

Conversational Exchange

A ConversationalGeek® Book

Sponsored by **mimecast**



Learn about:

- The history of Exchange development down to modern times.
- The various features of Exchange: messaging services, high availability, etc...
- The value of using a third-party solution like Mimecast (security/archive/continuity).

By J. Peter Bruzzese (Microsoft Office 365 MVP)

Sponsored by Mimecast

Mimecast delivers cloud-based email management for Microsoft Exchange and Microsoft Office 365, including archiving, continuity and security. By unifying disparate and fragmented email environments into one holistic solution that is always available from the cloud, Mimecast minimizes the risk and reduces cost and complexity, while providing total end-to-end control of email.

Founded in the United Kingdom in 2003, Mimecast serves more than 10,000 customers worldwide with millions of users and has offices in Europe, North America, Africa, Australia and the Channel Islands.

www.mimecast.com



Conversational Exchange

By J. Peter Bruzzese

Copyright© 2017



ConversationalGeek

Conversational Exchange

Published by Conversational Geek Inc.

www.conversationageek.com

All rights reserved. No part of this book shall be reproduced, stored in a retrieval system, or transmitted by any means, electronic, mechanical, photocopying, recording, or otherwise, without written permission from the publisher. No patent liability is assumed with respect to the use of the information contained herein. Although every precaution has been taken in the preparation of this book, the publisher and author assume no responsibility for errors or omissions. Nor is any liability assumed for damages resulting from the use of the information contained herein.

Trademarks

Conversational Geek, the Conversational Geek logo and J. the Geek are trademarks of Conversational Geek. All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. We cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

Warning and Disclaimer

Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied. The information provided is on an “as is” basis. The author and the publisher shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or programs accompanying it.

Additional Information

For general information on our other products and services, or how to create a custom Conversational Geek book for your business or organization, please visit our website at ConversationalGeek.com

Publisher Acknowledgments

Cast and Crew of Conversational Exchange (in 10 days!)

Author:	J. Peter Bruzzese
Lead Technical Editor:	Theresa Miller
Lead Expert Reviewer:	Janet Vargas
Expert Reviewers:	Phoummala Schmitt Lasse Pettersson, MVP Steve Ahlgrim Paul Robichaux, MVP

The “Conversational” Method

We have two objectives when we create a “Conversational” book: First, to make sure it’s written in a conversational tone so that it’s fun and easy to read. Second, to make sure you, the reader, can immediately take what you read and include it into your own conversations (personal or business-focused) with confidence.

These books are meant to increase your understanding of the subject. Terminology, conceptual ideas, trends in the market, and even fringe subject matter are brought together to ensure you can engage your customer, team, co-worker, friend and even the know-it-all Best Buy geek on a level playing field.

“Geek in the Mirror” Boxes

We infuse humor into our books through both cartoons and light banter from the author. When you see one of these boxes it’s the author stepping outside the dialog to speak directly to you. It might be an anecdote, it might be a personal experience or gut reaction and analysis, it might just be a sarcastic quip, but these “geek in the mirror” boxes are not to be skipped.



Greetings. They call me J. Within these boxes I can share just about anything on the subject at hand. Read ‘em!

Dedication

To my fellow Exchange bloggers, evangelizers, and MVPs (past and present) in the hopes that many more can be educated about that which we have come to know and love: Exchange

Forward

The country comedian Minnie Pearl used to say, "It costs a lot of money to look this cheap." Email is similar -- it takes a tremendous amount of complexity to make it seem so simple. Explaining that complexity would take a book longer than this one. I've worked on email technology for a third of a century, and I'm still learning.

In fact, as J. Peter Bruzzese makes clear, it would take a much bigger book than this to explain all the complexities of Microsoft Exchange alone. But J. Peter has taken on a rather harder task: explaining the most important things about Microsoft Exchange, as simply as possible, to someone who actually has a life beyond email.

He has done this with clarity and humor. If you want to learn about Exchange, but you're not particularly looking forward to it, this is the book for you. It gives enough information for you to convince most people that you're an Exchange expert, and enough real knowledge to be able to talk intelligently to the serious gurus. It's also a great foundation for growing into an Exchange guru yourself over time.

Microsoft Exchange is an incredibly important piece of software. It's not the only email system in the world, but it dominates the business email market. Given the importance of email in today's business world, a whole ecosystem of support and service has grown up around Exchange, with thousands of people making a fine living as Exchange experts. Whether you're hoping to become one, or just to talk with the one you married, this is the place to start.

Nathaniel S. Borenstein
Chief Scientist, Mimecast

Table of Contents

Author Biography	1
Acknowledgements	2
Introduction from the Author	4
Chapter 1: An Overview of Microsoft's Exchange	6
Chapter 2: Exchange Server Roles	23
Chapter 3: Database Management	37
Chapter 4: Recipient Management	51
Chapter 5: Regulatory Compliance	61
Chapter 6: High Availability and Site Resiliency	76
Chapter 7: Unified Messaging	91
Chapter 8: Exchange Virtualization	101
Chapter 9: Exchange Security	110
Chapter 10: Office 365 (Exchange Online)	119
Appendix: Basic Exchange Prerequisite Knowledge	127
Vendor Sponsor: Mimecast's Unified Email Management	140
Index	145

Author Biography

J. Peter Bruzzese has been working in the world of corporate networking since the early 90's as a teenager. He started in document processing departments of Manhattan-based legal and corporate banking firms like Goldman Sachs and Solomon Smith Barney. This was a time of great change in the computing and corporate world as the 'Internet' was just about to become a household name. J. Peter returned to school for networking on a leap of faith (with support from his wife) and he would eventually become one of the most certified IT professionals of his time. He passed certification exams for Microsoft networking, as well as exams through CompTIA, Novell, CIW and other vendors. In the late 90's he established two different networking schools (NetEssentials Training and LAN-Slide Technologies) before being asked to write his first book for Coriolis on Active Directory Design, which he co-wrote with Wayne Dipchan. He would go on to write and contribute to over a dozen titles sold internationally and translated into just as many languages.

More recently J. Peter Bruzzese formed an online company with Tim Duggan called ClipTraining. ClipTraining was formed in 2006 to provide training content through an online learning system that controls access and provides reporting for both individuals and corporate users. The goal is to assist in providing training and support for Windows, Office and other skills through short, task-based videos. J. Peter, as a noted expert on video creation (screencasts) and Microsoft technology has also worked with Pluralsight (formerly TrainSignal) in recent years to provide administrative training for Exchange 2010 and 2013.

As a result of his community effort J. Peter has been awarded the prestigious MVP Award for Microsoft Exchange several years in a row. He is also a Microsoft Certified Trainer (MCT). In addition to books, J. Peter writes for a variety of different in-print and online tech periodicals. He does product review work for TechGenix (MSEExchange.org). He has written the Enterprise Windows column for InfoWorld for 5+ years. And he speaks at technical conferences like TechMentor, MEC 2012/2014, FETC, Connections, TechEd and the Microsoft WPC.

Acknowledgements

As always, my wife Jennette deserves the greatest amount of appreciation. She has supported me from the beginning in this ever changing career.

There are truly too many to acknowledge and thank if I really start thinking about it. For this book I'm going to say focused on acknowledging those who connect with Exchange directly so as to narrow my focus a bit.

First off, I'd like to thank the Exchange Team (past and present). Your work is appreciated on multiple levels. I've always been proud to say I work with Exchange because it is such a stable and feature-rich product that continues to evolve in positive ways.

Others at Microsoft I'd like to thank include my MVP Lead Melissa Travers, Bharat Suneja, Ian Hameroff, Scott Scoll, and Jeff Mealiffe. I'd also like to thank Marissa Salazar, Navin Chand, Brian Shiers, Jake Zborowski, Jon Orton, and last but not least by any means, David Espinoza.

I'd also like to thank those folks at Waggener Edstrom who keep me informed as a journalist. You are much appreciated: Leigh Rosenwald, Kara Berman, Krista Valiante and others I've worked with over the years.

Next, I'd like to thank my fellow MVPs (including those who have moved on to work for Microsoft directly), as well as Exchange Rangers, who I've learned a great deal from. Your articles and insight have greatly improved my knowledge of Exchange.

More specifically I'd like to mention a few MVPs who I have bonded with and appreciate on a personal level. Tony Redmond, Paul Robichaux, Clint Boessen, Jaap Wesselius, Michel de Rooij, Jason Sherry, Jeff Guillet, Jim McBee, Lee Benjamin and Paul Cunningham.

I'd like to also mention Henrik Walther (one of the finest Exchange tech writers I've read and fellow MSeXchange.org contributor).

I'd like to thank InfoWorld and Galen Gruman, Eric Knorr and Ted Samson who I have thoroughly enjoyed working with for the past few years with my Enterprise Windows column. I'd like to thank the folks at TechGenix (aka MSEExchange.org) Sean Buttigieg, Michael Vella and Barbara Matysik-Magro from TechGenix (MSEExchange.org), Jay Gundotra from ENow, Ray Downes, Peter Melerud, Bhargav Shukla (MVP), and Jason Dover from KEMP Technologies, and Peter Bauer, Julian Martin, Steve McKenzie, Janet Vargas and Ani Hagopian from Mimecast. And a special thanks goes to Ed Liberman, David Davis and Scott Skinger.

I'd also like to mention my huge array of friends at Pluralsight including Aaron Skonnard, Chad Utley, Fritz Onion, Gosia Niklinski, Gary Eimerman, Joanna Beer, Lisa Szpunar, Sandy Moran and many, many others.

Thanks to all of you for working with me over the years.

I saved special acknowledgments for the end here. I'd like to thank all of the folks who worked with me on this book. Theresa Miller, you've proven to be a wonderful assistant to me over the time it took to create. And I very much appreciate all my expert reviewers including Janet Vargas, Phoummala Schmitt, Steve Ahlgrim, Lasse Pettersson (MVP) and Paul Robichaux (MVP). Thank you all again for your comments and suggestions to help make this book better.

I call the geek in the pictures J. because while my friends all know me as Peter when I write and speak I use J. Peter. So apparently the J. represents my geek side. He's what I see when I look in the mirror (although I realize I look nothing like him). An alter ego so to speak.

J. P. B.



Greetings! I'm J.

Introduction from the Author

Excerpt from a real conversation with my mother:

Mom: So, what is it you... do?

J. Peter: I assist in the planning and deployment of a messaging solution provided by Microsoft called Exchange (and Exchange Online through Office 365) where I'm responsible for designing enterprise grade messaging environments, at times on a global scale, for organizations that require aspects like regulatory compliance, high availability and Unified Messaging be taken into consideration, along with other unified communication options.

Mom: Ummm... what?

J. Peter: I do email stuff mom.

Mom: Oh... that's nice.

What exactly is Conversational Exchange all about?

It's really an Exchange primer. It's meant to be of help to all those IT admins who are non-Exchange admins looking to either grasp the concepts of Exchange or getting ready to jump into a more professional role with Exchange. It's also meant to assist IT decision makers who may not have time to be hands-on with Exchange but need to grasp all the concepts surrounding it.

It's also meant to assist all of the sales, marketing, and PR folks that work on products that help add value to Exchange (and Exchange Online) to grasp the concept of what Exchange is all about and understand the features and terminology that go along with working in this field.

And for those of you who are already Exchange experts it may help answer the question for your family and friends "what is it you do?" so that you don't have to respond "email stuff" any longer.

I try to make the explanations as easy to swallow as I can. It won't always be easy to grasp every point but don't stress about

that. Learn the concepts and keep moving forward. Like a puzzle, it will start to come together over time. It may be wise to take some time to research some of the terms and concepts.

If you read this book and feel ready to go to the next level, becoming an Administrator with Exchange, you might consider setting up a lab to work with it. Perhaps consider picking up a book (there are several that are really good like Paul Robichaux's "Microsoft Exchange Server 2013 Inside Out: Connectivity, Clients and UM" or one I worked on called "Mastering Exchange 2013").

If you like videos, I've created a ton of video training that you can watch at Pluralsight.com (www.pluralsight.com)

If you want to dive wholeheartedly into the world of Exchange, here are a few resources to get you started:

The Microsoft Exchange Team Blog:
<http://blogs.technet.com/b/exchange/>

TechNet: Exchange Server for IT Pros:
<http://technet.microsoft.com/en-us/exchange/>

Exchange MVP Tony Redmond's Blog: Thoughts of an Idle Mind: <http://thoughtsofanidlemind.com/> Note: Tony also has a blog with WindowsITPro called Exchange Unwashed: <http://windowsitpro.com/blog/tony-redmonds-exchange-unwashed-blog>

I could list out a ton of other blogs I read and various Twitter accounts I follow. Easiest way to learn it all is just to follow me @JPBruzese and I'll point off toward all of them over time. I also have my own personal Exchange blog called ExclusivelyExchange.com that I add information to from time to time.

