. In the Select Users, Contacts, Computers, Service Accounts, or Groups window, type **April**, and then click **OK** twice.
7. Open **Internet Explorer**. Type **https://lon-cas1.adatum.com/ecp**, and then press **Enter.**
8. Sign in as **Adatum\Administrator** with the password of **Pa$$w0rd**.
9. In the Exchange admin center, in the features pane, click **permissions**.
10. Click the **admin roles** tab.
11. Double-click **Discovery Management**.
12. In the Role Group window, ensure that **April Reagan** is in the **Members** list. Close the Role Group window.
13. Close the Exchange admin center.

MCT USE ONLY. STUDENT USE PROHIBITED

▶ Task 3: Perform an In-Place eDiscovery search

1. On LON-CL1, open **Internet Explorer**, and then type **https://lon-cas1.adatum.com/ecp**.
2. Sign in as **Adatum\April** with the password of **Pa$$w0rd**. On the page with language and time zone settings, click **save**.
3. In the Exchange admin center, in the feature pane, click **compliance management**.
4. On the toolbar, click **New**.
5. In the new in-place eDiscovery & hold window, in **Name**, type **Contoso search**, and then click **next**.
6. Click **Specify mailboxes to search**, and then click **Add**. In the Select Mailbox window, select **Aidan Delaney** and **Bill Malone**, and then click **add**, and then click **ok**. Click **next**.
7. Click **Filter based on criteria**, and then in **Keywords** type the following:

   (Contoso NEAR(3) tools) AND (confidential OR authorized)
8. Click **select message types**.
9. In the message types to search window, click **Select the message types to search**, and then click **Email**.
10. Click **ok**, and then click **next**.
11. On the **In-Place Hold settings** page, click **finish**.
12. Click **close**.
13. In the Exchange admin center, on the toolbar, click **Refresh**.
14. In the details pane on the right, ensure that the status is **Estimate Succeeded**. If it is not, wait one or two minutes, and then click **Refresh** again.
15. Click **Contoso search**. On the toolbar, click the arrow next to the **Search** icon, and then click **Preview search results**.
16. Ensure that you can see emails between Bill and Aidan that contain the words you searched for.
17. Close eDiscovery search preview window.

▶ Task 4: Configure In-Place Hold

1. In the Exchange admin center, in the feature pane, click **compliance management**.
2. On the toolbar, click **New**.
3. In the new in-place eDiscovery & hold window, in **Name**, type **ProjectX data preservation**, and then click **next**.
4. On the **mailboxes** page, click **Specify mailboxes to search**, and then click **Add**.
5. In the Select Mailbox window, add mailboxes for the following users: **Amr Zaki**, **Brad Sutton**, and **Ed Meadows**, and then click **ok**.
6. Click **next**.
7. On the **Search query** page, click **Filter based on criteria**. In **Keywords**, type **ProjectX**.
8. Click **select message types**.
9. In the Message Types to Search window, click **Select the message types to search**, and then click **Email**.
10. Click **ok**, and then click **next**.

MCT USE ONLY. STUDENT USE PROHIBITED

11. On the **In-Place Hold settings** page, click **Place content matching the search query in selected mailboxes on hold**, click **Specify number of days to hold items relative to their received date**, and then type **720**.
12. Click **finish**.
13. Click **close**.

Note: After you configure mailboxes for In-Place Hold, you can search for deleted or modified items in these mailboxes by using same procedure for eDiscovery search.

▶ Task 5: To prepare for the next module

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

1. On the host computer, start **Hyper-V® Manager**.
2. In the **Virtual Machines** list, right-click **20342B-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 to 3 for **20342B-LON-CAS1**, **20342B-LON-MBX1**, **20342B-LON-MBX2**, and **20342B-LON-CL1**.
5. In Hyper-V Manager, click **20342B-LON-DC1**, and in the Actions pane, click **Start**.
6. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
7. Sign in using the following credentials:
   - User name: **Adatum\Administrator**
   - Password: **Pa$$w0rd**
8. Repeat steps 5 to 7 for **20342B-LON-MBX1** and **20342B-LON-CAS1**.

Results: After completing this exercise, you will have configured eDiscovery.

MCT USE ONLY. STUDENT USE PROHIBITED

Exercise 4: Comparing Messaging Policy and Compliance Options

▶ Task 1: Discuss alternative solutions

1. Review the solution that you provided in Exercises 1, 2, and 3. Discuss with the students any alternative solutions that produce a result that is similar or the same.
2. Review the usage scenarios for DLP policies, In-Place Holds, and In-Place eDiscovery.
3. Which compliance situations does each of these technologies address?
4. Which technologies have overlapping features?
5. How should you decide which technology to use?

Results: After completing this exercise, students will have discussed alternative solutions for messaging policy and compliance options.

MCT USE ONLY. STUDENT USE PROHIBITED

Module 8: Designing and Implementing Administrative Security and Auditing

Lab: Designing and Implementing Administrative Security and Auditing

Exercise 1: Designing a Solution

▶ Task 1: Read and analyze scenario requirements

- Read the exercise scenario and analyze the requirements from a security perspective. Identify the permissions or audit requirements needed to satisfy the requirements.

▶ Task 2: Design a solution

Answers to questions:

1. You can use the following built-in role groups:
    - Organizational Management, to satisfy the requirement for "Exchange Organization Administrators." They do not require additional permissions and do not have any permission in the Active Directory® Domain Services (AD DS) directly.
    - Recipient Management, to satisfy the requirement for "delegated Exchange administrators."
    - Compliance Management, to satisfy the requirements of the audit department.
    - Delegated Setup, to satisfy the requirements for site administrators.
2. You need the following custom role group:
    - For Mailbox Managers, you need a custom management role group, and you need to add the "Mail Recipients" role to it. All built-in groups provide too much permission.
3. Mailbox audit logging for non-owner access must be enabled on all mailboxes. Administrator audit logging is enabled by default, so you do not need to configure it.

▶ Task 3: Discuss your solution with the class

- Present your proposed solution to the class. Discuss alternative solutions with other students and with the instructor.

MCT USE ONLY. STUDENT USE PROHIBITED

Exercise 2: Implementing Role-Based Access Control

▶ Task 1: Use the Recipient Management role group to delegate permissions to create and manage recipients

1. On LON-MBX1, open **Windows® Internet Explorer®**, and then connect to **https://LON-CAS1.adatum.com/ecp**. Sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.
2. In the Exchange Admin Center, in the feature pane, click **permissions**.
3. On the tabs, click **admin roles**, and then, in the list view, double-click **Recipient Management**.
4. In the Role Group window, in **Members**, click **Add**.
5. On the **Select Member** page, select **Brad Sutton**, click **add**, and then click **ok**.
6. In the Role Group window, click **save**.

▶ Task 2: Create a custom role group to allow only enabling and disabling of mailboxes

1. In the Exchange Admin Center, on the **admin roles** tab, on the toolbar, click **New**.
2. In the Role Group window, in **Name** type **Mailbox Managers**, in **Description** type **Enable and disable mailboxes**, and then, in **Roles**, click **Add**.
3. In the Select a Role window, select **Mail Recipients**, click **add**, and then click **ok**.
4. In the Role Group window, in **Members**, click **Add**.
5. On the **Select Member** page, select **Erwin Zischka**, click **add**, and then click **ok**.
6. In the Role Group window, click **save**.
7. Close Internet Explorer.

▶ Task 3: Verify that the administrators have permission to perform their tasks

1. Switch to **LON-CAS1**, open **Internet Explorer**, and then connect to **https://LON-CAS1.adatum.com/ecp**. Sign in as **Adatum\Brad** with the password **Pa$$w0rd**.
2. In the Exchange Admin Center, click **save**.
3. On the tabs, click **mailboxes**.
4. On the toolbar, click **New**. Click **User mailbox**.
5. In the User Mailbox window, in **Alias** type **Test**, and then click **New user**.
6. In **First name** type **Test**, in **Last name** type **Test**, and then in **User logon name** type **Test**. In **New password** and **Confirm password** type **Pa$$w0rd**, and then click **save**. This confirms that Brad can create new mailboxes.
7. In the feature pane, click **servers**. Note that you can see only servers on tabs.
8. In the list view, double-click **LON-MBX1**, and then verify that you cannot modify anything.
9. Close Internet Explorer.
10. Open **Internet Explorer**, and then connect to **https://LON-CAS1.adatum.com/ecp**. Sign in as **Adatum\Erwin** with the password **Pa$$w0rd**.
11. In the Exchange Admin Center, click **save**.
12. In the list view, click **Amr Zaki**. On the toolbar. click **more**, and then click **Disable**.

MCT USE ONLY. STUDENT USE PROHIBITED

13. In the warning window, click **yes**. Confirm that Amr Zaki's mailbox is gone from the list view.
14. Note that, on the feature pane, servers are not available to Erwin, because of his restricted permissions.
15. Close Internet Explorer.

MCT USE ONLY. STUDENT USE PROHIBITED

Exercise 3: Implementing Mailbox Audit Logging

▶ Task 1: Enable mailbox audit logging of non-owners on all mailboxes

1. On the LON-MBX1 virtual machine, click to the **Start** screen, and then click **Exchange Management Shell.**
2. In the Exchange Management Shell, at the **PS** prompt, type the following:

```
Add-MailboxPermission Tony -AccessRights fullaccess –user Administrator
```

3. In the Exchange Management Shell, at the **PS** prompt, type the following:

```
Get-Mailbox | Set-Mailbox -AuditEnabled $true
```

4. Minimize the Exchange Management Shell.

▶ Task 2: Delete items from a mailbox by using a different account

1. Open **Internet Explorer**, type **https://LON-CAS1.adatum.com/owa**, and then press Enter.
2. Sign in to **Outlook Web App** as **Adatum\Administrator** with the password **Pa$$w0rd**.
3. In Outlook Web App, in the upper right corner, click **Administrator**, and then click **Open another mailbox**.
4. In **open another mailbox dialog box**, type **Tony@adatum.com**, then select **Tony Smith**, click **open**, and then click **save**.
5. In Tony's Inbox, select a message, and then click **Delete**.
6. On the left navigation pane, right-click **Deleted Items**, and then click **empty**.
7. In the **empty folder** window, click **ok**.
8. Close Internet Explorer.