Chapter 1: An Overview of Microsoft's Exchange



*"Wait... did you just email me to ask how it's going?
Seriously?! I'm like... right here J!"*

Just breathe. This isn't going to hurt. You may be new to Exchange or even servers in general, but one thing is certain: You are not new to email. Email has become the underlying foundation for a new civilization based on global communication. According to recent estimates 2.3+ billion people use email. About 150 billion emails are sent per day. And most of that traffic is coming from the corporate world, according to the Radicati Group.

Most people understand that they can open a browser and go to their favorite browser-based email solution (Gmail, Yahoo, etc...) or open up an email application, like Microsoft Outlook, and send an email to a colleague in the cubicle next to them perhaps,

or on the other side of the globe... and it works. So long as the email address is accurate, they hit 'Send' and it... just... works. Unbelievable really. Especially since the capability to send any form of email is less than 50 years old. It's evolved a great deal over the years.

Three Interesting Facts about Email

The @ Symbol: Ray Tomlinson setup an email system in 1971 using ARPANET. Note: If you don't know what ARPANET is you need to read my book "Conversational Geek (in 7 days!)" Chapter 3: The Internet. Mr. Tomlinson used an @ symbol to distinguish the user from the machine they were working on.

MIME (Multi-Purpose Internet Mail Extensions): Proposed as a standard by Nathaniel Borenstein and Ned Freed in 1991 MIME extended the original capabilities of Internet email so that we could include more than just plain ASCII text. MIME is the "official Internet standard that defines the way that multimedia objects are labelled, compounded, and encoded for transport over the Internet" as explained by Dr. Borenstein. Thanks to MIME we can send all sorts of things in email like images, video, documents and so forth. The email client has the ability to discern how to read the email thanks to MIME headers. Dr. Borenstein can also be credited with sending out the first real MIME message with an "attachment" on March 11th 1992. To see it, go here: <http://www.guppylake.com/~nsb/mime.html>

Protocol Alphabet Soup: An Internet standard that email servers use for sending and receiving email is called SMTP (simple mail transfer protocol). SMTP uses port 25. Client applications, however, use SMTP only to send email. To connect to a mailbox and retrieve email client applications use either POP3 (Post Office Protocol using port 110) or IMAP (Internet Message Access Protocol using port 143). Later on in this chapter we will also discuss MAPI and Exchange Server.

You can Wikipedia the word Email if you are truly interested in this history but to stay on point here with Exchange we'll keep moving forward.

The question though is how does it work? Now toward the back of the book in the Appendix you'll see we have an item that reviews some of the underlying technology for networks and the Internet that will help explain some of the "how" from a deeper technical perspective. But it's obvious that in addition to the technology that allows email to travel from one place to another there must be some way that email is managed and organized, secured from start to finish (hopefully) and waiting for an end-user to pick it up.

Email works, in theory, like the sending of physical mail. If you have an email address it means you have an email mailbox where mail is delivered. When someone sends you a physical letter it must travel by various means (planes, trains, boats, trucks... you get the picture) but the reason it reaches your literal mailbox is not just because there are roads and such to allow transport. There are also centers for organized distribution of that mail. Post offices that help organize that mail. When you go to the post office and mail a letter you may be putting it in the Local box (for people within your town) or the Out of Town Box (for people anywhere else in the world).

Take this same illustration and place it in a large office that has a mailroom. You can send mail to others within your company and it is handled one way, or send it to persons elsewhere and it gets routed out.

In order for email to do the same thing for a company there must be servers that act as the post office or mail room for your organization. When you send an email to a person within your organization that email is sent from your computer so a special server that will handle its transport from that point forward. Now maybe the email is going to a recipient with a mailbox on that server. Maybe it's going to a recipient within your organization that has a mailbox on another email server. Maybe it's going to be sent out to another server in another company altogether where the receiving mailbox is located. The point is when you hit Send that email begins a journey that either completes when it reaches the mailbox of the recipient or, should something go wrong, gets bounced back (like a 'return to sender') with an NDR attached (non-delivery report).

So remember that mail room in your company? It handles the sending and receiving of mail at your site for your organization, but other sites have their own mail room (or multiple mail rooms depending on the company's size). In the email world that mail room is replaced by a special email server. Now there are a variety of different types of email servers in the world, but the one that has grown to become the market leader... the mother of all email servers... is Microsoft Exchange Server.

Microsoft Exchange Server is a messaging platform that provides email, scheduling and tools for collaboration. It's installed as a server-side application. In other words, you have a computer that is designed to be powerful enough to handle the workload that email will put upon it and this is your server. You install a Server OS on it (like Server 2012 R2) and quite honestly it looks and acts on screen like your normal Windows OS (if running Server 2012 it looks like Windows 8 for the most part). And you install Exchange Server on top of your Server OS in much the same way you install Office or Outlook on top of your Windows OS.

Am I oversimplifying the whole thing? Absolutely! So you Exchange experts reading this part and spiking out your blood pressure need to calm down just a bit. Folks reading this don't need to know how complex the process truly is. They don't need to know all the hoops we have to jump through making sure the AD schema is prepared and all the prerequisites are met before we do the install. We're not trying to make them Exchange admins (not yet anyway).

I can hear some of you now. "But... you're leaving out virtualization... you're leaving out the cloud... Office 365!... you're..." Breathe. I'm not leaving anything out. We'll get there.

Ok, so Exchange Server is a server-side messaging application that handles incoming and outgoing email for your organization. One of the things you can do with it is create mailboxes for everyone in your organization. How does it know who everyone is and how do you access these mailboxes? Well, to install Exchange Server you have to have an identity management system that works as a directory service. In the Microsoft world

we call this Active Directory. In order for you to log into your company network you get a username and password that is stored with Active Directory (along with other pertinent information like your address, mobile number, etc...). Exchange uses Active Directory in many ways but one of the key ways is to be able to create mailboxes for your people that connect to the AD network accounts. So when they log into the network and open Outlook up, if they have a mailbox on the Exchange Server they'll be able to send and receive email.

We mentioned earlier that you have POP/IMAP protocols that mail clients use to receive email and SMTP for mail clients to send email. However, Microsoft Outlook, when used within a company, uses MAPI to communicate with Microsoft Exchange. MAPI (Messaging Application Programming Interface) allows client applications to become messaging-aware and uses RPC (remote procedure calls) as its transport mechanism. In 2007 MAPI was also being called the Outlook-Exchange Transport Protocol (which was still just MAPI riding on RPC). With Exchange 2013 a change was made and MAPI connections are no longer supported. Instead RPC over HTTPS (or Outlook Anywhere) connections are supported for both internal and external client connectivity. Here is a great reference from Exchange MVP Tony Redmond entitled "Exchange 2013 focuses on RPC-over-HTTPS"

<http://windowsitpro.com/blog/exchange-2013-focuses-rpc-over-https>

Now before we go any further let's just take a step back and look at the history behind the product we know today as Exchange Server.

<p>Important Note to Reader: We haven't gotten too deep technically yet. But this next part is a bit overwhelming the first time through. Don't let it confuse you and don't give up. Many of the features I rattle off here really fast are covered in greater detail in later chapters in the book.</p>

The History of Exchange



A quick look at the history of Exchange will take you back about 20 years ago (1993) with the planned migration internally at Microsoft from a legacy XENIX-based system to a very early, beta version of Exchange. It wasn't until early 1996 that Exchange Server 4.0 was released to the public, a public already relying heavily on IBM/Lotus which dominated the messaging space. As they say, 'you've come a long way baby!'

Note: Some may wonder why 4.0 was the first version. Exchange MVP Lasse Petterson explains that “prior to Exchange 4.0 Microsoft had a product called Microsoft Mail, released in 1991, and the last version was 3.5” so that explains the 4.0 version number.

Exchange Server 4.0 was X.400 protocol based with support for X.500 directory services. Remember, this was before Active Directory was released (in 2000) so they still needed a directory service and the work they did on Exchange eventually helped with the creation of Active Directory (an LDAP based directory service that succeeds X.500).

With Exchange 4.0/5.0 and 5.5 there was only one mailbox database. Starting with 4.0 the Exchange Team developed single instance storage (SIS), which would provide for an efficient way to reduce disk space by not keeping more than one copy of a message. So if someone sent the same message to multiple people the message body and attachments would only be stored once. Obviously this would keep the database size down (which was smart because disk space wasn't cheap in those days).

The database management solution for Exchange is called the Extensible Storage Engine or ESE (aka JET Blue with JET standing for Joint Engine Technology) is a transaction-based database engine. It's been likened to a distant cousin of the database used in early versions of Access (JET Red) but it's not an Access database nor is it a SQL database. It's been specifically optimized over the years to store hierarchical data (ie. folders, messages, attachments) and to survive crashes so that upon recovery the data loss is minimal.

Ok... so Exchange 4.0 was out the door. With no time to spare they went right to work on version 4.1, which turned into version 4.5 and then ultimately became the 5.0 release (shipped in early 1997). There were some great new aspects to 5.0 like the implementation of SMTP and LDAP v2, as well as the first version of the web-based client (Exchange Web Access, or EWA) which we know today as Outlook Web App. There was a 16 GB limit on the database size.

On Exchange 5.0 CDs (if anyone has one) there was a file called EXGL32.DLL and if you renamed it to .AVI it was a video Easter egg that would credit the Exchange 5.0 team while having fun in the process.

Exchange 5.5, shipped at the end of 1997, continued to work off a separate directory service. It was sold in two editions: Standard (which maintained the 16 GB limit of the database size) and Enterprise (which allowed up to 16 TB). There were some additional feature differences between the two editions including transport connectors and clustering options.



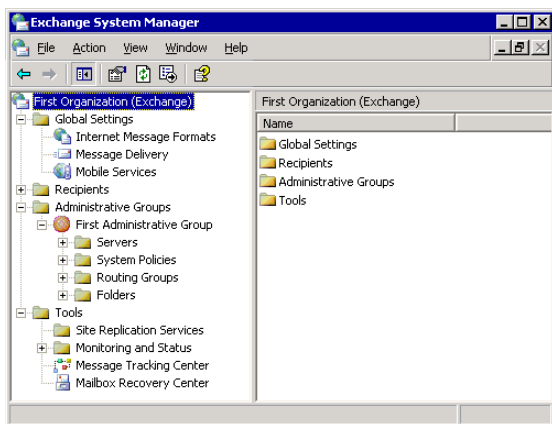
Incidentally, it was right around this time that I became proficient with Exchange and passed my certification for it in January of 1999 allowing me to become MCSE+I certified. The +I didn't impress anyone in 1999 and it impresses even fewer today.

A term you may be familiar with is “email storm” which occurs

when there is such a huge load of email traffic generated that the email servers go down (like a DDoS attack). This can occur when there is a spike in Reply All and distribution lists. One such email storm that is well known at Microsoft occurred on October 14, 1997 with Exchange 5.5. It involved a distribution list called Bedlam DL3, which had about 13,000 email addresses in it. One Microsoft employee asked to be removed from the list (to all) and others responded with “Me too!” and supposedly 15 million emails were sent in the process, causing a crash. Read all about it here on the Exchange Team Blog:

<http://blogs.technet.com/b/exchange/archive/2004/04/08/109626.aspx>

Exchange 2000 (v6.0) was released in November 2000 and this was the first version that dropped a separate directory service and now relied upon Active Directory. One of the features I liked with 2000 was an Instant Messenger feature but this didn't remain in the product for long (it was moved over to Office Live Communications Server and yanked from Exchange 2003). One pain about this release was the migration from Exchange 5.5 where you had this Active Directory Connector (ADC) which was anything but fun (although it did work). Exchange 2000 had us focused on the ability to create multiple storage groups where we could put multiple databases. We'll discuss the evolution of the Exchange database and architecture in Chapter 3.



The Exchange System Manager in Exchange 2003

Exchange 2003 (v6.5) was released in September of 2003 and included features like RPC/HTTPS (now known as Outlook Anywhere), cached mode and ActiveSync (a key piece for mobile client connectivity to Exchange). Spam was becoming a nightmare for admins and Exchange 2003 added some basic filtering features like connection filtering, recipient filtering and Sender ID filtering.



To combat spam Exchange Admins often had to look at 3rd party options such as Sybari Antigen for Exchange. Today you may know this product as Microsoft's ForeFront Protection for Exchange Server.