▶ Task 3: Verify that the activity is logged

1. Open **Internet Explorer**, and then type **https://LON-CAS1.adatum.com/ecp**.
2. Sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.
3. In the Exchange admin center, in the feature pane, click **compliance management**.
4. On the tabs, click **auditing**.
5. Click **Run a non-owner mailbox access report**.
6. In the **Search for access by** drop-down box, select **All non-owners**, and then click **Search**.
7. In **Search results**, click **Tony Smith**, and then notice in the report that the Administrator performed a soft-delete operation in the mailbox.
8. Click **Close**, and then close Internet Explorer.

MCT USE ONLY. STUDENT USE PROHIBITED

Exercise 4: Using Administrative Audit Logging to Troubleshoot

▶ Task 1: Change the Exchange configuration

1. On the LON-CAS1 virtual machine, click to the **Start** screen and then click **Exchange Management Shell**.
2. In the Exchange Management Shell, at the **PS** prompt, type the following and then press Enter:

```
E:\LabFiles\Mod08\Mod08Ex4.bat
```

3. Close the Exchange Management Shell.

▶ Task 2: Troubleshoot what happened by using the administrative audit log

1. On the LON-MBX1 virtual machine, maximize the Exchange Management Shell.
2. In the Exchange Management Shell, at the **PS** prompt, type the following:

```
Search-AdminAuditLog -ObjectIds "April Reagan"
```

3. If only one log entry is listed, wait one minute and run the command again.
4. Locate the log entry with the **CmdletName Disable-Mailbox**. Note that the **Caller** name is **Adatum.com/Managers/Ed Meadows**. Because Ed Meadows is not an Exchange administrator of A. Datum Ltd. that you know of, you need to investigate further.
5. At the **PS** prompt, type the following:

```
Get-Mailbox "Ed Meadows"
```

6. You find out that the account for Ed Meadows does not exist anymore. The mailbox must have been removed.
7. At the **PS** prompt, type the following:

```
Search-AdminAuditLog –StartDate yesterday –EndDate tomorrow | Sort-object rundate
```

   In this command, the dates are written in the mm/dd/yyyy format. So May 7, 2013 would be written as 05/07/2013.

8. Review the last several items in the audit log. Notice that there is a **Remove-Mailbox** cmdlet used to remove Ed Meadows' mailbox. Verify that the caller was Ed himself. The next entry shows that Ed modified April's account. The next entry shows that an account was added to the Organizational Management role, and the caller field shows the user account that did this: Don Funk. Therefore, you need to talk to Don Funk to find out why he added Ed to the Organization Management group.

MCT USE ONLY. STUDENT USE PROHIBITED

Exercise 5: Implementing and Testing Active Directory Split Permissions

▶ Task 1: Configure Active Directory split permissions for Exchange Server

1. On LON-MBX1, in the 20342B-LON-MBX1 on localhost - Virtual Machine Connection window, in the menu bar, click **Media**, select **DVD Drive**, and then click **Insert Disk**.
2. Navigate to **D:\Program Files\Microsoft Learning\20342\Drives\ExchangeServer2013CU1.iso**, and then click **Open**.
3. Open the **Windows PowerShell**® command-line interface.
4. Type **D:** and then press Enter.
5. Type **Setup /PrepareAD /ActiveDirectorySplitPermissions: true /IAcceptExchangeServerLicenseTerms**, and then press Enter.
6. Wait until the process finishes, and then close Windows PowerShell.

▶ Task 2: Verify that the Exchange Server administrators cannot change objects directly in AD DS

1. On LON-MBX1, open **Internet Explorer**, and then connect to **https://LON-CAS1.adatum.com/ecp**. Sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.
2. In the tabs, click **mailboxes**, and then on the toolbar click **New**. Click **User mailbox**.
3. In the **User Mailbox** dialog box, in the **Alias** field type **Ales**, and then click **New user**. Note that all of the fields for creating a new user, such as **First name**, **Last name**, and **User logon name**, are grayed out. Therefore, even though this administrator is a Domain Admin, you cannot create a user object in Microsoft® Exchange Server anymore through their account.
4. In the **User Mailbox** dialog box, click **Existing user**, and then click **browse**.
5. In the **Select User** dialog box, double-click **Ales Ruzicka**, and then click **save**. Note that a mailbox for Ales Ruzicka is created.
6. In the Exchange admin center, in the tabs click **groups**.
7. On the list view, double-click **IT**.
8. In the Security Group window, in the left pane, click **membership**, and then in **Members** click **Add**.
9. In the **Select Members** dialog box, double-click **Ales Ruzicka**, and then click **ok**.
10. In the Security Group window, click **save**. An error appears that says, "You don't have sufficient permissions." This is because you cannot manage groups anymore from Exchange Server.
11. Click **ok** to close the error message, and then click **cancel**.

MCT USE ONLY STUDENT USE PROHIBITED

▶ Task 3: To prepare for the next module

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

1. On the host computer, start **Microsoft Hyper-V® Manager**.
2. In the **Virtual Machines** list, right-click **20342B-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20342B-LON-CAS1** and **20342B-LON-MBX1**.
5. In Hyper-V Manager, click **20342B-LON-DC1**, and in the Actions pane, click **Start**.
6. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
7. Sign in using the following credentials:
   - User name: **Adatum\Administrator**
   - Password: **Pa$$w0rd**
8. Repeat steps 5 to 7 for **20342B-LON-MBX1** and **20342B-LON-CAS1**.

Results: In this exercise, you configured Active Directory split permissions.

MCT USE ONLY. STUDENT USE PROHIBITED

MCT USE ONLY. STUDENT USE PROHIBITED

Module 9: Managing Exchange Server 2013 with Exchange Management Shell

Lab: Managing Microsoft Exchange Server 2013 by Using Exchange Management Shell

Exercise 1: Exploring the Exchange Management Shell

▶ Task 1: Import the Exchange Management Shell module into the Windows PowerShell ISE

1. Sign in to the **LON-CAS1** virtual machine with the user name of **Adatum\Administrator** and the password **Pa$$w0rd**.
2. On the **Start** screen, type **Windows PowerShell ISE**, right-click **Windows PowerShell ISE**, and then click **Run as administrator**.
3. Click **View**, and then click **Show Script Pane**.
4. In the Console pane, run **Import-Module 'C:\Program Files\Microsoft\Exchange Server \V15\bin\RemoteExchange.ps1'; Connect-ExchangeServer -auto -ClientApplication:ManagementShell**.

▶ Task 2: Generate a Table view of the Deleted Item retention settings of all mailbox databases in the Exchange organization

1. On LON-CAS1, run the following in the Windows PowerShell® Integrated Scripting Environment (ISE) Console pane:

```
Get-MailboxDatabase
```

2. In the Console pane, run the following:

```
Get-MailboxDatabase | Format-Table Name, DeletedItemRetention
```

MCT USE ONLY. STUDENT USE PROHIBITED

▶ Task 3: Create a Windows PowerShell job to return the five most recent events from the Application Event log on LON-CAS1 and LON-MBX1

1. On LON-CAS1, run the following command in the Windows PowerShell ISE Console pane:

```
$job = Start-Job –ScriptBlock { Get-EventLog –LogName Application –Newest 5 –Computer
LON-CAS1, LON-MBX1}
```

2. Run the following in the Console pane:

```
Get-Job
```

3. Run the following in the Console pane:

```
Receive-Job $job
```

Results: After completing this lab, you will have used the Exchange Management Shell and performed basic management tasks.

MCT USE ONLY. STUDENT USE PROHIBITED

Exercise 2: Using Exchange Management Shell to Manage Recipients

▶ Task 1: Create a list of all of the cmdlets that are available to manage mailbox objects

1. On LON-CAS1, run the following in the Windows PowerShell ISE Console pane:

```
Get-Command
```

2. Run the following in the Console pane:

```
Get-Command *Mailbox*
```

▶ Task 2: Modify a script to create new mailboxes from a CSV file

1. On LON-CAS1, browse to the **Start** screen, and then type **Notepad.exe**. Press Enter.
2. In the Notepad window, on the **File** menu, click **Open**. Navigate to **E:\Labfiles\Mod09\labfiles\AddUsers.Csv**. You will need to click the **File type** drop-down list, and then click **All Files**.
3. Modify the line beginning with Jim to read: **Jim Daly,Jim,Daly,Jim Daly,JimD,New York,JimD@adatum.com**.
4. Click **File**, and then click **Save As**.
5. In the Save As window type **AddConsultants.csv** in the **File name** field. Then click the **File type** drop-down list, and then click **All Files**. Click **Save**.
6. Close Notepad.
7. In the Windows PowerShell ISE, click **File**, and then click **Open**. Navigate to **E:\Labfiles\Mod09\labfiles\AddUsers.ps1**. Click **Open**.
8. Modify the line beginning with $OU to read: **$OU = "Consultants"**.
9. Modify the line beginning with $password to read: **$password = ConvertTo-SecureString "Pa$$w0rd!" –AsPlainText –Force**.
10. Modify the line beginning Import-CSV to read: **Import-CSV E:\Labfiles \Mod09\Labfiles\AddConsultants.CSV | Where-Object {New-Mailbox –Alias $_.SamAccountName `**.
11. Modify the line beginning -OrganizationalUnit to read: **-OrganizationalUnit $OU –UserPrincipalName $_.UserPrincipalName –Password $password `**.
12. Click the **File** menu, and then click **Save As**.
13. In the Save As window, type **AddConsultants.ps1** in the **File name** field, and then click **Save**.
14. Press F5 to run the AddConsultants.ps1 script.
15. In the Console pane, type the following to verify that accounts for Darren Waite, Ioannis Xylara, and Marko Zajc are created:

```
Get-ADUser –Filter * –SearchBase "OU=consultants,DC=Adatum,DC=com"
```

MCT USE ONLY STUDENT USE PROHIBITED

▶ Task 3: Configure settings on mailboxes by using the Exchange Management Shell

1. On LON-CAS1, run the following in the Windows PowerShell ISE Console pane:

```
Get-Mailbox –OrganizationalUnit Research | Set-Mailbox -Office "Research"
```

2. Run the following in the Console pane:

```
Get-DistributionGroupMember IT | Set-Mailbox –Mailtip “If you require IT assistance
please contact the Help Desk.”
```

Results: After completing this lab, you will have used the Exchange Management Shell and performed basic recipient management tasks.