To see a list of build numbers and release dates for Exchange Server from 4.0 to 2007 SP3 go here:

<http://support.microsoft.com/kb/158530/en-us>

From Exchange 2007 forward...the start of something new

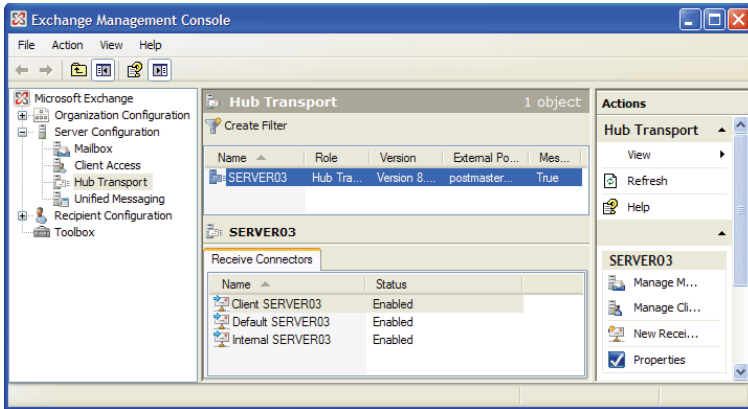


The Exchange Team was working on some awesome new features with the next version of Exchange (v7.0). The v7.0 had a lot of “proof of concept code” that didn’t get a release however. Instead they took things to the next level and went in a new direction with v8.0 (released as Exchange 2007).

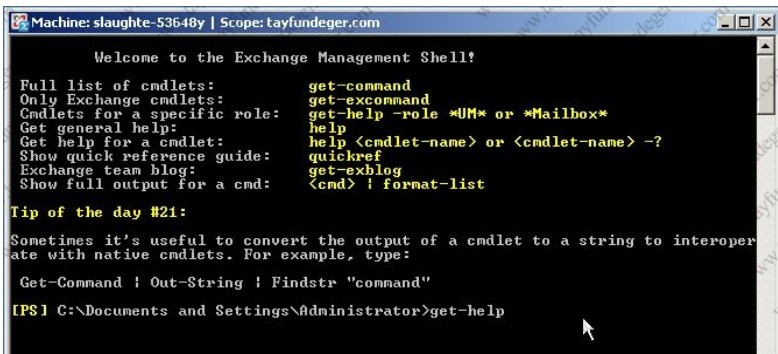
Exchange 2007 is actually v8.0 but because the Office team was releasing their 12 wave (and to sync with them) they call it E12 now. That may sound confusing but the first version was 4.0 so why should it matter to us that they jumped from 8 to 12?

Exchange 2007 was shipped in December 2006. This version of Exchange had a variety of new features that have carried through to the Exchange 2013 version (although some have morphed a bit into better features). Some of the new features included:

- **Server Roles:** Exchange 2007 introduced 5 server roles. There were 4 internal roles: Mailbox, Hub Transport, Client Access and Unified Messaging. And 1 external, perimeter-based role called the Edge Transport. We'll discuss these further in Chapter 2.
- **Continuous Replication:** This feature allowed a copy of the active database to be copied and then transaction logs to be shipped to provide different levels of availability. There were 2 initial CR types and a third added with SP1 including: Local Continuous Replication (LCR), Cluster Continuous Replication (which had clustering features for automatic failover support), and Standby Continuous Replication. These will be discussed further in Chapter 6. Note: there was also a legacy clustering option called Single Copy Clusters (SCC)
- **Unified Messaging:** A way to have voicemail go into your Inbox by connecting to your existing PBX/IP-PBX system. Also a set of auto attendant features built right into Exchange. We'll discuss this further in Chapter 7.
- **Exchange Management Shell (EMS):** A new command-line/scripting language based on PowerShell was introduced and in some cases you could only perform things through the EMS. The Exchange Management Console (EMC) provided a GUI based administration method as well. Fun Fact: PowerShell was originally code named Monad prior to official release.



The Exchange Management Console



The Exchange Management Shell (EMS)

With Exchange 2007 a line was drawn in the sand that required x64 hardware running 64-bit versions of Server. This would require companies to purchase new hardware if they didn't have an x64 server to use. Another interesting change was made with Exchange 2007 upgrades and that was the inability to perform an in-place upgrade from a legacy version of Exchange (2003 and lower) directly to 2007. This remains the case with Exchange 2010 and 2013. You have to install the new Exchange server into an existing environment, coexist for a period of time, and transition over to the new when ready.

Although some administrator's might balk at the inability to do an in-place upgrade, environment coexistence is often a welcomed change for administrators because it simplifies

migrations and gives us an excuse to update our systems to ensure performance for the newer implementation of Exchange.

The term migration is typically used when referring to a move from one messaging system or a legacy Exchange system to a more modern version of Exchange that cannot coexist with the other system. For example, if you are trying to move from Lotus Notes to Exchange that is considered a migration. However, coexistence with a transition is an upgrade that takes you from a legacy version of Exchange to a modern version so long as the two can coexist. For example, Exchange 2007/2010 to Exchange 2013 would involve a period of coexistence (where you have both server types in your organization at the same time), you move mailboxes from legacy to modern at the needed pace (smaller organizations might do it in a weekend, others may take months), and then when all mailboxes have been transitioned over you decommission your legacy Exchange servers and have now “upgraded” so to speak. But again, there is no in-place upgrade.

The 2 editions of Exchange still existed with 2007. In the Standard you could have 5 databases in up to 5 storage groups. The Enterprise edition supported up to 50 databases and up to 50 storage groups.

Exchange 2010

Exchange 2010 (v14) was released to manufacturing (RTM'd) in May of 2009 and released officially in November of 2009. Here were some of the major features that we still have with Exchange 2013:

Database Availability Groups (DAG): Building off the continuous replication options with 2007 the 3 CR options were boiled down into 1 solution. Note: We use DAG in Exchange 2013 as well, although it's constantly being improved upon by the Exchange Team.

Personal Archive: Hoping to help eliminate the proliferation of PSTs and as a result of storage becoming cheaper a personal archive feature was added so that admins could keep the Inbox

on higher performance disk and an archive on lower performing storage (if necessary).

Storage Groups are dropped, as is Single Instance Storage (SIS) which improves performance greatly but creates the possibility of storage bloat. However, by this point storage is very inexpensive and you can design for better efficiency to avoid the lack of SIS from having that great an impact.

For a list of Exchange Server 2010 (back to 4.0) server build numbers and release dates: [http://technet.microsoft.com/en-us/library/hh135098\(v=exchg.141\).aspx](http://technet.microsoft.com/en-us/library/hh135098(v=exchg.141).aspx)

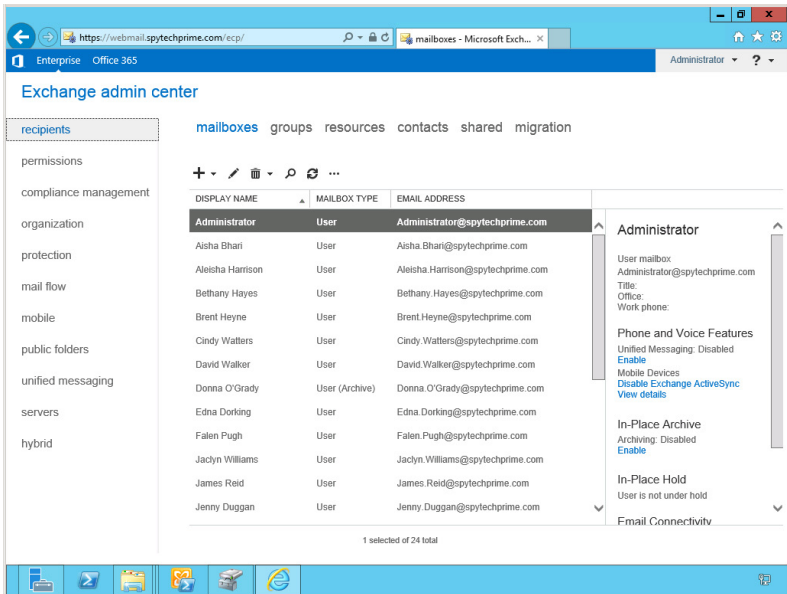
Exchange 2013

Released to manufacturing in December of 2012, Exchange 2013 (v15) is the latest version of Exchange Server. Exchange 2013 has a handful of important “what’s new” features including the following:

- Exchange Administration Center: The EAC brings to mind the quote, "You're riding it, dude!" Yes, I cannot help but think of the line from "Finding Nemo" whenever I think of the new admin console for Exchange. But seriously, the GUI-based EMC (Exchange Management Console) and the Web-based ECP (Exchange Control Panel, released with Exchange 2010) has been replaced by a single Web-based UI. Ordinarily I don't like Web-based consoles for administration; they always feel clunky and unfriendly. Plus, it has that Metro look, which leaves me cold. But to be honest I've come to really appreciate the Exchange Admin Center because of its ease of use and the fact that I can access it easily from a browser.
- Exchange architecture revisions: Exchange 2007 and 2010 are broken into five server roles, mainly to address performance issues like CPU performance, which would suffer if Exchange were running as one monolithic application. But Microsoft has made progress on the performance side, so Exchange 2013 has just two roles:

Client Access server role and Mailbox server role. The Mailbox server role includes all the typical server components (including unified messaging), and the Client Access server role handles all the authentication, redirection, and proxy services.

- A new managed store: The store service has been completely rewritten in managed code (C#). Exchange 2013 continues to use ESE as the database engine. But now each database runs in its own dedicated worker process, so a hung process in one database will not cause problems in other databases. Fast Search (an add-on to SharePoint 2010) is also integrated into the managed store for improved search and indexing.
- Modern public folders: In previous versions of Exchange you had to have a public folder database for public folders, but now you can create a public folder mailboxes, which means they use regular mailbox databases. In turn, this means they can be made part of a database availability group for disaster recovery.
- DLP (data loss prevention): DLP is new in Exchange 2013's transport rules that warn or prevents users when they may be violating policies meant to prevent disclosure of sensitive data like a credit card number or Social Security number in an email. The built-in DLP policies are based on regulatory standards.
- Outlook Web App enhancements: One awesome feature is support for offline access, which lets users write messages in their browser when offline, and have the messages delivered when they connect to the Internet.
- Built-in anti-malware: Exchange has had anti-spam capabilities for quite some time; as of Exchange 2007 you could even choose whether to turn on anti-spam in the Edge role or in the Hub Transport role. Exchange 2013 extends anti-spam to a broader set of antimalware capabilities but it's a still a first attempt at this.



The Exchange Admin Center in 2013

For Exchange 2013 releases (including cumulative update releases) go here: <http://technet.microsoft.com/en-us/library/hh135098.aspx>

Hosted Exchange, Exchange Online and Office 365

Terms you will hear in modern Exchange deployments include on-premise (where you install your Exchange Server in your own environment), hosted Exchange (where a service provider manages your Exchange server), cloud-based (where you go with an Office 365/Exchange Online solution) and Hybrid (where you combine both).

Note the options shown in the free, online tool called the Deployment Assistant: <http://technet.microsoft.com/en-US/exdeploy2013/>

The Exchange Server Deployment Assistant is the IT pro's source for Exchange deployment technical guidance. Tell us what kind of deployment you're interested in, answer a few questions about your environment, and then view Exchange deployment instructions created just for you.



The Exchange Deployment Assistant

Hosted Exchange comes in many forms. In some cases the provider will place your organization's email on the same server as other companies and this is called a multi-tenant scenario. You will not be aware of the other companies using the same server however but you share the power of that server and typically are provided minimalist tools to access and manage your end-users. In other cases you might have a virtual or dedicated server with a full Exchange deployment that you can have complete control over.

Exchange Online (a part of Office 365) offers Exchange as a cloud service with Microsoft as the provider. You essentially choose from a variety of subscription options (with different features and prices) and manage your Exchange through the O365 admin tools and the Exchange Admin Center (EAC).

Exchange Online has been designed to make it easier for companies who are not ready to jump all in with a cloud-based email solution to tie their on-premise Exchange with their cloud-based Exchange online and form a Hybrid solution. It sounds easier in theory than it is in real life so when an Exchange admin tells you they are working on a hybrid configuration it's ok to raise your eyebrows and say "Impressive... how's that going?"

The Big Takeaways


This chapter has way too much to digest in one sitting. But it's meant to be that way. You were hit with a great deal of information all at once. The roller coaster of development over

20 years of this massive solution with so many features that one chapter couldn't present it all. But here we are... finished with Chapter/Day 1.

The big takeaway? Exchange is ever evolving. Knowing a bit about that process will help you learn about the here and now. We're going to use this same approach with other subjects like Server Roles and High Availability. Walk you through the history so that you can learn the process from its inception and build on it from that point.

The goal here was to introduce you to a ton of new vocabulary all at once. In the chapters ahead we will take pieces of this discussion and break it down into more digestible bites. So don't lose steam now... push on to Chapter 2!

Chapter 2: Exchange Server Roles



*“J!!! You have to come back with me to 2003!
Server roles are reverting back to 2003!”*


If you have a small environment you can install Exchange Server on one Windows Server and it will handle all the work. But as your organization grows it may be wiser to install additional Exchange Servers to handle more mailboxes. And wouldn't it be nice if you could install Exchange Server in pieces so that you can really optimize performance on the servers and distribute server tasks? Well... that line of thinking is what led to the ability to install Exchange as roles, not just one monolithic solution.

In this chapter the evolution of these roles will be explained so you can know what is available with each flavor of Exchange you may be dealing with. Let's walk through in stages starting with Exchange 2003.