▶ To prepare for the next module

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

1. On the host computer, start **Microsoft® Hyper-V® Manager**.
2. In the **Virtual Machines** list, right-click **20342B-LON-CAS1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20342B-LON-MBX1** and **20342B-LON-DC1**.
5. In Hyper-V Manager, click **20342B-LON-DC1**, and in the Actions pane, click **Start**.
6. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
7. Sign in using the following credentials:
   - User name: **Adatum\Administrator**
   - Password: **Pa$$w0rd**
8. Repeat steps 5 to 7 for **20342B-LON-MBX1**, **20342B-LON-CAS1**, **20342B-LON-CAS2**, **20342B-LON-CL1**, **20342B-TREY-DC1**, and **20342B-TREY-EX1**. Sign in to the **TREY** virtual machines as **TreyResearch\Administrator** with the password of **Pa$$w0rd**.

MCT USE ONLY. STUDENT USE PROHIBITED

Module 10: Designing and Implementing Integration with Microsoft Exchange Online

Lab: Designing Integration with Exchange Online

Exercise 1: Designing Integration with Microsoft Exchange Online

▶ Task 1: Read and analyze the scenario requirements

- Read the exercise scenario, and then analyze the requirements from an integration perspective. Identify the configurations needed to satisfy the requirements.

▶ Task 2: Design a solution

Answers to questions:

- You need the following components:
    - The Directory Synchronization tool, in order to implement Active Directory® synchronization between the Active Directory Domain Services (AD DS) domain and Microsoft® Exchange Online.
    - Active Directory Federation Services (AD FS), in order to implement single sign-on (SSO).
    - The Hybrid Configuration Wizard, in order to configure a hybrid deployment that satisfies the rest of the requirements.
- The only existing server that you can use for message transport and Autodiscover is LON-CAS1, because it has the Client Access server role configured and is located in a site that is connected to the Internet. Before you can run the Hybrid Configuration Wizard, you need to plan for the following:
    - You must upgrade the Exchange Online tenant version to the latest version, because the version is currently less than 15.0.000.0.
    - You must publish the LON-CAS1 Client Access server in Microsoft Forefront® Threat Management Gateway, in order to resolve Autodiscover requests.
    - You need to install a publicly trusted digital certificate on LON-CAS1, in order to make Autodiscover and Microsoft Exchange Server federated delegation work with Exchange Online.
    - Because A. Datum currently uses a non-Microsoft email relay, you must either configure LON-CAS1 or replace the non-Microsoft email relay with an Edge Transport server to route messages between the Exchange Server on-premises and Exchange Online.
    - If A. Datum were running Microsoft Exchange Server 2010 instead of Microsoft Exchange Server 2013, you would need to add an Exchange Server 2013 Client Access and Mailbox server role to the Exchange Server organization. You would also need to consider all the other requirements in the previous answer.

MCT USE ONLY. STUDENT USE PROHIBITED

▶ Task 3: Discuss your solution with the class

- Present your proposed solution to the class. Discuss alternative solutions with other students and with the instructor.

Results: After completing this exercise, you will have successfully:

Read and analyzed the scenario requirements.

Designed a solution.

Discussed your solution with the class.

MCT USE ONLY. STUDENT USE PROHIBITED

Module 11: Designing and Implementing Messaging Coexistence

Lab: Implementing Messaging Coexistence

Exercise 1: Implementing Message Routing Coexistence

▶ Task 1: Implementing a cross-forest certification authority (CA) trust

1. On LON-DC1, open **Server Manager**, click **Tools**, and then click **Certification Authority**.
2. In the certsrv console, right-click **Adatum-LON-DC1-CA**, and then click **Properties**.
3. On the **General** tab, click **View Certificate**.
4. In the Certificate window, click the **Details** tab.
5. Click **Copy to File**.
6. In the Certificate Export Wizard, click **Next**.
7. On the **Export File Format** page, ensure that **DER encoded binary X.509 (.CER)** is selected, and then click **Next**.
8. On the **File to Export** page, type **C:\AdatumRoot.cer**, and then click **Next**.
9. Click **Finish**, and then click **OK** three times.
10. Close the certsrv console.
11. Switch to **TREY-DC1**.
12. Click **Start**, navigate to **Administrative Tools**, and then click **Certification Authority**.
13. In the certsrv console, right click **TreyResearch-TREY-DC1-CA**, and then click **Properties**.
14. On the **General** tab, click **View Certificate**.
15. In the Certificate window, click the **Details** tab.
16. Click **Copy to File**.
17. In the Certificate Export Wizard, click **Next**.
18. On the **Export File Format** page, ensure that **DER encoded binary X.509 (.CER)** is selected, and then click **Next**.
19. On the **File to Export** page, type **C:\TreyRoot.cer**, and then click **Next**.
20. Click **Finish**, and then click **OK** three times.
21. Close the certsrv console.
22. On Trey-DC1, open **Windows® Explorer**. Navigate to **C:**, right-click the file **TreyRoot.cer**, and then click **Copy**.
23. In the address bar, type **\\172.16.0.10\C$**, and then press Enter.
24. If prompted, in **user name** type **Adatum\Administrator**, and then in **password** type **Pa$$w0rd**.
25. Right-click in the **C:** folder, and then click **Paste**.
26. Right-click the **AdatumRoot.cer** file, and then click **Copy**.
27. Navigate to **C:**, right-click in the window, and then click **Paste**.

MCT USE ONLY. STUDENT USE PROHIBITED

28. On Trey-DC1, click **Start**, select **Administrative Tools**, and then click **Group Policy Management**.
29. In the Group Policy Management window, expand **Forest: treyresearch.net**, expand **Domains**, expand **treyresearch.net**, right-click **Default Domain Policy**, and then click **Edit**.
30. In Group Policy Management Editor, expand **Computer Configuration**, expand **Policies**, expand **Windows Settings**, expand **Security Settings**, expand **Public Key Policies**, and then right-click **Trusted Root Certification Authorities**. Select **Import**.
31. On the **Welcome to the Certificate Import Wizard** page, click **Next**.
32. On the **File to Import** page, click **Browse**.
33. Navigate to **C:\AdatumRoot.cer**, select the file, and then click **Open**.
34. Click **Next**.
35. On the **Certificate Store** page, click **Next**, and then click **Finish**.
36. Click **OK**.
37. Close Group Policy Management Editor.
38. Close Group Policy Management Console.
39. Switch to **LON-DC1**.
40. In Windows® 8 Server Manager, click **Tools**, and then click **Group Policy Management**.
41. In the Group Policy Management window, expand **Forest: Adatum.com**, expand **Domains**, expand **Adatum.com**, right-click **Default Domain Policy**, and then click **Edit**.
42. In Group Policy Management Editor, expand **Computer Configuration**, expand **Policies**, expand **Windows Settings**, expand **Security Settings**, expand **Public Key Policies**, and then right-click **Trusted Root Certification Authorites**. Select **Import**.
43. On the **Welcome to the Certificate Import Wizard** page, click **Next**.
44. On the **File to Import** page, click **Browse**.
45. Navigate to **C:\TreyRoot.cer**, select the file, and then click **Open**.
46. Click **Next**.
47. On the **Certificate Store** page, click **Next**, and then click **Finish**.
48. Click **OK**.
49. Close Group Policy Management Editor.
50. Close Group Policy Management Console.
51. Switch to **LON-CAS1**. Click **Windows PowerShell**®, type **gpupdate /force**, and then press Enter. After Group Policy is refreshed, close the Windows PowerShell window.
52. Switch to **TREY-EX1**. Open **Windows PowerShell**, type **gpupdate /force**, and then press Enter. After Group Policy is refreshed, close the Windows PowerShell window.

▶ Task 2: Creating conditional forwarders and mail exchanger (MX) resource records

1. On LON-DC1, open **Server Manager**, click **Tools**, and then click **DNS**.
2. In the DNS Manager, expand **Forward Lookup Zones**, and then click the **Adatum.com** zone object.
3. Right-click **Adatum.com**, and then click **New Mail Exchanger (MX)**.
4. In the New Resource Record window, in **Fully qualified domain name (FQDN) of mail server**, type **lon-cas1.adatum.com**, and then click **OK**.

MCT USE ONLY. STUDENT USE PROHIBITED

5. Select and then right-click the **Conditional Forwarders** node, and then click **New Conditional Forwarder**.
6. In the New Conditional Forwarder window, in **DNS Domain**, type **treyresearch.net**.
7. Click **Click here to add an IP Address or DNS Name**, type **172.16.0.100**, press Enter, and then click **OK**.
8. Close DNS Manager.
9. Switch to **TREY-DC1**. Click **Start**, navigate to **Administrative Tools**, and then click **DNS**.
10. In the DNS Manager console, expand **Forward Lookup Zones**, and then click the **treyresearch.net** zone object.
11. Right-click **treyresearch.net**, and then click **New Mail Exchanger (MX)**.
12. In the New Resource Record window, in **Fully qualified domain name (FQDN) of mail server**, type **trey-ex1.treyresearch.net**, and then click **OK**.
13. Select and then right-click the **Conditional Forwarders** node, and then click **New Conditional Forwarder**.
14. In the New Conditional Forwarder window, in **DNS Domain**, type **Adatum.com**.
15. Click **Click here to add an IP Address or DNS Name**, type **172.16.0.10**, press Enter, and then click **OK**.
16. Close DNS Manager.
17. On Trey-DC1, from the task bar, open the **Windows PowerShell** command-line interface.
18. In Windows PowerShell, type **ping adatum.com**, and then press Enter.
19. Ensure that you receive a reply from 172.16.0.10. Close Windows PowerShell. (**Note**: If you do not receive a ping reply, type **ipconfig /flushdns**, and then try again).
20. Switch to **LON-DC1**.
21. From the task bar, open **Windows PowerShell**.
22. In Windows PowerShell, type **ping treyresearch.net**, and then press Enter.
23. Ensure that you receive a reply from 172.16.0.100.
24. Close Windows PowerShell.
25. Switch to **LON-CAS1**.
26. Open **Windows Internet Explorer**®, type **https://trey-ex1.treyresearch.net/owa**, and then press Enter.
27. Ensure that you do not receive a certificate trust warning message in Internet Explorer, and that Microsoft® Outlook® Web App opens.
28. Close Internet Explorer.
29. Switch to **TREY-EX1**.
30. Open **Internet Explorer**, type **https://lon-cas1.adatum.com/owa**, and then press Enter.
31. Ensure that you do not receive a certificate warning message and that Outlook Web App opens.
32. Close Internet Explorer.