Exchange 2003 “Roles”

I can just hear some of the Exchange vets now saying “ahem... pre-Exchange 2007 didn’t have roles”. Not in the modern sense but The Exchange Team Blog said “Exchange 2003 provided primitive server roles called BackEnd server and FrontEnd server”. <http://tinyurl.com/okbsgyw>

But let’s not fight over semantics. The fact is with Exchange 2003 we were able to configure the front-end Exchange server to handle client requests and then proxy them back to the appropriate back-end Exchange server where the mailboxes resided.



Client Connectivity to Mailboxes with Exchange 2003

Note: Prior to Exchange 2003 the savvy administrator found ways to create their own roles and ensure system performance. For example, a mailbox server without databases and the use of DNS could be turned into what we might call an Edge or Hub Transport server today.

With front-end servers the internal clients connected to their mailbox using Outlook and MAPI but they connected directly to the mailbox servers (the back-end servers). The external clients used the front-end as more of a proxy that could handle RPC over HTTPS (or Outlook Anywhere), Outlook Web Access (OWA), POP/IMAP or ActiveSync connections.

What does it mean to proxy? When an end-user goes to access their mailbox the front-end server contacts Active Directory (specifically a global catalog server in the domain) and locates which back-end server contains the user’s mailbox. That action

of moving requests from the front to the back is what is meant by the word proxy. Now that is the general functionality of the front-end server but the exact functionality varies depending on the protocol used and the action being performed. No point in getting that deep into it, but good for you to know because the word 'proxy' is going to come up again in this discussion.

In addition to front-end/back-end for client access and mailboxes there was another concern with Exchange 2003 and that was the actual transport of messages. Within an Exchange organization (depending on its size and number of locations) you would have servers organized by routing groups. These groups would have connectors between them. Bridgehead servers would handle message transfer from one routing group to another (or to an external messaging system).

Ok, so now let's take things to the next level.

Exchange 2007 Server Roles

Like I said at the outset, if you have a small or medium sized company with a few hundred mailboxes you can install all your Exchange required roles for 2007 on a single physical server. But with the larger enterprises ranging into thousands or tens of thousands of mailboxes, multiple office locations there was a need to provide a more flexible deployment approach.

In addition, scaling up (which is necessary when your organization grows) was not easy to do with hardware that was not as powerful as what we have today. So scalability and flexibility, although with performance were all drivers for server roles to be built in to the deployment options.

The result was a breakdown of server roles into 5 roles. 3 of these roles (the Mailbox, Client Access and Hub Transport) are required and the other two (Edge Transport and Unified Messaging) are optional. In smaller environments you can install all 3 required roles on one box. Or you can install them on separate servers.

In addition to flexible deployment options this also allows you to improve hardware utilization because the binaries installed are only what you have chosen. In other words, you don't install the whole huge solution on one server, just the bits you need. And only the services for those options will run. This makes the servers easier to configure, secure, maintain and size for hardware.

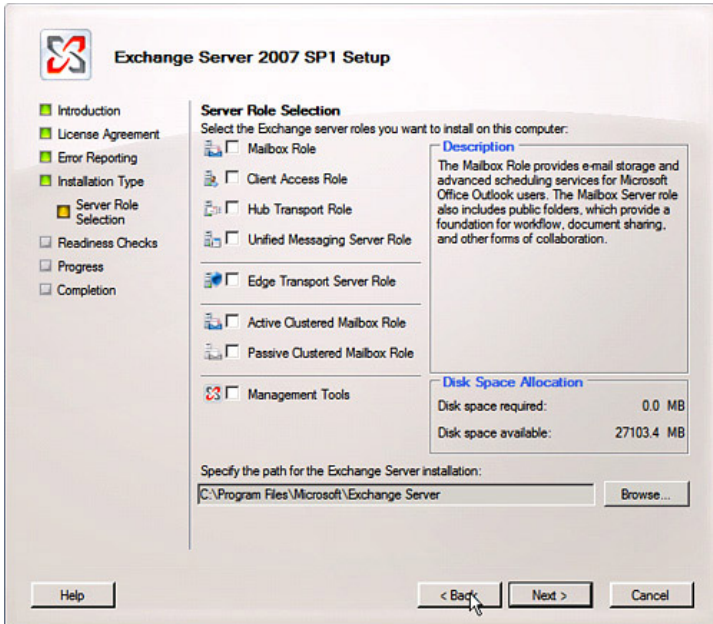
Trivia: During the beta for Exchange 2007 there was a 6th role planned. It's true! It was the Public Folder role (which was rolled into the Mailbox role). And the Hub Transport role was originally called the Bridgehead role because its function was similar to the bridgehead server functionality with routing groups in Exchange 2003.

A Closer Look at Server Roles

Let's start with the 3 required roles for an Exchange 2007 installation and then address the 4th internal role and close it out with a discussion of the Edge role that resides in the perimeter.

The Mailbox server role as its name implies hosts the mailbox databases as well as any public folder databases, while also providing MAPI access to Outlook clients. The Mailbox Role is ordinarily installed with other roles on a single server, such as the Hub Transport Role, Client Access Server Role and the Unified Messaging Server Role, as you can see in the graphic.

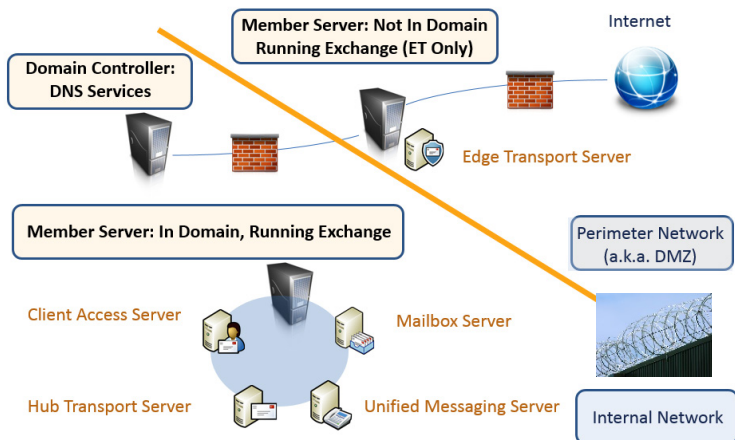
The Client Access server role is similar to the front-end servers in Exchange 2000 or 2003. The Client Access server supports the use of Outlook Web Access for access using a web browser as well as Exchange ActiveSync for mobile devices. POP and IMAP are also implemented here, while Outlook Anywhere support allows Outlook clients to connect from outside the corporate network without the use of a VPN.



Server Role Installation Option for Exchange 2007

The Hub Transport Server Role will route messages within an Exchange organization and is similar to the Bridgehead server found in Exchange 2000/2003. It can also be configured to route external email in lieu of the optional Edge Transport Server Role. The Hub Transport Server is reliant on the presence of Active Directory to have a logical infrastructure in place for the routing of internal messages.

The Unified Messaging server role provides voice over IP capabilities to an Exchange Server in order to integrate e-mail, voicemail and incoming faxes as part of providing a ‘universal’ Inbox. Outlook Voice Access (OVA) also opens the door to multiple access interfaces such as the phone. Given the potential complexity of telephony infrastructure such as IP-PBX and VoIP gateways, a telephony expert is suggested for the installation and configuration of the Unified Messaging server role.



Server Roles in Action in Exchange 2007 and 2010

The optional Edge Transport server role is meant to be the last hop for mail going out of your organization and the first hop for mail coming in. It acts like a smart host and sits on the perimeter and is not part of Active Directory. In addition, an Exchange Server configured as an Edge Transport server role cannot also be configured as any other role. Its main task is to sit on the perimeter of the network to provide security in the form of anti-spam/malware filtering agents and the implementation of organizational transport rules for an organization.

Terminology Note: A smart host is an email message transfer agent (or MTA) that allows an SMTP server to route email to an intermediate mail server instead of directly to a recipient's server.

Terminology Note: The perimeter network (or DMZ, demilitarized zone) is not required for your network or Exchange. Some like to have multiple firewalls with servers in-between that handle anti-spam, anti-virus and other protective pieces to ensure greater security. The Edge Transport role sits in that perimeter however if you didn't want to use one but wanted the anti-spam/malware capabilities you can enable these agents on the Hub Transport server.

Exchange 2010 Server Roles

For the most part the server roles in Exchange 2010 remain exactly the same as 2007. There were some improvements made however.

Let's focus our attention primarily on the Client Access (CAS) role. Remember with 2003 we said the front-end was primarily a proxy back to the mailbox servers? Well, that meant the mailbox servers were still doing a lot of the work. With the CAS role in Exchange 2007 the CAS really helped to offload a lot of the load from the mailbox server although internal clients still connected directly to the Mailbox role. With Exchange 2010 that changes thanks to a new service called the RPC Client Access service (MSEExchangeRPC) making the CAS a true middle-tier solution.

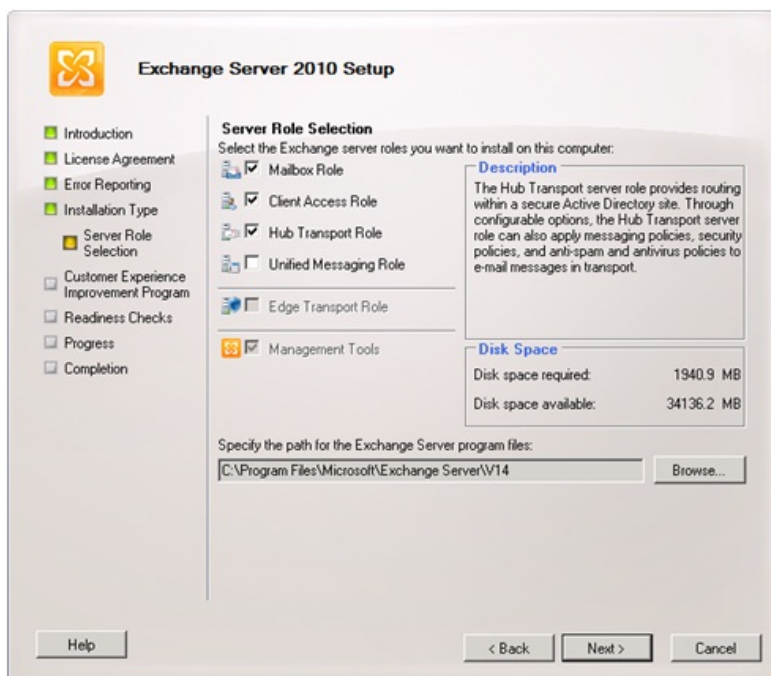
In Exchange 2010 the CAS role handles both external and internal connections to the Mailbox role (with the exception of Public Folder connections). This takes a ton of pressure off the Mailbox server and allows it to handle more concurrent connections.

The CAS role handles the following connections and services:

- Outlook Web App: Allows you to access email through a Web browser (including IE, Firefox, Safari and Chrome)
- Exchange ActiveSync: Allows you to synchronize your data between your mobile device or smart phone and Exchange
 - Note: There are varying levels of ActiveSync support in devices and one key security element is remote wipe, which is not available for all devices
- Outlook Anywhere: Allows you to connect to your Exchange mailbox externally using Outlook (RPC over HTTP) without going through a VPN connection
- POP/IMAP: Mail clients other than Outlook that connect with POP or IMAP are supported through the CAS role

- Availability Service: Shows free/busy data to Outlook 2007/2010 users
- Autodiscover Service: Helps Outlook clients and some mobile phones to automatically receive profile settings and locate Exchange services

We'll discuss high availability in another chapter later on, but it's good to note that with the CAS role being so important you want to ensure you don't lose it. If your CAS server goes down, email goes down. So you want to have more than one (just in case) and you can tie them together as a CAS array within a single site. Your CAS array should be load-balanced with either Microsoft software load-balancing (aka NLB) or a 3rd party appliance based load balancer.



Server Role Installation Option for Exchange 2010

There were other important changes in Exchange 2010 especially with regard to the changes in high availability for the Mailbox server role, but we'll save those for later on.

As you can see in the screenshot, the role install hasn't changed all that much between 2007 and 2010 with the exception of the missing active/passive mailbox roles.

With Exchange 2010 SP1 however, the Exchange Team started encouraging Typical installations of Exchange, which included the three primary roles (MB, CAS and HT) on all internal servers rather than breaking up the roles.