MCT USE ONLY. STUDENT USE PROHIBITED

▶ Task 3: Establishing a forest trust

1. On LON-DC1, open **Server Manager**, click **Tools**, and then click **Active Directory Domains and Trusts**.
2. Right click the **Adatum.com** domain object, and then click **Properties**.
3. In the Adatum.com Properties window, click the **Trusts** tab.
4. Click **New Trust**.
5. In the New Trust Wizard, on the **Welcome** page, click **Next**.
6. On the **Trust Name** page, in **Name**, type **treyresearch.net**, and then click **Next**.
7. On the **Trust Type** page, click **Forest trust**, and then click **Next**.
8. On the **Direction of Trust** page, ensure that **Two-way** is selected, and then click **Next**.
9. On the **Sides of Trust** page, click **Both this domain and the specified domain**, and then click **Next**.
10. On the **User Name and Password** page, in **User name** type **Administrator@treyresearch.net**, and then in **Password** type **Pa$$w0rd**. Click **Next**.
11. On the **Outgoing Trust Authentication Level – Local Forest** page, click **Forest-wide authentication**, and then click **Next**.
12. On the **Outgoing Trust Authentication Level – Specified Forest** page, click **Forest-wide authentication**, and then click **Next**.
13. On the **Trust Selections Complete** page, click **Next**.
14. On the **Trust Creation Complete** page, click **Next**.
15. On the **Confirm Outgoing Trust** page, click **Yes, confirm the outgoing trust**, and then click **Next**.
16. On the **Confirm Incoming Trust** page, click **Yes, confirm the incoming trust**, and then click **Next**.
17. On the **Completing the New Trust Wizard** page, click **Finish**.
18. In the Adatum.com Properties window, click **OK**.
19. Close the Active Directory Domains and Trusts console.

▶ Task 4: Create send and receive connectors

1. On LON-CAS1, open the Exchange admin center at **https://LON-CAS1.adatum.com/ecp**, and then log on as **Adatum\Administrator** with the password **Pa$$w0rd**.
2. In the feature pane, click **mail flow**.
3. Click the **send connectors** tab.
4. Click **New**.
5. In the new send connector window, in **Name**, type **treyresearch.net**. In **Type**, click **Partner**. Click **next**.
6. On the **Network settings** page, click **next**.
7. On the **Address space** page, click **Add**.
8. In the Address Space window, in **Full Qualified Domain Name (FQDN)**, type **treyresearch.net**, and then click **save**.
9. Click **next**.
10. On the **Source server** page, click **Add**.

MCT USE ONLY. STUDENT USE PROHIBITED

11. In the Select a Server window, click **LON-MBX1**, and then click **add**.
12. Click **ok**.
13. Click **finish**.
14. Double click **treyresearch.net** connector.
15. In Exchange Send Connector window select option **Proxy through client access server.**
16. Click **save**.
17. Click the **receive connectors** tab.
18. In the **Select server** drop-down box select **LON-CAS1.adatum.com.**
19. Click **New**.
20. In the new receive connector window, in **Name**, type **treyresearch.net**, and then click **Partner**. Click **next**.
21. On the **network adapter bindings** page, click **next**.
22. On the **Remote network settings** page, click **Add**.
23. On the **add IP address** page, type **172.16.0.101**, and then click **save**.
24. Select the range **0.0.0.0-255.255.255.255** and then click **Remove**.
25. Click **finish**.
26. In Exchange admin center, click **servers** and then click **certificates** tab.
27. In **Select server** drop-down list select **LON-CAS1.Adatum.com**.
28. Double click **Webmail.adatum.com** certificate.
29. In Webmail.adatum.com window, click services.
30. Select **SMTP** checkbox.
31. Click **save**. On prompt, click **Yes**.
32. Open **Exchange Management Shell** on LON-CAS1.
33. Type: **Set-TransportConfig –TLSSendDomainSecureList adatum.com** and press Enter.
34. Type: **Set-TransportConfig –TLSReceiveDomainSecureList treyresearch.net** and press Enter.
35. Switch to **TREY-EX1**, and then open the **Exchange Management Console**.
36. Expand **Microsoft Exchange On-Premises**, expand **Organization Configuration**, and then click **Hub Transport**.
37. Click the **Send Connectors** tab.
38. In the Actions pane, click **New Send Connector**.
39. In the New Send Connector window, in **Name**, type **Adatum Send Connector**.
40. In the **Select the intended use for this Send connector** list, click **Partner**, and then click **Next**.
41. On the **Address space** page, click **Add**.
42. In the **Address space** (for example, contoso.com), type **adatum.com**, click **OK**, and then click **Next**.
43. On the **Network settings** page, click **Next**.
44. On the **Source Server** page, ensure that **TREY-EX1** is listed, and then click **Next**.

MCT USE ONLY. STUDENT USE PROHIBITED

45. On the **New Connector** page, click **New**, and then click **Finish**.
46. Click **Server Configuration**.
47. On the **Exchange Certificates** tab, click **Trey Mail Certificate**, and click **Assign Services to Certificate**.
48. On the **Select Servers** page, click **Next**.
49. On the **Select Services** page, select the **Simple Mail Transfer Protocol (SMTP)** check box, and click **Next**.
50. Click **Assign**. When prompted to overwrite existing SMTP certificate click **Yes**, and then click **Finish**.
51. In the Exchange Management Console, expand **Server Configuration**, click **Hub Transport**, and then, in the Hub Transport pane, click **TREY-EX1**.
52. In the Actions pane, click **New Receive Connector**.
53. In the New Receive Connector window, in **Name**, type **Adatum.com Receive Connector**.
54. In the **Select the intended use for this Receive connecto**r list, click **Partner**, and then click **Next**.
55. On the **Local Network settings** page, click **Next**.
56. On the **Remote Network settings** page, click the **Remove** icon (the red X) to delete the entry, and then click **Add**.
57. In **Address or address range**, type **172.16.0.20**, click **OK**, and then click **Next**.
58. On the **New Connector** page, click **New**, and then click **Finish**.
59. Open **Exchange Management Shell** on TREY-EX1.
60. Type: **Set-TransportConfig –TLSSendDomainSecureList treyresearch.net** and press Enter.
61. Type: **Set-TransportConfig –TLSReceiveDomainSecureList adatum.com** and press Enter.
62. Close Exchange Management Shell.

▶ Task 5: Test the domain security between adatum and Trey Research

1. On LON-CL1, open **Microsoft Outlook 2013** and complete the profile creation (if needed) by clicking **Next** three times and then click **Finish**. If the First things first window appears, click **Ask me later** and then click **Accept**.
2. Open a command prompt and type **gpupdate /force** to refresh Group Policy.
3. Click **New Email**.
4. In the **To** field, type **Administrator@treyresearch.net**.
5. In **Subject**, type **test secure email**.
6. In the message body, type some text.
7. Click **Send**.
8. Open **Internet Explorer**.
9. Type **https://trey-ex1.treyresearch.net/owa**.
10. Sign in as **Treyresearch\Administrator** with the password **Pa$$w0rd**.
11. Ensure that you receive the message from the Adatum administrator.
12. Click **Reply**.
13. Type some text, and then click **Send**.

MCT USE ONLY. STUDENT USE PROHIBITED

14. Log off from Outlook Web App.
15. Switch to **Outlook 2013**.
16. Ensure that you receive the message from the Trey Research administrator. Also, ensure that the message has a green check mark. Click the green check mark, read the text and click **Close**. (**Note**: If you don't receive any messages within 1-2 minutes, go to the next step. Otherwise, proceed directly to step 21)
17. On TREY-EX1 machine, in the Exchange Management Shell, type **Restart-service MSExchangeTransport** and press Enter.
18. On LON-CAS1, in the Exchange Management Shell, type **Restart-service MSExchangeFrontEndTransport** and press Enter.
19. On LON-MBX1 machine, in the Exchange Management Shell, type the following:
    - **Restart-service MSExchangeSubmission**, and press Enter,
    - **Restart-service MSExchangeDelivery**, and press Enter,
    - **Restart-service MSExchangeTransport**, and press Enter.
20. Wait for a minute or two and verify if messages are delivered.
21. Using Outlook 2013, send a few email messages with whatever content you like to **Cindy@treyresearch.net**.
22. Close Outlook 2013.
23. Open **Internet Explorer**.
24. Navigate to **https://trey-ex1.treyresearch.net/owa**.
25. Sign in as **Treyresearch\cindy** with the password **Pa$$w0rd**. Accept default values on the regional and language page.
26. Ensure that you receive the messages from Adatum administrator.
27. Close Internet Explorer.

Results: After completing this exercise, you will have successfully implemented message routing coexistence.

MCT USE ONLY. STUDENT USE PROHIBITED

Exercise 2: Migrating User Mailboxes

▶ Task 1: Prepare the source server for a cross-forest mailbox move

1. On TREY-EX1, open **Exchange Management Shell**.
2. Type **Get-WebServicesVirtualDirectory | FL**, and then press Enter.
3. Verify that the value of the **Name** attribute is **EWS (Default Web Site)** and that the value of the **MRSProxyEnabled** attribute is **false**.
4. Type **Set-WebServicesVirtualDirectory -Identity "EWS (Default Web Site)" -MRSProxyEnabled $true**, and then press Enter.
5. Click **Start**, navigate to **Administrative Tools**, and then click **Internet Information Services (IIS) Manager**.
6. Expand **Trey-EX1**, expand **Sites**, and then click **Default Web Site**.
7. Double-click **Handler Mappings**. Scroll through the list, and verify the presence of ***.svc** entries in the **Path** column. (**Note**: If you find entries for *.svc, proceed directly to step 15. If not, go to step 8.).
8. Open a Command Prompt window as an Administrator.
9. Change the path to **C:\Windows\Microsoft.Net\Framework\v3.0 \Windows Communication Foundation**.
10. Type **ServiceModelReg.exe –r**, and then press Enter.
11. Type **Y**, and then press Enter.
12. Switch to **IIS**.
13. Click **Default Web Site**.
14. Double-click **Handler Mappings**. Scroll through the list, and search for the ***.svc** in the **Path** column. You should find entries for ***.svc**.
15. Switch to the **Exchange Management Shell**.
16. Type **Test-MRSHealth**, and then press Enter.
17. Ensure that all three tests have the value **True** in the **Passed** row.