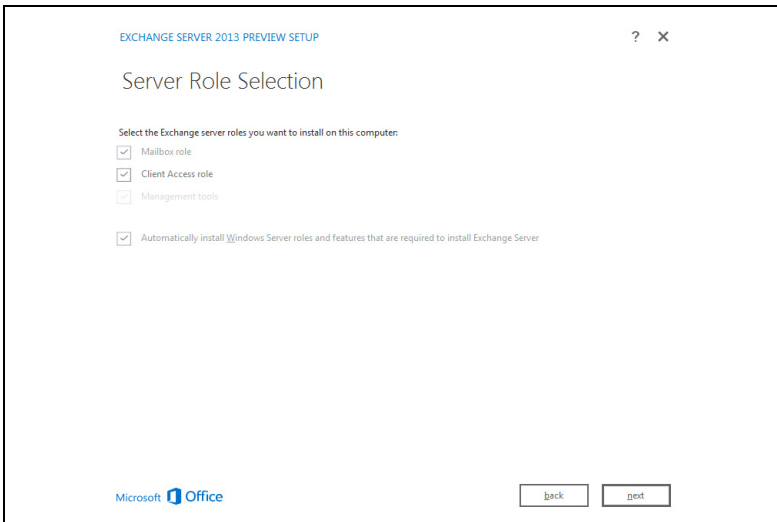
Exchange 2013 Server Roles

Ok... get ready... were going to blow your mind now (and explain the cartoon at the start of this chapter). One thing you should know about Exchange at this point is that it is ever evolving (remember... key takeaway from Chapter 1). And as hardware evolves and performance improves it was decided there wasn't a need for 5 server roles any longer. So, with Exchange 2013 things went from 4 internal roles to 2. And the Edge didn't get an update in 2013 at all.

So what we have is the Mailbox role that handles the primary part of the Hub services, the Mailbox database and the UM role as well. And then we have the Client Access that handles authentication, redirection and proxy services with support for all the typical access protocols: HTTP, POP, IMAP and SMTP.

In addition, these changes bring with it a change to client connectivity. RPC is no longer supported as a direct access protocol. All Outlook connectivity will be done through RPC over HTTPS (aka Outlook Anywhere). Immediately the upside to this is fewer namespaces needed for the connectivity. But this change combined with other adjustments in how clients connect will hopefully eliminate client issues (like the need to restart Outlook at times). Keep in mind that only Outlook 2007 and later is supported with Exchange 2013.

The kicker here is that these two roles closely resemble their 2003 counterparts in many ways. The CAS in 2013 (like the front-end in 2003) proxies/redirects connections back to the Mailbox server. The Mailbox server has the mailbox databases (and all the mailboxes logically) which also holds the public folders.



Server Role Installation Option for Exchange 2013

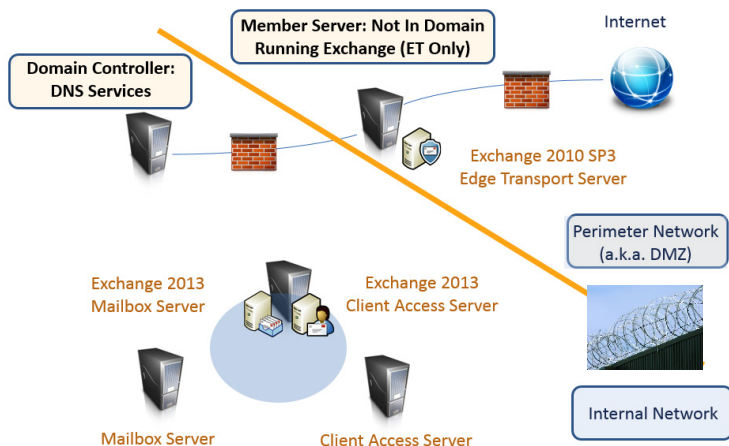
Trivia: Some refer to the CAS in Exchange 2013 as the Client Access Front-End (or “café”) server role. Personally I think it would have been better to give the role the new name to help folks grasp its new (ahem... old) purpose.

What happened to the Unified Messaging server role? Well, it isn't gone, it's now installed with the Mailbox role, so it's completely wrapped into that role now.

What about the Hub Transport server role? Don't we NEED transport? Yep, and so those features have been split up between the CAS and Mailbox (with the Mailbox getting the majority in the split).

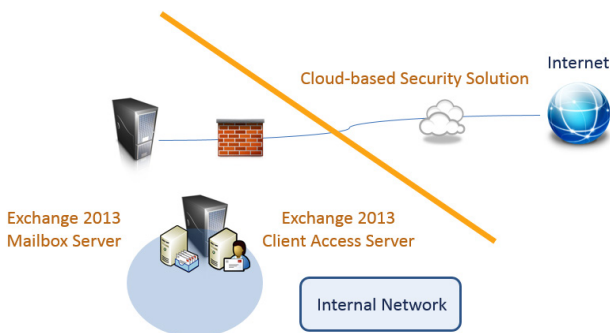
And the Edge again? It wasn't updated in Exchange 2013. In other words you can still use the 2010 SP3 version if you like, or

go with another option for your anti-spam solution (perhaps an enterprise grade option in the cloud like Mimecast) but the Edge was not a major focus in the development of Exchange 2013.



Server Roles in Action in Exchange 2013

There is talk of an update with SP1 but one MVP (who shall remain nameless) used the word “lame” to describe it. Maybe the Edge wasn’t deployed enough to be worth development time. Perhaps Microsoft has other plans in mind. They recently dropped the majority of their Forefront product line, including Threat Management Gateway (TMG) which was oftentimes paired up with the Edge to provide anti-virus/malware and greater security.



Alternative to Edge Transport Role with Exchange 2013

The Exchange 2013 Transport Pipeline and Mailflow

The Microsoft Exchange Team says the mail flow process occurs through the “transport pipeline” which is made up of three services. These services aid in transport on our Client Access and Mailbox servers (which may exist on the same server). The Client Access has the Front End Transport service while the Mailbox server has the Hub Transport service and Mailbox Transport service (which is made up of two services).

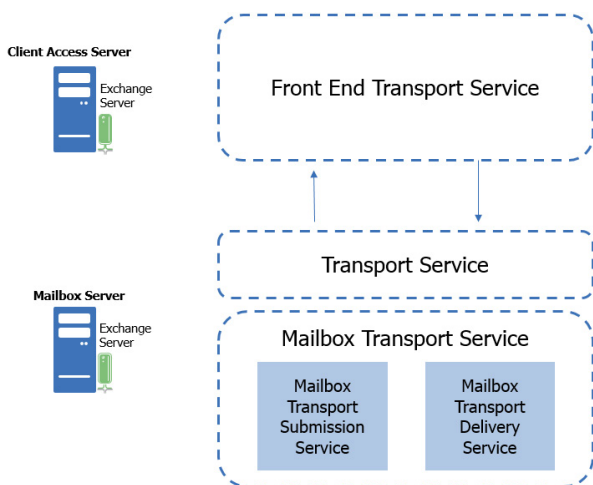
The Front End Transport service on the Client Access server handles the flow of mail from the Mailbox server (specifically the Hub Transport service on the Mailbox server side) to the outside world. The Hub Transport service (or just Transport service) handles routing from the Front End Transport service to the Mailbox Transport service as well as between other servers within the organization internally. The Mailbox Transport service handles mail transport between the Hub Transport service and the mailbox database.

Going back to the Front End Transport service it’s basically a stateless proxy for inbound and outbound traffic with no traffic being queued as a result of that service, however, as mentioned, it can be used to filter traffic. That filtering can be based upon connections, domains, senders and recipients. It does inspect message content however. Inspection of the content itself can be done by the Transport service as it handle SMTP mail flow from the Front End Transport service to the Mailbox Transport service and into the database.

Note: One important point regarding the Front End Transport service is that mail inbound and outbound to the Internet through an Edge will bypass this service. The Edge communicates directly with the Transport service on the Mailbox server.

The two services that make up the Mailbox Transport service include the Mailbox Transport Submission service and the Mailbox Transport Delivery service. The Delivery side accepts messages from the Transport service and delivers them using RPC to the mailbox database. And the Submission service

receives, through RPC, messages from the local mailbox database and passes it to the Transport service.



The Exchange 2013 Transport Pipeline

I promised when I started writing this book that I wouldn't over-involve the reader with too much deep technical content. And even though it's going to look like I'm breaking that promise, believe me... this is just the tip of the discussion on the transport pipeline and how all of this works. I just felt it was essential for you to see where the Hub Transport role really went.

Deeper Dive: If you feel like seeing the full graphic of mail flow and the transport pipeline you can go here and scroll down: <http://tinyurl.com/cw4976k>


The Big Takeaways

Ok... keep breathing, you're doing fine. The sum up from this chapter:

- In 2000/2003 we had “primitive” roles in that we could configure a front-end server for client access, a back-end server for mailboxes and a bridgehead to assist with transport.
- With 2007 and 2010 that changed into 5 official roles. 4 were internal (Mailbox, Client Access, Hub Transport and Unified Messaging) and one was external (perimeter/DMZ) called the Edge Transport.
- In 2007 the CAS received more responsibility for client connectivity and lightened the load on the Mailbox. In 2010 it received all internal/external client connectivity load and that really lightened the load on the Mailbox server.
- In 2013 we go back to the style of roles from the days of 2003. We have a Client Access (front-end, or CAFÉ) and a Mailbox role. The Edge is ignored (so far) and transport services are broken up between the two remaining roles. UM is now automatically installed with the MB role.

There it is... you know everything you need to know to have a conversation about Exchange Server Roles. Next up... a breakdown of the database, one of my favorite subjects. So don't lose steam now... push on to Chapter 3!

Chapter 3: Database Management



“It’s not that difficult... Exchange uses fault-tolerant, transaction-based databases to store messages... and... you look like you don’t feel so good, are you ok?”

At its core... Exchange is a database management server. It has to keep track of all the mailboxes you configure for people in your organization. It puts email in those mailboxes and retains all that information in an organized way. The Exchange Server is all about sending and receiving email, and storing it for end-users to read. In addition to email, calendaring items, contacts, tasks and so forth all need to be housed somewhere and in the end it all goes into the Exchange database (and that database needs an engine to help manage it).

It’s important for you to visually realize that every email that goes to your Mailbox server must go into a database, and this creates challenges because of the huge variety of messages Exchange handles. From the tiny one-line emails to the monster MB emails with video attachments.

The I/O (input/output) profile of a Mailbox server is not predictable it's very random. Read/write that occurs between memory and disk is substantial. At times there may be waves of messages, other times may be idle. So it's essential to understand how the databases work for us to understand how Exchange, at the core, works.

Dr. Carl Sagan said "You have to know the past to understand the present". We've seen that in previous chapters and this one is going to be a review of the past as well to help build your understand up to modern times.

Ok, so let's get started...!

Legacy Exchange Databases

You might recall we talked about the Extensible Storage Engine (ESE) in Chapter 1. This engine has morphed over the years and been improved upon. One benefit of ESE was Atomic, Consistent, Isolated, Durable (ACID) transactions to make for more reliable and recoverable data management.

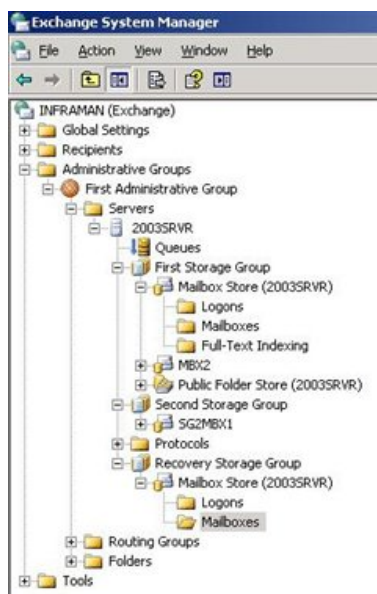
With early flavors of Exchange (4.0, 5.0 and 5.5) there were three databases:

- Dir.edb: Held the directory information for all mailboxes
- Pub.edb: The Public Folder database
- Priv.edb: The mailboxes (all of them) were in this single file database

The maximum database file size was 16 GB so that was a bit of a limitation on how much your Exchange Server could handle.

And then with Exchange 2000 (and later 2003) things changed a bit. With Active Directory now handling the directory information (ie. usernames and passwords for mailbox users) we didn't need a dir.edb database anymore.

Another change was the ability to create multiple databases (called mailbox stores) and store them in containers called 'storage groups'. At that time we also saw the introduction of .stm files, which were paired up with database (.edb) files to provide content in native MIME format. If you recall what MIME does from Chapter 1 the .stm files were simply making things faster for content conversion requirements.



Storage Groups in Exchange 2003

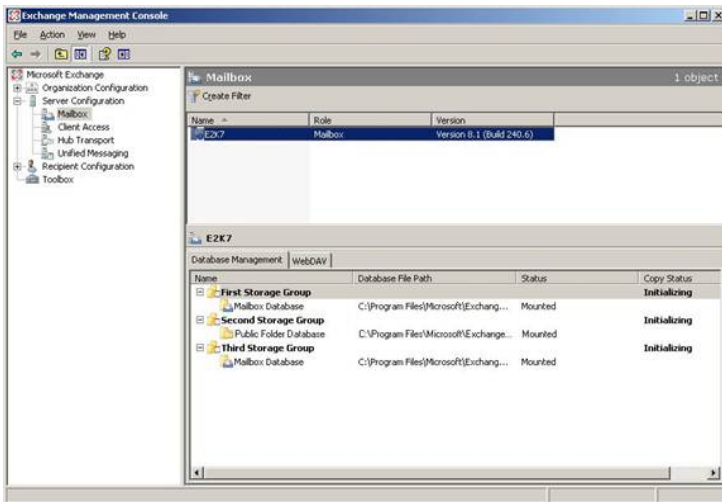
Note: Storage groups may sound like a good idea because you can put multiple databases in them, which is great compared to only having one. But what was odd is that the transaction logs for the databases were merged together in that one storage group. We haven't discussed transaction logs just yet but you'll see soon enough why that might lead to frustration.