▶ Task 2: Prepare the object for the move

1. On LON-CAS1, from the **Start** screen, click **Exchange Management Shell**.
2. Change the path to "**C:\Program Files\Microsoft\Exchange Server\v15\scripts"**.
3. Type **$Local = Get-Credential**, and then press Enter.
4. In the Windows PowerShell Credential window, in **User name** type **Adatum\Administrator** and in **Password** type **Pa$$w0rd**. Click **OK**.
5. Type **$Remote= Get-Credential,** and then press Enter.
6. In the Windows PowerShell Credential window, for **User name** type **Treyresearch\Administrator** and for **Password** type **Pa$$w0rd**. Click **OK**.
7. Type **.\Prepare-MoveRequest.Ps1 -Identity Cindy@treyresearch.net -RemoteForestDomainController trey-dc1.treyresearch.net -RemoteForestCredential $Remote -LocalForestDomainController lon-dc1.adatum.com -LocalForestCredential $Local -TargetMailUserOU "OU=IT,dc=adatum,dc=com"**, and then press Enter.
8. Ensure that you receive the message **1 mailbox(es) ready to move**.

MCT USE ONLY. STUDENT USE PROHIBITED

9. Switch to **LON-DC1**. Open **Active Directory Users and Computers**.
10. Click the **IT** organizational unit.
11. Ensure that there is an object called **Cindy White** there and that it is disabled.
12. Close Active Directory Users and Computers.

▶ Task 3: Move the user object from Treyresearch to Adatum

1. On LON-CAS1, open the **Exchange** admin center.
2. In the feature pane, click **recipients**.
3. Click the **migration** tab.
4. Click the arrow next to the **New** icon (plus sign (+)).
5. Click **Move to this forest**.
6. In the new cross-forest mailbox move window, click **Add**.
7. In the Select Mail User window, click **Cindy White**, click **add**, and then click **ok**.
8. Click **next**.
9. On the **Enter on-premises account credentials** page, type **Treyresearch\administrator** for the Source forest administrator name(domain\administrator name) and **Pa$$w0rd** for the password. Click **next**.
10. On the **Confirm the migration endpoint** page, type **trey-ex1.treyresearch.net** in **Remote MRS Proxy Server** text box, and then click **next**. (**Note**: If you get an error that connection to trey-ex1.treyresearch.net cannot be made, restart TREY-EX1 machine and then try again.)
11. On the **Move configuration** page, in **New migration batch name**, type **Cindy**.
12. In the **Target database** section click **browse**.
13. In the Select Mailbox Database window click **Mailbox Database 1** and click **add->**. Click **ok**.
14. Click **More options**.
15. In both text boxes, type **10**.
16. Click **next**.
17. On the **Start the batch** page, ensure that Administrator is listed in the text box. If not, click **browse**, and then click **Administrator,** and then click **OK**.
18. Ensure that **Automatically start the batch** is selected. Click **new**.
19. Wait until **Status of Cindy object** becomes **Synced**. You can click **Refresh** a few times. It may take a few minutes.
20. In the tasks pane, click **Cindy**, and then click **Complete this migration batch**.
21. In the warning window, click **yes**.
22. Wait until the status of the Cindy object becomes **Completed**. It may take a few minutes. (**Note**: If you don't get status **Completed** in 5-6 minutes, restart the same services as in Exercise 1, Task 5, steps 17, 18 and 19, and then repeat this task.)

MCT USE ONLY. STUDENT USE PROHIBITED

▶ Task 4: Validate the move

1. Switch to **LON-DC1**. Open **Server Manager**, click **Tools**, and then click **Active Directory Users and Computers**.
2. Navigate to the **IT** organizational unit.
3. Right-click **Cindy White**, and then click **Reset Password**.
4. Type **Pa$$w0rd** in both text boxes, and then clear the **User must change password at next logon** check box.
5. Click **OK** twice.
6. Right-click the **Cindy White** user object, and then click **Enable Account**.
7. Click **OK**.
8. Open **Internet Explorer**.
9. Type **https://lon-cas1.adatum.com/owa**, and then press Enter.
10. Sign in as **Adatum\Cindy** with the password **Pa$$w0rd**.
11. Ensure that you sign in, and that you see all messages that this user received while they were in Trey Research.
12. Close Outlook Web App.
13. Switch to **TREY-EX1**. Open the **Exchange Management Console**.
14. Expand **Recipient Configuration**, and then click **Mailbox**.
15. Ensure that **Cindy White** is not there anymore.

▶ Task 5: To prepare for the next module

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

1. On the host computer, start **Microsoft Hyper-V® Manager**.
2. In the **Virtual Machines** list, right-click **20342B-LON-CAS1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20342B-LON-MBX1**, **20342B-LON-DC1**, **20342B-LON-CL1**, **20342B-TREY-DC1**, and **20342B-TREY-EX1**.

Results: After completing this exercise, students will have moved a mailbox between Microsoft Exchange Server organizations.

MCT USE ONLY. STUDENT USE PROHIBITED

Module 12: Designing and Implementing Microsoft Exchange Server Upgrades

Lab: Upgrading from Exchange Server 2010 to Exchange Server 2013

Exercise 1: Documenting the Exchange Server 2010 Organization

▶ Task 1: Document the Exchange Organization configuration

1. On TREY-EX1, click **Start**, click **All Programs**, click **Microsoft Exchange Server 2010**, and then click **Exchange Management Console**.
2. In the Exchange Management Console, expand **Microsoft Exchange On-Premises** (trey-ex1.treyresearch.net), expand **Organization Configuration**, and then click **Mailbox**.
3. On the **Database Management** tab, double-click **Mailbox Database 1**.
4. On the **Limits** tab, document the following settings, and then click **OK**.
   - Issue warning at (MB): 1945
   - Prohibit Send at (MB): 2048
   - Prohibit send and receive at (MB): Not configured
5. On the **Retention Policy Tags** tab, document the settings for the following retention policy tags:
   - TreyResearch – Business Critical
   - Tag type: Personal Tag
   - Age limit: 1095 days
   - Action: Move to Archive
   - TreyResearch – Default Delete
   - All other folders in the mailbox
   - Age limit: 1825 days
   - Action: Delete and Allow Recovery
   - TreyResearch – DefaultMovetoArchive
   - Tag type: All other folders in the mailbox
   - Age limit: 365 days
   - Action: Move to Archive
   - TreyResearch – Deleted Items

     Tag type: Deleted Items
   - Age limit: 30 day
   - Action: Permanently Delete

MCT USE ONLY. STUDENT USE PROHIBITED

6. On the **Retention Policies** tab, document the retention policy tags assigned to the following retention policies. Also, identify the organizational units for the recipients to which the policy is applied.
   - Default Policy
   - Retention Policy Tags: Never Delete, TreyResearch - Default Delete, TreyResearch – Deleted Items
   - Organizational Units: Production, IT, Research, Sales
   - Executive Group Retention
   - Retention Policy Tags: TreyResearch – Business Critical, TreyResearch – DefaultMovetoArchive, TreyResearch – Deleted Items

     Organizational Units: Executives
7. On the **Offline Address Book** tab, what are the **Generation Server** and **Distribution Mechanism** settings?

   TREY-EX1, Web-Based, Public Folders
8. Under Organization Configuration, click **Client Access**.
9. On the **Outlook Web App Mailbox Policies** tab, what are the **Change Password** settings for the **Default** and **Executives Policy**?
   - Default: Password changes disabled.
   - Executives Policy: Password changes enabled.
10. On the **Exchange ActiveSync Mailbox Policies** tab, what are the password settings for the **Executives EAS Policy**?
    - Require password: enabled
    - Enable password recovery: enabled
    - Allow simple password: enabled
    - Minimum password length: 6
    - Enforce password history: 0
11. Under **Organization Configuration**, click **Hub Transport**.
12. On the **Transport Rules** tab, double-click the **E-Mail Disclaimer** rule, and then document the settings.
    - Conditions: apply to all messages
    - Actions: attach a disclaimer.
    - Exceptions: None
13. On the **Journal Rules** tab, double-click the **Research Department Message Journaling** rule, and then document the settings.
    - Journal report email address: ResearchJournal@treyresearch.net
    - Scope: Global – all messages
    - Recipient: Research@treyresearch.net

MCT USE ONLY. STUDENT USE PROHIBITED

14. On the **Send Connectors** tab, document the settings for the **Internet Send Connector**.
    - Address Space: *
    - Network: Use DNS MX records
    - Source Server: TREY-EX1

▶ Task 2: Document the Exchange Server configuration

1. Under **Server Configuration**, click **Client Access**.
2. Double-click **owa (Default Web Site)**. Document the **External URL** and **Authentication** settings.
    - External URL: https://mail.treyresearch.net/owa
    - Authentication: Forms based authentication, user name only, logon domain is TreyResearch.net.
3. Under **Server Configuration**, click **Hub Transport**.
4. Document the permission group configuration for the **Default TREY-EX1** Receive Connector.
    - Permission groups include Anonymous users, Exchange users, Exchange servers, Legacy Exchange Servers

▶ Task 3: Document the public folder configuration

1. On TREY-EX1, click **Start**, click **All Programs**, click **Microsoft Exchange Server 2010**, and then click **Exchange Management Shell**.
2. In the Exchange Management Shell, type **Get-PublicFolder –recurse**, and press Enter. Document the public folder structure.
    - Departments public folder with IT, Research, and Sales subfolders.
3. Type **Get-PublicFolder –recurse | Get-PublicFolderClientPermission**, and press Enter. Document the public folder client permissions for the IT, Research, and Sales public folders:
    - IT: Default – Reviewer, Administrator – Owner, IT – Publishing Editor
    - Research: Default – None, Administrator – Owner, Research – Publishing Editor
    - Sales: Default – Reviewer, Administrator – Owner, Sales – Publishing Editor
4. Type **Get-PublicFolderStatistics** and press Enter. Document the item count in the IT, Research, and Sales public folders.
    - IT: 2
    - Sales: 2
    - Research: 2

▶ Task 4: Document the Exchange recipient configuration

1. In the Exchange Management Console, expand **Recipient Configuration** and then click **Mailbox**. Document the total number of mailboxes: 233
2. Double-click **Anders Riis**, who is a member of the Executive team, and document the following settings:
    - Archive mailbox: enabled
    - Retention policy (click **Mailbox Settings**, **Messaging Records Management**, **Properties**): Executive Group Retention
    - Storage Quota (click **Mailbox Settings**, **Storage Quotas**, **Properties**): Default

MCT USE ONLY. STUDENT USE PROHIBITED

- Outlook Web App mailbox policy (click **Mailbox Features**, **Outlook Web App**, **Properties**): Executives Policy
- Microsoft Exchange ActiveSync® mailbox policy (click **Mailbox Features**, **Exchange ActiveSync**, **Properties**): Executives EAS Policy

3. Click **Cancel**.
4. Double-click **Aaron Nicholls**, who is a member of the Production team, and then document the following settings:
   - Archive mailbox: disabled
   - Retention policy (click **Mailbox Settings**, **Messaging Records Management**, **Properties**): Default Policy
   - Storage Quota (click **Mailbox Settings**, **Storage Quotas**, **Properties**): Default
   - Outlook Web App mailbox policy (click **Mailbox Features**, **Outlook Web App**, **Properties**): None
   - Exchange ActiveSync mailbox policy (click **Mailbox Features**, **Exchange ActiveSync**, **Properties**): Default
5. Click **Cancel**.
6. Double-click **April Stewart**, who is a member of the Research team, and document the following settings:
   - Archive mailbox: disabled
   - Retention policy (click **Mailbox Settings, Messaging Records Management, Properties**): Default Policy
   - Storage Quota (click **Mailbox Settings**, **Storage Quotas**, **Properties**): Issue warning at 4000 MB, Prohibit send at 5000 MB)
   - Outlook Web App mailbox policy (click **Mailbox Features, Outlook Web App**, **Properties**): None
   - Exchange ActiveSync mailbox policy (click **Mailbox Features**, **Exchange ActiveSync**, **Properties**): Default
7. Click **Cancel**.
8. Right-click **Research Journal Mailbox**, and then click **Manage Full Access Permission**. Document the user mailbox with full access, and then click Cancel.
   - TREYRESEARCH\MailboxAuditor
9. Double-click **Mailbox Auditor**. On the **Member Of** tab, document the groups that the Mailbox Auditor account belongs to, and then click **Cancel**.
   - Discovery Management
10. Double-click **Kai Axford**. On the **Mailbox Settings** tab, double-click**Messaging Records Managemen**t. Verify that the **Enable Litigation Hold** check box is selected. Click **Cancel**.

Results: After completing this exercise, you will have documented the Microsoft® Exchange Server 2010 organization.

MCT USE ONLY. STUDENT USE PROHIBITED

Exercise 2: Deploying Exchange Server 2013

▶ Task 1: Preparing AD DS for the Exchange Server 2013 deployment

1. On TREY-EX13, in the Virtual Machine Connection window, click the **Media** menu, select **DVD Drive**, and then click **Insert Disk**.
2. Navigate to **D:\Program Files\Microsoft Learning\20342\Drives\ExchangeServer2013CU1.iso**, and then click **Open**.
3. Click to the desktop and then, on the task bar, click **Windows PowerShell**.
4. Type **Install-WindowsFeature RSAT-ADDS**, and then press Enter.
5. Wait for the feature to be installed, and then type **D:**, and then press Enter.
6. Type the following command, and then press Enter:

```
.\Setup.exe /PrepareAD /IAcceptExchangeServerLicenseTerms
```

7. Wait until the process completes.

▶ Task 2: Install Exchange Server 2013

1. On TREY-EX13, in the Windows® PowerShell® window, type **.\setup.exe** and press Enter.
2. On the **Check for Updates?** page, click **Don't check for updates right now**, and then click **next**. Wait until setup copies files and initializes the setup process.
3. On the **Introduction** page, click **next**.
4. On the **License Agreement** page, click **I accept the terms in the license agreement**, and then click **next**.
5. On the **Recommended Settings** page, click **next**.
6. On the **Server Role Selection** page, select **Mailbox role** and **Client Access role**, and then click **next**.
7. On the **Installation Space and Location** page, accept the default values, and click **next**.
8. On the **Malware Protection Settings** page, make sure that **No** is selected, and then click **next**.
9. On the **Readiness Checks** page, ensure that all prerequisites are met, and then click **install**.
10. Wait until the installation completes. It can take 30 to 40 minutes to finish. On the **Setup Completed** page, click **finish**.
11. Restart **TREY-EX13**, and sign in as **TreyResearch\Administrator** with the password **Pa$$w0rd.**

▶ Task 3: Verify a successful installation

1. On TREY-EX13, from the **Start** screen, open the **Exchange Management Shell**.
2. Type **Get-MailboxDatabase | Set-MailboxDatabase –Name EX13MDB1**, and then press Enter. This command renames the default mailbox database created during the Exchange Server installation.
3. Type **$password = Read-Host "Enter password" –AsSecureString**, and then press Enter.
4. Type **Pa$$w0rd**, and then press Enter.
5. Type **New-Mailbox -UserPrincipalName EX13Test@treyresearch.net -Alias EX13Test -Database EX13MDB1 -Name EX13Test -Password $password**, and then press Enter.
6. From the **Start** screen, click **Internet Explorer**.
7. In the Address bar, type **https://TREY-EX13.TreyResearch.net/owa**, and then press Enter.

MCT USE ONLY. STUDENT USE PROHIBITED

8. Sign in as **TreyResearch\EX13Test** with the password **Pa$$w0rd**.
9. At the **Outlook Web App** page, click **save**. Verify that Microsoft Outlook® Web App opens.
10. Click **new mail**.
11. Prepare a message to **Aaron Nicholls** with a subject of **Test from Exchange 2013**. Click **Send**.
12. On TREY-EX1, open **Windows Internet Explorer**® and connect to **https://TREY-EX1.treyresearch.net/owa**.
13. Sign in as **Aaron** using the password **Pa$$w0rd**. Verify that the email from the EX13Test account is received in the inbox.
14. Reply to the message.



Note: If you receive an error message that the server operation timed out, click **Close**.

15. On Trey-EX13, verify that EX13Test receives the reply from Aaron.
16. Close Internet Explorer.

Results: After completing this exercise, you will have deployed an Exchange 2013 server in the Trey Research Exchange organization.

MCT USE ONLY. STUDENT USE PROHIBITED

Exercise 3: Upgrading from Exchange Server 2010 to Exchange Server 2013

▶ Task 1: Move the Administrator mailbox to Exchange Server 2013

1. On TREY-EX13, in the Exchange Management Shell, type **New-MoveRequest –id Administrator –TargetDatabase EX13MDB1**, and then press Enter.
2. Wait a minute, and then type **Get-MoveRequest**, and then press Enter. Verify that the move request for the Administrator account has completed. If it is not complete, wait another minute and then run the command again.
3. Open **Internet Explorer** and connect to **https://TREY-EX13.TreyResearch.net/ecp**.
4. Sign in as **TreyResearch\Administrator** using the password **Pa$$w0rd**.
5. Verify that the Administrator can now access the Exchange Administration Center (EAC).
6. Connect to **https://TREY-EX13.TreyResearch.net/owa**.
7. In Outlook Web App, send a message to Aaron.
8. On TREY-EX1, in Outlook Web App, verify that Aaron receives the message.

▶ Task 2: Obtain a certificate for the Exchange 2013 server

1. On TREY-EX13, connect to **https://TREY-EX13.TreyResearch.net/ecp**.
2. If required, sign in as **TreyResearch\administrator** with the password **Pa$$w0rd**.
3. In the EAC, in the left navigation pane, click **servers**.
4. In the right pane, click **certificates**.
5. Click **New** to create a new certificate request.
6. In the Exchange Certificate – Windows Internet Explorer window, in the new Exchange certificate wizard, select **Create a request for a certificate from a certification authority**, and then click **next**.
7. In the **Friendly** name for this certificate, type **mail.TreyResearch.net**, and click **next**.
8. On the page with the option for using wildcard certificates, do not make any changes, and click **next**.
9. Click **browse**.
10. In the Select a Server window, click **TREY-EX13**, and click **ok**.
11. Click **next**.
12. On the next page, click **Outlook Web App (when accessed from the Internet)**, and then click the **Edit** icon.
13. In the **Specify the domains for the above Access type**, enter **mail.TreyResearch.net**, and click **OK**.
14. Repeat steps 12 and 13 for items where **<not specified>** is in the **DOMAIN** column.
15. Click **next**.
16. On the next page, make sure that you have the following names in the list: **mail.TreyResearch.net**, **TREY-EX13.TreyResearch.net**, **AutoDiscover.TreyResearch.net, TREY-EX13**, and **TreyResearch.net**, and then click **next**.
17. On the next page, fill in the following fields as follows:
    a. Organization name: **TreyResearch**
    b. Department name: **IT**
    c. City/Locality: **London**

MCT USE ONLY. STUDENT USE PROHIBITED

d. State/Province: **England**

e. Country/Region name: **United Kingdom**

18. Click **next**.
19. On the next page, type **\\TREY-EX13\C$\users\administrator.treyresearch\downloads \certreq.req,** and click **finish**. Close Internet Explorer.
20. On TREY-EX13, open **File Explorer**, and navigate to **C:\users\administrator.treyresearch\downloads**.
21. Right-click **CertReq.req**, and then click **Open with**.
22. In the **Windows** dialog box, click **More options**, and then click **Notepad**.
23. In the CertReq.req – Notepad window, press Ctrl+A to select all the text, and then press Ctrl+C to copy and save the text to the clipboard. Close Notepad.
24. Click to the **Start** screen, and then click **Internet Explorer**.
25. Connect to **http://TREY-DC1.TreyResearch.net/certsrv**.
26. Sign in as **Administrator**, using the password **Pa$$w0rd**.
27. The browser displays a message that it does not support the generation of certificate requests. Press F12.
28. In the **Browser Mode** drop down list, click **Internet Explorer 10 Compatibility View**. Close the bottom tab.
29. On the **Welcome** page, click **Request a certificate**.
30. On the **Request a Certificate** page, click **advanced certificate request**.
31. On the **Advanced Certificate Request** page, click **Submit a certificate request by using a base-64-encoded CMC or PKCS#10 file, or submit a renewal request by using a base-64-encoded CMC or PKCS#7 file**.
32. On the **Submit a Certificate Request or Renewal Request** page, click in the **Saved Request** field, and then press Ctrl+V to paste the certificate request information into the field.
33. In the **Certificate Template** drop-down list box, click **Trey Web**, and then click **Submit**. If a Web Access Confirmation dialog box appears, click **Yes**.
34. On the **Certificate Issued** page, click **Download certificate**.
35. In the **File Download** dialog box, click **Save**.
36. On TREY-EX13, connect to the EAC. Sign in as **TreyResearch\Administrator** using the password **Pa$$w0rd**.
37. Click **servers**, and then click **certificates**.
38. Click on **mail.TreyResearch.net**, and then click **Complete**.
39. Type **\\TREY-EX13\C$\users\administrator.treyresearch\downloads\certnew.cer** and click **ok**.
40. Click **mail.TreyResearch.net**, and click **Edit** on the toolbar.
41. Click **services**.
42. Select **IIS**, and click **save**.