Back in the day, the max size of 16 GB left our users with mailbox sizes as small as 50-100MB in size and heavily reliant on .pst files (another can of worms discussed in chapter 5). Depending on the number of users in your organization you could end up with many mailbox servers to administer and maintain. More databases per server was a welcomed change. We could use fewer servers while still maintaining smaller database sizes which was important for achieving respectable database backup and recovery times.

Modern Exchange Databases (2007/10/13)

Exchange 2007 actually yanked the .stm file from Exchange and took us back to the single .edb file for storing content. The storage group maximum number of databases was increased to 50 but overall it was clearer with 2007 that the focus was shifting to the database and away from the storage group. The direction was to have one database per storage group, so that eliminated the need for storage groups.



Storage Groups in Exchange 2007

And as a result, with Exchange 2010 we no longer had to worry about storage groups. And the ESE database engine received enhancements that improved I/O by 70% (meaning, Exchange 2010 can read/write emails to disk 70% faster than 2007 using the same engine). These improvements included increasing the page size from 8KB to 32KB, storing header data in a single database (DB) table, and compressing attachments. In turn, because of these optimizations, you actually have more options for using cheaper disks for your Exchange server.

Note: With older versions of Exchange due to the high I/O requirements many organizations used dedicated SAN to ensure excellent performance for their Exchange environments. With the I/O changes in Exchange 2007 and higher we are able to consider options such as Direct Attached Storage (DAS) and SATA disk which helps the corporate pocketbook.

What do we mean by cheaper disks? Well, think about all the work your Mailbox server has to perform. People in your company are constantly tapping it for their email. To get the best performance out of it in the past you would want to make sure you had the most expensive, highest performing disks. These were usually SAN or RAID arrays. Don't stress too much about the words. Type SAN or RAID array in an image search in Google and you'll see what we mean. They were high-performance, but they were budget busters.

With the enhancements in Exchange 2010 and lower I/O this allows for lower cost SATA disks or JBOD storage (Just a Bunch of Disks) which has had slow adoption but is a great option. And another fact is that with Exchange 2010 you can mount up to 100 databases.

Note: One important point to keep in mind is that Microsoft removes Single Instance Storage with Exchange 2010. The idea behind SIS is when a message is sent to a bunch of people (perhaps with a large file included) the original message is stored once. SIS is replaced by database compression technology and new tools to help administrators to purge mailboxes and reduce the overall size of the database. But dropping SIS supposedly helped improve performance a great deal.

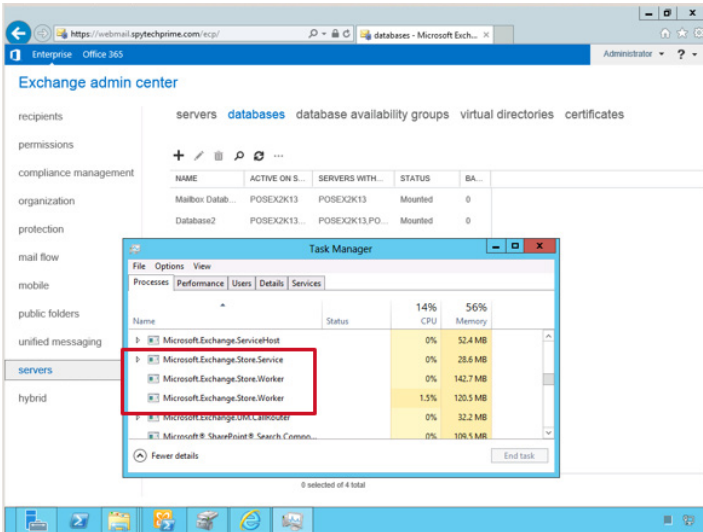
With Exchange 2013 additional enhancements have been made to improve the performance of the ESE database. Although with the RTM version of Exchange 2013 the number of mounted databases was back to 50 (like with Exchange 2007), with Cumulative Update 2 (CU2) the support for a maximum of 100 mounted mailbox databases per server was added (for those who have the Enterprise Edition license).

Note: What's a Cumulative Update? Well, you may know these as hotfixes or rollup updates. Starting with Exchange 2013 the Microsoft Exchange Team adheres to a scheduled delivery model (every quarter, give or take) where a CU is released for the product. And at times these updates (and feature improvements) are rolled into a Service Pack (SP).

Now for years some folks have been wanting a change to SQL as the underlying database for Exchange but in the words of Exchange Team expert Ross Smith IV "SQL squeals like a pig where ESE is easy". My guess is that he means that ESE still outperforms SQL with regard to the type of transactions Exchange requires.

Now in legacy versions of Exchange the information store was a single process (store.exe). With Exchange 2013 the information store has been completely rewritten in C# and renamed the Managed Store. This new store has two processes, the Microsoft.Exchange.Store.Service.exe and the Microsoft.Exchange.Store.Worker.exe process. With each mounted database you have another Microsoft.Exchange.Store.Worker.exe process started up. So, each mount-request will create a new worker process which exits when a database is successfully dismounted. This means that the process of one database does not necessarily impact another process/database when e.g. it hangs.

So if you have 40 mounted databases there will be 41 processes working to support these. You can see the worker processes if you take a look at Task Manager on your Exchange server and have multiple databases mounted. Note in the figure below that the Store.Service.exe is present and used to help with managing the Store.Worker.exe processes.




Task Manager and the Store.Worker and Store.Service Processes

This helps isolate single database issues without impacting other active databases running on the same server. Database failover and physical disk handling have been improved, reducing IOPS utilization by + 50% and now supporting disk capacity up to 8TB. ESE has also been enhanced with deeper checkpoint depth for both active and passive database copies.

How the Database Really Works

Try to visualize the flow of email into your server. It's sent from a person's email client. The 0's and 1's are flying around the Internet until they arrive at your Mailbox server.



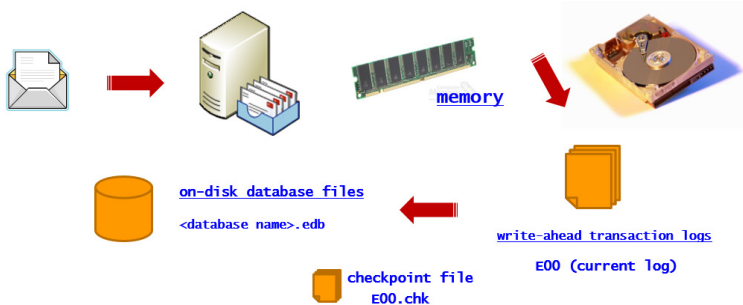
Email Entering the Mailbox Server

So what happens next? The email enters the Mailbox server, it goes through the memory and it's written first to transaction logs. Transaction logs don't do anything. They are 1 MB in size (1024 KB), a reduction from 5 MBs, which we had in legacy versions of Exchange (pre-Exchange 2007). Depending on how busy the Exchange Mailbox server is the data is also written to the database (.edb file) which is one monolithic file. Now the email message has been written to two locations: the database where the user will check in using his email client and retrieve his mail, and the transaction logs where the email is broken up, depending on its size, into 1 MB chunks.

Note: New email is not the only reason for transaction logs to be created. For every 'transaction' that occurs on the server (new email, deleted email, a change to an email message, a modified attachment...), that information is written to a transaction log.

Transaction logs are created in a log stream; in other words, they follow a sequential manner, kind of like a factory line where the current log is E00.log and then when it fills up it gets renamed and moved over. Although the current log written to might look like an E00.log, renamed logs might look more like E000000002E.log. The current log is not committed to the database and does not have its name changed until after it is filled to the full 1 MB capacity. Then, it is closed out.

There is a checkpoint file to keep track of which transaction logs have been added or committed to the database so that none are missed. If the database ever corrupts, and the transaction logs are safe, you can restore a backed up database and using the logs, the restored database can be made current by applying the changes recorded in the logs since the last backup.



The Database Process at Work

Reserve logs (.jrs) also exist (10 of them), just in case the disk space runs out and you need a little space (although a few MB of extra transaction log space won't buy you much). If the drive you are using for your database runs out of space the database dismounts and the store processes will stop and your Exchange server will be dead in the water until you clear up space.

MailboxDatabase.edb	4/2/2010 10:13 AM	EDB File	1,146,944 KB
tmp.edb	4/2/2010 9:07 AM	EDB File	8,256 KB
E01res00001.jrs	2/1/2010 2:06 PM	JRS File	1,024 KB
E01res00002.jrs	2/1/2010 2:06 PM	JRS File	1,024 KB
E01.chk	4/2/2010 11:52 AM	Recovered File Frag...	8 KB
E01	4/3/2010 9:07 AM	Text Document	1,024 KB
E01tmp	4/3/2010 9:07 AM	Text Document	0 KB
E01000000A0	2/12/2010 6:08 PM	Text Document	1,024 KB
E01000000A1	2/12/2010 6:10 PM	Text Document	1,024 KB

The Database, Transaction Logs, Reserve Logs and Checkpoint

We mentioned earlier the problem with storage groups being the transaction logs and that is because the transaction logs are intertwined within a storage group. In the event you have 3 databases in the same storage group, the transaction logs will continue to use the log stream approach for all three at once. This might have a negative impact on performance and disaster recovery at some point which is why the recommendation to use one database for each storage group was given, and why we no longer use storage groups in Exchange.

In the past, it was best practice is to move your database and transaction logs off the drive that holds the system files (basically, where you have installed Exchange, let's say the C:\ drive) and

then separate the database from the transaction logs by putting them on separate volumes backed by different physical disks. To go one step further, if you can place your databases on a form of striped volume (if redundancy is provided some other way) or a striped volume with parity (a RAID 5 setup) to enable fault tolerance, and if you can mirror your transaction logs, you can achieve the best-practice level of storage for your mailbox servers. It's also called "database per log isolation". This is still considered best practice for stand-alone (not using high availability) deployments. However, if you are using high availability (which is covered in Chapter 6) isolation of disks and logs is not required.

It's important to note what happens to your database and transaction logs over time. The database continues to grow. Microsoft wants you to keep it to under 200 GB but they have a maximum size of 2 TBs. The transaction logs grow too. At 1 MB each you can have thousands and thousands of them. But when you perform a full or incremental backup of your Exchange store databases logs are truncated (or removed).



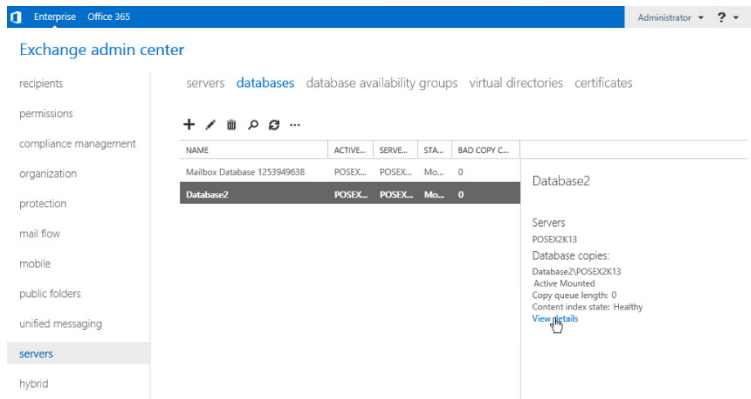
Note that logs typically purge daily after successful database backups. If a backup fails then the logs will not clear, so it is important to over allocate any disk that your logs will reside on. If you run out of disk space your database will shut down, so supersize your drives to ensure database uptime. Fortunately SATA disk is an option for new versions of Exchange as this makes disks cheap enough to accommodate your Exchange sizing needs.

Concepts of Database Management

One thing that makes Office 365 (or Exchange Online) very interesting is that you don't have a way to manage your databases. You cannot access the properties of your databases and configure settings and options because this is all handled by Microsoft. For

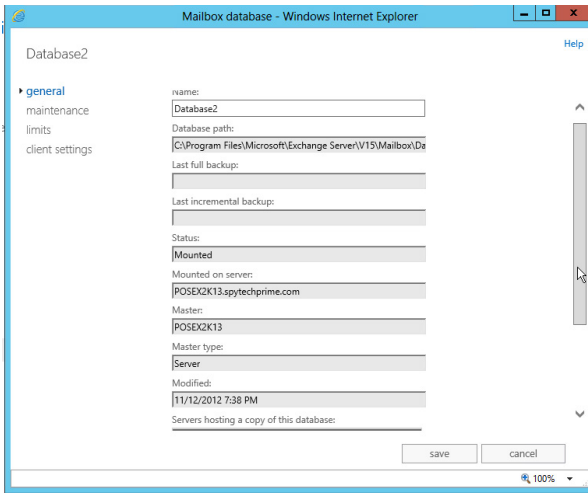
some administrators that is a bit stifling because they want to get their hands under the hood and alter settings.