MCT USE ONLY. STUDENT USE PROHIBITED

▶ Task 3: Change the Client Access configuration to use Exchange Server 2013

1. In the EAC, click **servers** in the left pane, and then click **virtual directories**.
2. Double-click the **owa (Default Web Site)** virtual directory located on TREY-EX13.
3. On the **owa (Default Web Site)** page, in the **External URL** field, type **https://mail.treyresearch.net/owa**. Click **save**, and then click **ok** to dismiss the warning dialog box.
4. Double-click the **ecp (Default Web Site)** virtual directory located on TREY-EX13.
5. On the **ecp (Default Web Site)** page, in the **External URL** field, type **https://mail.treyresearch.net/ecp**. Click **save**.
6. Double-click the **EWS (Default Web Site)** virtual directory located on TREY-EX13.
7. On the **EWS (Default Web Site)** page, in the **External URL** field, type **https://mail.treyresearch.net/EWS/Exchange.asmx**. Click **save**.
8. Double-click the **Microsoft-Server-ActiveSync (Default Web Site)** virtual directory located on TREY-EX13.
9. On the **Microsoft-Server-ActiveSync (Default Web Site)** page, in the **External URL** field, type **https://mail.treyresearch.net/Microsoft-Server-ActiveSync**. Click **save**.
10. Double-click the **OAB (Default Web Site)** virtual directory located on TREY-EX13.
11. On the **OAB (Default Web Site)** page, in the **External URL** field, type **https://mail.treyresearch.net/OAB**. Click **save**.
12. Close Internet Explorer.
13. In the Exchange Management Shell, type **iisreset**, and then press Enter. Wait for the service to restart.
14. On TREY-DC1, click **Start**, click **Administrative Tools**, and then click **DNS**.
15. Expand **Forward Lookup Zones**, and then expand **TreyResearch.net**.
16. Double-click the **Mail** resource record.
17. Change the **IP address** to **172.16.0.102**, and then click **OK**.
18. In the navigation pane, right-click **TREY-DC1**, and click **Clear Cache**. Close the DNS Manager.
19. On TREY-EX13, in the Exchange Management Shell, type **nslookup mail.treyresearch.net** and press Enter.
20. Verify that the command returns the IP address **172.16.0.102**.
21. Type **ipconfig /flushdns** and press Enter.
22. Open **Internet Explorer** and connect to **https://mail.treyresearch.net/owa**. Verify that the **Exchange Server 2013 Outlook Web App** page appears.
23. Sign in as **TreyResearch\Administrator** using the password **Pa$$w0rd**. Verify that the Administrator can access the Exchange 2013 mailbox. Close Internet Explorer.
24. Open **Internet Explorer** and connect to **https://mail.treyresearch.net/owa**.
25. Sign in as **TreyResearch\Aaron** using the password **Pa$$w0rd**. Verify that the Aaron can access his Exchange 2010 mailbox. Close Internet Explorer.

MCT USE ONLY. STUDENT USE PROHIBITED

▶ Task 4: Move the IT department mailboxes to Exchange Server 2013

1. On TREY-EX13, in the Exchange Management Shell, type the following command, and then press Enter.

```
get-mailbox –organizationalunit IT | New-MoveRequest –TargetDatabase EX13MDB1
```

2. Leave the Exchange Management Shell open.

▶ Task 5: Configure and verify message routing

1. Connect to **Outlook Web App** and sign in as **TreyResearch\Administrator** using the password **Pa$$w0rd**.
2. Create a new mail for Kai Axford with the subject heading **Message before migration**. Kai Axford is a member of the Research group, which has a journaling rule configured. Kai's mailbox also has litigation hold enabled. Send the message.
3. Connect to the EAC. In the left pane, click on **mail flow**, and then click **delivery reports**.
4. Beside **Mailbox to search**, click **browse**. Click **Administrator**, and then click **OK**.
5. Beside **Search for messages sent to**, click **select users**. Click **Kai Axford**, click **add**, and then click **ok**.
6. Click **search**.
7. Double-click the message sent to Kai Axford and verify that the message was delivered successfully. Click **close**.
8. Beside **Mailbox to search**, click **browse**. Click **Kai Axford**, and then click **OK**.
9. Click **Search for messages received from**, and then click **select a user**. Click **Administrator**, click **add**, and then click **ok**.
10. Click **search**.
11. Double-click the message sent to Kai Axford and verify that the message was tracked successfully. Click **close**.
12. Click **send connectors**.
13. Double-click **Internet Send Connector**. On the **scoping** tab, under Source server, click **Add**.
14. Click **TREY-EX13**, click **add**, and then click **ok**.
15. Click **save**.
16. Click **receive connectors**.
17. In the **Select server** list, click **TREY-EX13.TreyResearch.net**.
18. Double-click **Default Frontend Trey-EX13**.
19. On the **security** tab, verify that the **Anonymous users** check box is selected. This receive connector is configured by default to allow anonymous connections from external SMTP servers. Click **cancel**.

▶ Task 6: Move the remaining mailboxes to Exchange Server 2013

1. In the Exchange Management Shell, type **Get-MoveRequest**, and press Enter. Verify that all mailboxes have been successfully moved. If the accounts show Queued, wait for 5 minutes and rerun the command until it shows Completed.
2. Type **Get-Mailbox –Arbitration -Database "Mailbox Database 1" | New-MoveRequest –TargetDatabase EX13MDB1**, and press Enter.

MCT USE ONLY. STUDENT USE PROHIBITED

3. In the EAC, click **recipients**, and then click the **migration** tab.
4. Click **New**, and then click **Move to a different database**.
5. In the new local mailbox move window, under **Select the users that you want to move**, click **Add**.
6. In the Select Mailbox window, sort the view by clicking the **Database** column.
7. Press the Shift key, and then select all mailboxes in the **Mailbox Database 1** database. Click **add**, and then click **ok**.
8. Click **next**.
9. In the New Migration Batch window, in the **New migration batch name** box, type **CompleteMigration**.
10. Under **Target database**, click **browse**, click **EX13MDB1**, click **add**, and then click **ok**.
11. Under **Target archive database**, click **browse**, click **EX13MDB1**, click **add**, and then click **ok**.
12. Click **next**, and then click **new**.
13. Under **Mailbox status**, click **view details**. Review the information, and click **close**.
14. On the **migration** tab, click **Status For All Batches**.
15. Review the information, then in the Status for All Batches window, click **Close**.
16. The migration will take some time to finish.

▶ Task 7: Migrate public folders to Exchange Server 2013

1. On TREY-EX13, in the Exchange Management Shell, type **Get-PublicFolderMigrationRequest** and press Enter. This command verifies that there are no outstanding public folder migration requests.
2. Type **Get-Mailbox -PublicFolder** and press Enter. This command verifies that there are no public folder mailboxes on the Exchange 2013 Mailbox server.
3. On TREY-EX1, open **Windows Explorer**.
4. On the **C:** drive, create a folder named **Migration**.
5. Right-click the folder, click **Share with**, and then click **Specific people**.
6. Click **Share**, and then click **Done**.
7. In the Windows Explorer path bar, type **\\TREY-EX13\C$\Program Files\Microsoft \Exchange Server\v15\Scripts** and press Enter.
8. Copy the following files from the Scripts folder on TREY-EX13 to the **C:\Migration** folder on TREY-EX1.
   - Export-PublicFolderStatistics.ps1
   - Export-PublicFolderStatistics.strings.psd1
   - PublicFolderToMailboxMapGenerator.ps1
   - PublicFolderToMailboxMapGenerator.strings.psd1
9. Open the **Exchange Management Shell**, type **cd C:\Migration**, and then press Enter.
10. Type **.\Export-PublicFolderStatistics.ps1 PFStats.csv TREY-EX1** and press Enter. This command exports the public folder statistics to a .CSV file.
11. In the **C:\Migration** folder, right-click **PFStats.csv**, click **Open**, click **Select a program from a list of installed programs**, click **OK**, and then click **Notepad**. Review the information and close the file.

MCT USE ONLY. STUDENT USE PROHIBITED

12. In the Exchange Management Shell, type **.\PublicFolderToMailboxMapGenerator.ps1 2000 C:\Migration\PFStats.csv PFtoMBXMap.csv** and press Enter.

Note: The value "2000" in the previous command specifies the maximum public folder mailbox size in bytes planned for the Exchange Server 2013 environment. This number does not set a limit on the mailbox size; it is only a value used by the script to determine how many public folder mailboxes will be required. In a production environment, this value would be much larger. The smaller number is used here so that the script will require more than one public folder mailbox on Exchange Server 2013.