What kind of settings and such can we alter? Well, one thing we've mentioned in this chapter is 'mounted' databases. You can mount or dismount a database. When mounted it is online and accessible by users to get their mail. When dismounted it's offline and inaccessible and maintenance tasks can be performed on the database.



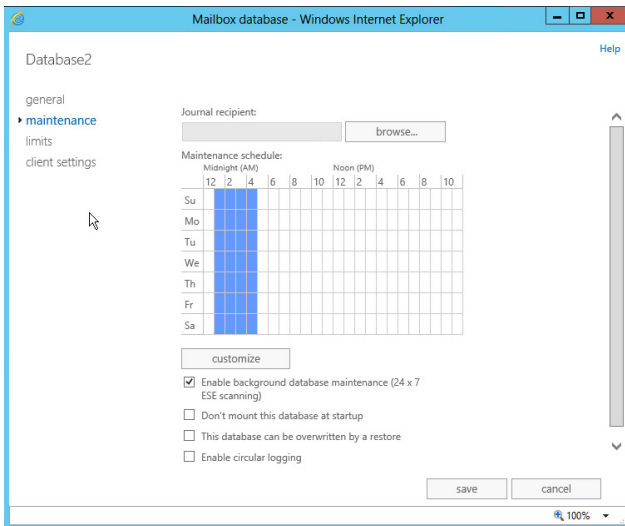
Configuration of Databases in Exchange 2013

If we go into the Properties of a database in Exchange 2013 we see four options: general, maintenance, limits and client settings. The general tab gives information about the database location, information about backups and so forth. You can see the last backup (full or incremental) and other important information.



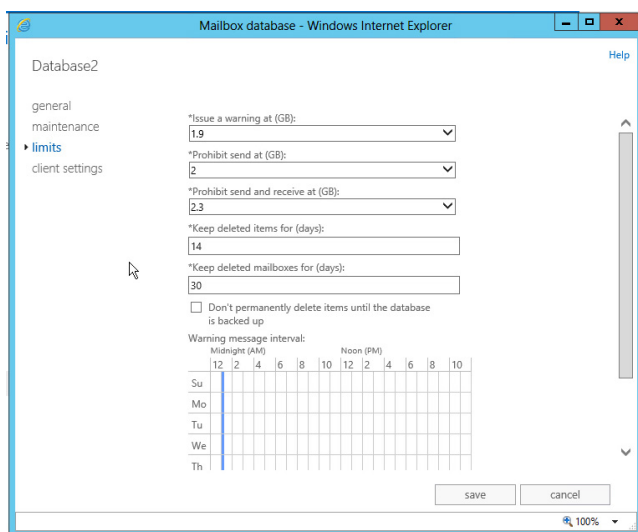
The General tab of Mailbox Database Properties

The maintenance tab allows you to configure standard journaling (discussed in Chapter 5: Regulatory Compliance), set the schedule for the database maintenance time and configure a few other options, including enabling circular logging.



The Maintenance tab of Mailbox Database Properties

Circular logging allows transaction logs to be overwritten or purged. This is fine if you don't need to worry about the transaction logs being retained (for example, if you use replication as a backup solution) but otherwise you don't want to enable circular logging because it eliminates your logs, which are helpful if you have to restore a backup of your database due to corruption or disk failure.



The Limits tab of Mailbox Database Properties

The limits tab allows you to specify the amount of storage the mailboxes in the database are allowed to have. You can see in the figure above that there are defaults, but these can be adjusted. You can see there are three stages: issue a warning, prohibit send, prohibit send and receive. When the final limit is reached, the server will stop receiving email from others and end-users will not be able to send.

You can also see that there are deleted items and deleted mailboxes settings. When a user deletes something from their Inbox it goes in the Deleted Items folder. If they delete it from that folder it isn't gone yet. They have 14 days (by default, unless altered here) to recover that deleted item before it is purged. So, deleted item recovery is a great option for mailbox item level

recovery. The same is true of deleted mailbox retention. If an employee quits or is fired and the administrator is told to delete their mailbox, that mailbox is still available for restoration for 30 days (by default). If the person returns within that timeframe you can restore the mailbox without looking to a backup.

The final tab, client settings, has the option to configure the ‘offline address book’. An offline address book is very helpful if an end-user is working and needs to send emails but cannot connect to the Exchange server directly (maybe they are on a plane) and cannot access the live Global Address List (or GAL). The offline address list can allow them to still find email addresses in that list. Note: the entire GAL may not be in that list, that’s why it is configurable. The admin can determine what that address book contains.

The Big Takeaways

The Exchange Mailbox server is all about the databases. Mailboxes live in those databases and managed thanks to the Extensible Storage Engine (ESE). All mail coming in and out is handled by the Mailbox server and mail items, calendar items, contacts and tasks are all stored in the database.


Database architecture has evolved and been enhanced over the course of Exchange’s lifetime. The current iteration is designed to allow for cheaper storage solutions (like JBOD) no longer requiring the expensive SAN solutions we needed in times past (although many admins still prefer the higher performance disks).

Email comes into the server and through memory into transaction logs and into the database .edb file. A checkpoint file keeps track of which logs have been entered and which ones haven’t (to ensure recovery if there is a crash). Those logs are truncated upon full/incremental backup (or purged automatically if you enable circular logging on the database).

The database has properties that you can configure including limits and deleted item/mailbox retention times for the database.

And there you have it. Time to move on to Chapter 4!

Chapter 4: Recipient Management



I changed all my passwords to “incorrect”. This way if I forget it will tell me “Your password is incorrect”

There are so many different types of recipients for an Exchange server. That may sound odd because you may be thinking the only recipient type is a user who is assigned a mailbox. But there are many recipient types, some of which have mailboxes and some that do not.

But let’s not jump too far ahead of ourselves. This chapter is going to focus primarily on the most obvious of recipient types, the user mailbox. Then we will review other recipient types and conclude with a discussion of public folders (now called ‘modern’ public folders).

Ok, so let’s get started!

User Mailbox Management

User mailboxes are the most common recipient type. Each mailbox is associated with an Active Directory account.

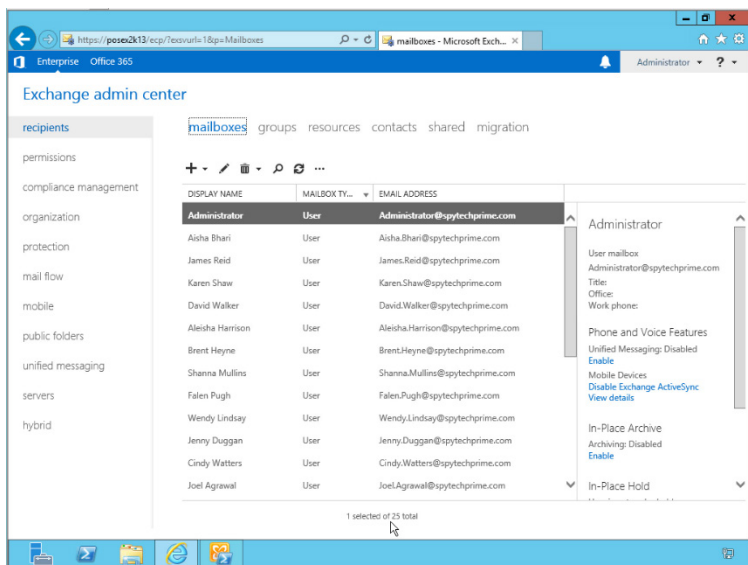
The account in Active Directory may already exist and you create a mailbox that connects to that account. Or perhaps you have to create both at the same time. Now which method you choose may depend upon whether or not the user was created prior to your installation of Exchange, which may be the case. Or it could be that you don't have permissions to access or create accounts within Active Directory and so you cannot create the end-users, you have to wait for the AD admins to perform that task. But either way, these two must co-exist (the AD account and the mailbox) to have a user mailbox in Exchange.

When you create the mailbox for an end user you can determine which database is going to hold that mailbox. If you have more than one database, you can put the mailbox in the one that is best for that end-user. Sometimes you need to balance out your databases and you can move mailboxes when you need to from one database to another.

Creating these accounts is super easy. No matter which version of Exchange you are using it's not rocket science to create new user mailboxes (or delete them if you so desire). And the real beauty of the user mailbox is in how easily it is for end-users to connect to their mailbox once complete. Because they are logged into their computer on the network with their AD credentials, when they open up Outlook (let's assume all modern stuff here... Exchange 2013 and Outlook 2013 so I don't have to explain pre-Autodiscover) the application can automatically find the Exchange server thanks to a service called Autodiscover. Credentials are passed and the mailbox is automatically located and the end user is connected up.

Seriously, this is an awesome feeling the first time you install an Exchange server, create a mailbox and see that user open Outlook and connect. It makes me feel good with every new installation. And even better is when you can send and receive email internally as well as externally (which requires a bit more

configuration to ensure mail can flow externally... but let's not get too far ahead of ourselves).



User Mailboxes in Exchange 2013

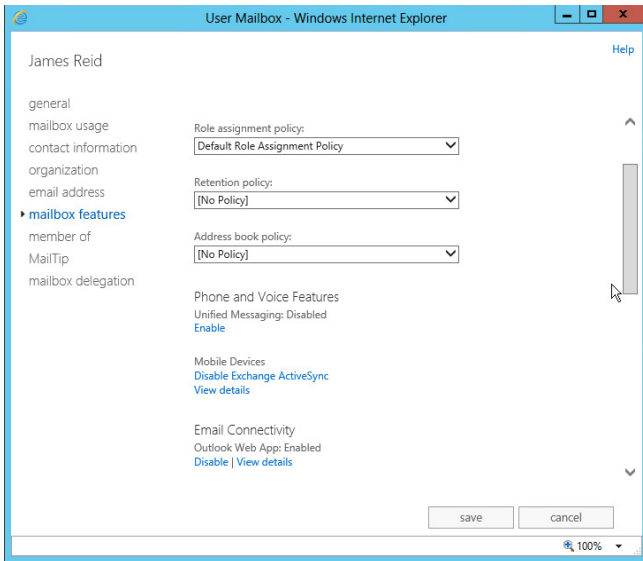
Ok, so you've created this user mailbox. They can connect to it through Outlook and send email (internally only to start). Now what? What configuration is available to you?

There are so many things you can configure with user mailboxes. You can configure them one at a time through the Exchange Admin Center (or select multiple mailboxes and configure in bulk) or you can use the Exchange Management Shell (the Powershell command line method if you recall from Chapter 1).

You might configure the following:

- Individual mailbox quotas (that take precedence over the ones set for the mailbox database)
- Deleted item retention settings (that also take precedence over the ones set for the mailbox database)

- Message size restrictions (send and receive maximum message sizes)
- Mailbox features (enable/disable) like Unified Messaging, mobile device support, Outlook Web App (OWA), etc... as well as policy choices



- Send As and Send on Behalf Of or Full Access mailbox delegation settings



Thu 11/22/2012 10:56 AM

Cindy Watters on behalf of James Reid

re: send on behalf of test

Note: Let's just dive deeper on that last one. Imagine you have a VP of Sales who has an assistant that needs to be able to send mail out but it has to look as if that mail came from the VP of Sales. If you give the assistant Send As permissions they will be able to do that through their Outlook. However, if you want the assistant to be able to send email but make it clear that it is coming from the delegate, you give them Send on Behalf Of permissions and in the From: line of the email it will be clear that

the email was sent by the assistant, on behalf of the VP of Sales. Obviously if you want a delegate to have full permission to open and use the VP of Sales mailbox as if it were their own, you can provide Full Access permissions to the assistant.

Other Recipient Types

We're not going to deep dive into all of these so if they don't all make perfect sense don't worry about it right now. Exchange admins may never even use some of these recipient types – never ever. Like this first one. Linked mailboxes are only used in specific deployment scenarios and so, as one example, if you aren't in an environment with Exchange deployed in a resource forest, you may never need a linked mailbox.

- **Linked Mailbox:** Accessed by users in a separate, trusted forest (often used when Exchange is deployed in a resource forest)
- **Groups:** There are several different types of groups you can create with Exchange and the two most common are distribution groups and dynamic distribution groups.
 - Distribution groups have a static membership (that you configure manually) and when you send an email to that group's email address, it goes to all the members in the group.
 - Dynamic distribution groups are based on criteria through filters so that persons are added or removed from the group based on their attributes. The membership of the group is derived at the time any given message is sent. So imagine a group where the criteria is location (Orlando) and a person named Sue is transferred to the Orlando branch office. The moment their criteria is changed in AD to reflect that their location is now Orlando, they are automatically added to that dynamic distribution group. Should their location change, they will be removed automatically as well.