13. In the **C:\Migration** folder, right-click **PFtoMBXMap.csv**, click **Open with**, and click **Notepad**.
14. Edit the target mailbox names by adding a PF to the mailbox name. For example, Mailbox1 should be changed to PFMailbox1. After changing all three mailbox names, save and close the file.
15. On TREY-EX13, in the Exchange Management Shell, type **New-Mailbox –PublicFolder PFMailbox1 –HoldForMigration** and press Enter.
16. Type **New-Mailbox –PublicFolder PFMailbox2** and press Enter.
17. Type **New-Mailbox –PublicFolder PFMailbox3** and press Enter.
18. Type **New-PublicFolderMigrationRequest –SourceDatabase (Get-PublicFolderDatabase –Server TREY-EX1) –CSVData (Get-Content \\TREY-EX1\Migration\PFToMBXMap.csv -Encoding Byte)** and press Enter.
19. Wait a minute, and then type **Get-PublicFolderMigrationRequest | Get-PublicFolderMigrationRequestStatistics** and press Enter. Verify that the StatusDetail is displayed as AutoSuspended, and PercentComplete is set to 95. If these values are not displayed, wait another minute and run the command again.
20. Type **Set-OrganizationConfig -PublicFoldersLockedForMigration:$true** and press Enter.
21. Type **Set-PublicFolderMigrationRequest -Identity \PublicFolderMigration -PreventCompletion:$false** and press Enter.
22. Type **Resume-PublicFolderMigrationRequest -Identity \PublicFolderMigration** and press Enter.
23. This request can take several minutes to finish. You can continue with the next steps while the migration finishes.

▶ Task 8: Verifying the upgrade of compliance features

1. On TREY-EX13, in the EAC, on the **recipients** tab, click **migration**.
2. Click **Status for all batches**. Verify that 189 mailboxes have been migrated. Click **close**.

Note: If not all of the mailboxes have been migrated, then you will need to wait until the migration is complete before attempting the following steps. You may see an error that there is one or more failed migrations. Verify that the failed mailboxes are system mailboxes. If that is the case, this failure can be ignored because it only indicates that the mailbox was already part of a move request.

3. Connect to **Outlook Web App** and, if required, sign in as **TreyResearch\Administrator** using the password **Pa$$w0rd**.

MCT USE ONLY. STUDENT USE PROHIBITED

4. Create a new mail for Kai Axford with the subject heading **Message after migration**. Kai Axford is a member of the Research group, which has a journaling rule configured. Kai's mailbox also has litigation hold enabled. Send the message.
5. On TREY-EX1, open **Internet Explorer** and connect to **https://mail.treyresearch.net/owa**.
6. Sign in as **Treyresearch\Kai** using the password **Pa$$w0rd**. Click **save**.
7. Verify that the message from the Administrator arrived and that it includes the email disclaimer configured by the transport rule configured in Exchange Server 2010.
8. Delete the two messages from the Administrator with the subjects **Message before migration** and **Message after migration**.
9. Right-click the **Deleted Items** folder and click **empty**. Click **ok**.
10. Right-click the **Deleted Items** folder and click **recover deleted items**.
11. In the recover deleted items window, hold the Ctrl key, click both messages, right-click and click **purge**. Click **ok**.
12. Close the Internet Explorer windows.
13. On TREY-EX13, in the Exchange Management Shell, type **Get-mailbox Discover* | FL Hidden*** and press Enter. Verify that the DiscoverySearchMailbox is hidden from the address lists.
14. Type **Set-mailbox Discover* -HiddenFromAddressListsEnabled $false** and press Enter. This step is required so that the Mailbox Auditor can open the DiscoverySearchMailbox from Outlook Web App.
15. On TREY-EX1, open **Internet Explorer** and connect to **https://mail.treyresearch.net/owa**.
16. Sign in as **Treyresearch\MailboxAuditor** using the password **Pa$$w0rd**. Click **save**.
17. In the top right corner, click **Mailbox Auditor**, and click **Open another mailbox**.
18. Type **Research Journal**, and press Enter. Click **open**, and then click **save**.
19. In the Research Journal Mailbox, verify that the two messages sent to Kai Axford are listed. Kai is a member of the Research group, and the messages sent to any member of the Research group are journaled to this mailbox.
20. On TREY-EX1, on the **Mailbox Auditor** tab, change the URL to **https://mail.treyresearch.net/ecp**.
21. Click **compliance management**. Since the Mailbox Auditor account is a member of the Discovery Management role group, the **in-place eDiscovery & hold** tab is available.
22. Click **New**. In the new in-place eDiscovery & hold window, type **Search Kai's mailbox** as the **Name**, and then click **next**.
23. On the **Mailboxes** page, verify that **Specify mailboxes to search** is selected, and then click **Add**.
24. Click **Kai Axford**, click **add**, click **ok**, and then click **next**.
25. On the **Search query** page, verify that **Include all user mailbox content** is selected, and then click **next**.
26. On the **In-Place Hold settings** page, click **finish**, and then click **close**.
27. Click the arrow beside the search icon, and click **Copy search results**.
28. Select the **Send me mail when the copy is completed** check box, and then click **Browse**.
29. Click **Discovery Search Mailbox**, click **ok**, and then click **Copy**. Click **ok**.
30. Click **refresh**, and check the status of the search. Wait a moment, and then click **refresh** again. Repeat until the search status is **Search Succeeded**. Close Internet Explorer.

MCT USE ONLY. STUDENT USE PROHIBITED

31. Open **Internet Explorer** and connect to **Outlook Web App**. Sign in as **TreyResearch\MailboxAuditor** using the password **Pa$$w0rd**.
32. Click **Mailbox Auditor**, and click **Open another mailbox**.
33. Type **Discover** and press Enter. Click **open**, and then click **save**.
34. In the Discovery Search Mailbox, expand the **Search Kai's mailbox** folder and subfolders.
35. Verify that the two messages purged by Kai are in the **Purges** folder. Kai's mailbox was placed on Litigation Hold in Exchange Server 2010, and the hold and all saved messages were retained during the migration.

▶ Task 9: Verifying additional upgrade components (optional)

- Review the configuration options that you documented in Exercise 1. If time permits, verify that all of the options have been upgraded to Exchange Server 2013.

Results: After completing this exercise, you will have completed the upgrade of all data and functionality to the Exchange 2013 server.

MCT USE ONLY. STUDENT USE PROHIBITED

Exercise 4: Removing Exchange Server 2010

▶ Task 1: Removing Exchange Server components

1. On TREY-EX1, in the Exchange Management Console, under **Organization Configuration**, click **Mailbox**.
2. On the **Offline Address Book** tab, verify that two versions of the Offline Address Book are listed. When Exchange Server 2013 is installed, a new Offline Address Book is created that is only distributed through the web-based mechanism.
3. Click **Default Offline Address Book**, click **Remove**, and then click **Yes**.
4. On TREY-EX13, in the Exchange Management Shell, type **Get-PublicFolderMigrationRequest | Get-PublicFolderMigrationRequestStatistics** and press Enter. Verify that the public folder migration is complete.
5. If the public folder migration is not complete, type **Suspend-PublicFolderMigrationRequest -Identity \PublicFolderMigration** and press Enter.
6. Type **Resume-PublicFolderMigrationRequest -Identity \PublicFolderMigration** and press Enter.
7. Wait a few minutes, and then type **Get-PublicFolderMigrationRequest | Get-PublicFolderMigrationRequestStatistics** and press Enter. Verify that the public folder migration is complete.
8. On TREY-EX1, open the **Exchange Management Shell**, type **Get-Mailbox –server TREY-EX1** and press Enter. Verify that no mailboxes are listed on the server.
9. Type **Get-Mailbox –server TREY-EX1 -arbitration** and press Enter. Verify that no arbitration mailboxes are listed on the server.
10. Type **Get-PublicFolder** and press Enter. Verify that the command returns an error. When you completed the public folder migration, the public folders on TREY-EX1 were no longer available.
11. On TREY-EX1, in the Exchange Management Console, under **Organization Configuration**, click **Hub Transport**.
12. On the **Send Connectors** tab, double-click **Internet Send Connector**.
13. On the **Source Server** tab, remove **TREY-EX1**. Click **OK**.
14. In the Exchange Management Console, under **Organization Configuration**, click **Mailbox**.
15. On the **Database Management** tab, right-click **Public Folder Database 1**, and then click **Dismount Database**. Click **Yes**.
16. Right-click **Mailbox Database 1** and click **Dismount Database**. Click **Yes**.
17. Right-click **Mailbox Database 1** and click **Remove**. Click **Yes**. Click **OK**.
18. Right-click **Public Folder Database 1**, and click **Remove**. Click **Yes**.
19. If you get an error message that the public folder still contains public folders, complete the following steps.
20. In the error message, click **OK**.
21. Click **Start**, point to **Administrative Tools**, and then click **ADSI Edit**.
22. Right-click **ADSI Edit**, and click **Connect To**.
23. In the **Connection Settings** dialog box, under **Select a well known Naming Context**, click **Configuration**. Click **OK**.

MCT USE ONLY. STUDENT USE PROHIBITED

24. Expand **CN=Configuration [TREY-DC1.TreyResearch.Net]**, **CN=Configuration,DC=TreyResearch,DC=net**, **CN=Services**, **CN=Microsoft Exchange**, **CN=TreyResearchOrg**, **CN=Administrative Groups**, **CN=Exchange Administrative Group (FYDIBOHF23SPDLT)**, and **CN=Databases**.
25. Right-click **CN=Public Folder Database 1**, and click **Delete**. Click **Yes** twice, and then close **ADSI Edit**.

▶ Task 2: Remove the Exchange 2010 server

1. On TREY-EX1, close the Exchange Management Console.
2. If you get a warning message, you must close all dialog boxes before closing the Exchange Management Console, click **OK**, and then complete the following three steps.
3. Right-click the task bar, and click **Start Task Manager**.
4. Click **Exchange Management Console**, and click **End Task**.
5. Click **End Now**, and then close Windows Task Manager.
6. Close the Exchange Management Shell.
7. Click **Start**, and then click **Control Panel**.
8. Click **Uninstall a program**.
9. Click **Microsoft Exchange Server 2010**, and then click **Uninstall**.
10. On the **Exchange Maintenance Mode** page, click **Next**.
11. On the **Server Role Selection** page, clear all check boxes, and then click **Next**.
12. On the **Readiness Checks** page, click **Uninstall**.
13. When the uninstallation finishes, click **Finish**.

Results: After completing this exercise, you will have removed Exchange Server 2010 from the Exchange organization.

▶ Shut down the virtual machines

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

1. On the host computer, start **Microsoft Hyper-V® Manager**.
2. In the **Virtual Machines** list, right-click **20342B-TREY-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20342B-TREY-EX1** and **20342B-TREY-EX13**.

MCT USE ONLY. STUDENT USE PROHIBITED

Notes

MCT USE ONLY. STUDENT USE PROHIBITED

Notes

MCT USE ONLY. STUDENT USE PROHIBITED