- Resource Mailboxes: These are great for keeping track of equipment and scheduling the use of company equipment as well as meeting locations. This mailbox is created solely for scheduling, not for sending email and such. There are two types:
 - Equipment mailboxes can be used for scheduling and keeping track of projectors, laptops, even company cars or whatever needs to be managed
 - Room mailboxes can be used for scheduling meeting locations (like the conference room, the auditorium, the lab, the training room and so forth)

- Mail Contacts: These are contacts that have an object in Active Directory that is mail-enabled, but the email address is external (meaning there is no on-premise mailbox... the mailbox is handled by someone else... Gmail, Yahoo, whoever) so you can find them in the Global Address List (GAL) but they are contacts, not user mailboxes. You might use these for outside contractors that work for your company so that users can locate their emails easily in the GAL.

- Mail Users: These are users who have accounts in Active Directory, so they can log into your network, onto an AD domain, but they don't have a mailbox on-premise. They have external mailboxes. These might be for temp workers who you don't want to give mailboxes to but need to be able to easily email, or for other workers that fit in needed a login to the domain but not a mailbox.

- Shared Mailboxes: These are usually set up for collaboration purposes. Like a sales@companyname.com or an info@companyname.com so that multiple people can be given Full Access, Send As or Send on Behalf Of access and then persons can be allowed to monitor and/or send email from the common account.


Note: In addition to shared mailboxes there are two other collaboration mailbox types: site mailboxes and public folder mailboxes. Site mailboxes are new to Exchange and they require both Exchange 2013 and SharePoint 2013 to allow for a new form of collaboration through the Outlook 2013 client. They aren't widely used yet however.

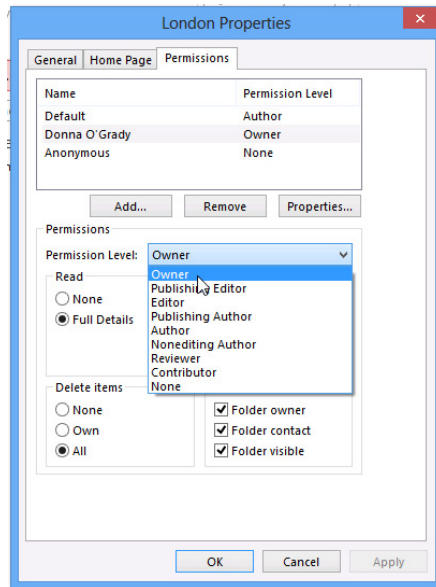
Modern Public Folder Mailboxes

In every version of Exchange up to Exchange 2013 when you wanted to use public folders for collaboration within your organization you had to create a separate database for it. Now I might be jumping ahead here because I'm assuming you know what public folders are, but you might not depending on your experience and the work environments you've been in, so let me back up here.

A public folder structure is usually started with a clear purpose in mind so that end-users can collaborate easily or share information easily with other within their organization. Perhaps it is designed by location to start (as you can see from the figure to the right).

Typically an admin will start the top-level folders and assign permissions to others to manage them going forward. Oftentimes public folders sprawl out of control and become the dumping grounds for all sorts of material that goes out of date. Preventing users from creating top-level public folders and requiring them to request public folders through a Service Desk can prevent this feature from sprawl and turning into a dumping ground.

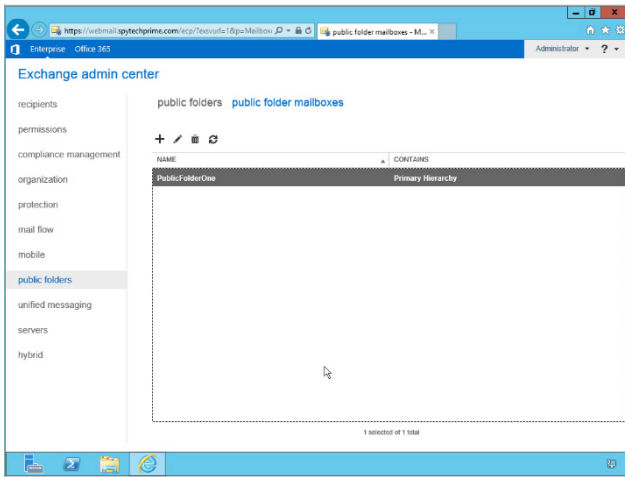




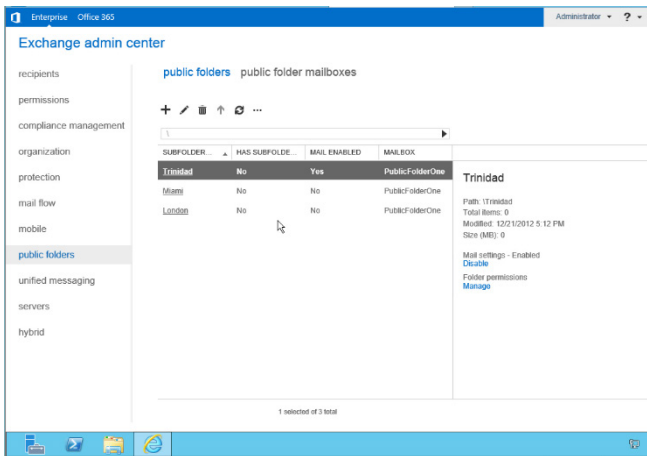
Public Folder Permissions

In legacy versions of Exchange one of the benefits to the way public folders were designed is that you could provide for availability and faster access to them by creating replicas. So imagine a large organization with a branch office in London and one in Trinidad. Let's say you had a top tier folder for each location. You could have an Exchange server in each location with a public folder database on those Exchange servers. Perhaps the original public folder for London was created on the Exchange server in London. You could create a 'replica' of that public folder and place it on the server in Trinidad. This would accomplish two things. There would now be a duplicate of that folder and if folks in Trinidad wanted to access it they could get to it faster because it is now local, instead of going over the wire to the London server for that folder and its contents.

Starting with Exchange 2013 it doesn't work like that anymore. Now you create a public folder mailbox and select a mailbox database to host it. Then you add public folders into that public folder mailbox. So the separate database for public folders is gone as is the separate replication architecture for replicas.



Public Folder Mailbox in Exchange 2013



Public Folders in Exchange 2013

As with legacy public folders you can see from the figure above that you can have subfolders nested within your top tier folders and you can mail-enable a public folder so that folks can email items to it directly. You can also configure folder permissions through the EAC as well as through the Outlook client.

When designing your public folder top-tier layout you will want to put folders in mailboxes that are close to the people who will access them often. With modern connectivity speeds this may not be a problem for folks to access a public folder located in another location (but that will depend on the connection speeds of the remote location). On the plus side however, because the public folders are now in a normal mailbox database they come under the availability features (aka DAG) that we will discuss soon in Chapter 6.

Now, if you have an existing public folder structure and you need to migrate over to this new structure it's going to take some research and work on your part.

The Big Takeaways

There are a lot of different recipient types but the most commonly used is the user mailbox. There are a lot of configuration options for user mailboxes including individual quotas, deleted item retention, features enabled/disabled and more.

Additional recipient types include linked, shared, and resource mailboxes as well as contacts, groups and more. Each has its own function within your Exchange environment (they've thought of pretty much everything).

Collaboration mailboxes like shared mailboxes and site mailboxes (when configured to work with SharePoint 2013) help within your organization. Public folder mailboxes ("modern public folders") are a legacy collaboration solution with a modern twist to it. Now we create a public folder mailbox in a mailbox database (as opposed to a public folder database) and add our public folders (and/or subfolders) to that PF mailbox.

Now it's time to get into some legalese with you folks in Chapter 5 and Regulatory Compliance.

Chapter 5: Regulatory Compliance



*“Blah, blah, blah... Sarbanes-Oxley...
blah, blah, blah... sexual harassment...
eDiscovery... blah, blah, blah... prison time...”*

What do you think of if I say Enron, WorldCom or Tyco? These were highly publicized corporate and accounting scandals. They exposed a lack of proper supervision and a lack of strong regulation against the fraudulent practices taking place in major corporations. As a result billions were lost and public confidence was shaken. With Enron, one of the world’s largest accounting firms, Arthur Andersen, collapsed due to evidence that was brought up from company email.

To combat this unfortunate trend toward fraud the result has been a slew of new laws that improve government auditing of businesses. These problems, combined with issues like employee harassment cases (sexual or bullying) has led to an increase in regulations, many of which apply to communications within an organization – email being a primary form of auditable communication. Note: It’s impossible to track and audit every

conversation (some are discretely held in the shadows), but email is another story. That's where the Exchange admin comes in.

What is Regulatory Compliance?

Basically, there are a variety of regulations that currently exist including:

- Sarbanes-Oxley Act of 2002 (SOX)
- Security Exchange Commission Rule 17a-4 (SEC Rule 17 A-4)
- Gramm-Leach-Bliley Act (Financial Modernization Act)
- Financial Institution Privacy Protection Act of 2003
- Health Insurance Portability and Accountability Act of 1996 (HIPAA)
- Uniting and Strengthening America by Providing Appropriate Tools Required to Intercept and Obstruct Terrorism Act of 2001 (aka the Patriot Act)
- European Union Data Protection Directive (EUDPD)

Some of these you may be aware of. These are just the highlights, there are many more in effect. And different countries have their own versions of these laws.

How do these laws affect you? How do you determine what you need to do for your organization? Well, every company is different. If you are in the healthcare, government, financial fields than there are regulations that apply to you that may not apply to other organizations. For example, you may be obligated to retain communication data longer than another company. It's essential that you have a legal advisor (or team of advisors) to make sure you are clear on what you are obligated to do.

Now while some organizations have a legal person or team to assist (or even have compliance officer positions) smaller businesses need some help in this regard. To assist there is a site called

Business.USA (<http://business.usa.gov/>) that provides smaller businesses with resources.

Let's answer a quick question... whose side are you on in a legal case against members of your organization? Answer: the company's side. Not the individual members, even if you know them and like them. Remember, if a person went outside the bounds of the law and, as an example, sexually harassed someone in the company, you need to be able to protect the company. The company must be able to supply the communications to either exonerate the person of wrongdoing, or provide transparency to show that they don't condone the behavior. If they don't provide such communications they can be held liable for not being compliant with the regulations. So if you are responsible for the Exchange environment you need to know the technologies involved and know when they truly meet enterprise grade regulatory compliance, and when they don't.

There is one line from Sarbanes-Oxley that is worth keeping in mind, not to scare the reader but to educate as to the seriousness of complying with investigations. It says in Section 802(a) of the SOX, 18 U.S.C. § 1519:

“ Whoever knowingly alters, destroys, mutilates, conceals, covers up, falsifies, or makes a false entry in any record, document, or tangible object with the intent to impede, obstruct, or influence the investigation or proper administration of any matter within the jurisdiction of any department or agency of the United States or any case filed under title 11, or in relation to or contemplation of any such matter or case, shall be fined under this title, imprisoned not more than 20 years, or both.”

Ahem...<cough><cough>... yikes! Now just so you know, I have not heard of any of my fellow Exchange admins going to jail but it's good to see the seriousness of this part of an Exchange admins job with regard to complying with company policy and legal requirements.

What it Boils Down To

So what is regulatory compliance? Following the laws for your business. And in the end what does that mean? Primarily, it means being able to discover all relevant communications (email, IM, etc...) for the period of time stipulated by your corporate policy, which should be based on the legal requirements. It also means being able to provide through auditing that no tampering has taken place (ie. to prove if an email was read or not by an individual, to prove that someone else didn't send an email that is inappropriate or illegal, etc...).

Focusing on the first part, let's say your organization is responsible for discovering relevant messages for 7 years. This is not a backup (although in times past that's all we had) and it's not availability (which is just the current state of your existing operation) but it is retention of relevant email data (and by relevant we mean no spam/junk). It's an archive of your data.

In the past we may have relied on backups of the data, but those have not proven to be as reliable as an official archive nor are backups easy to search and discover information on so they prove time consuming in the event of litigation against your organization.

So, if you were asked to provide all emails from Joe User to Jane User from 3 years ago between the months of March and July would you be able to do it? Without an enterprise grade archive solution, one that captures all relevant email and makes it tamper proof, and one that provides impressive discovery tools that might be a ginormous task.

Note: Some companies that do not have the strict legal requirements to retain data have begun to establish corporate policies that openly state that they retain nothing. 15 day retention and that's it. While this may be legal in their case, it may not be wise. In the case of litigation your accuser will have their proof, but you won't. Why let the opposition be the only one with cards at the table?... legally speaking that is.

Exchange Regulatory Compliance Features

Every flavor of Exchange in the past 10 years has improved upon the regulatory compliance features available to administrators.

Here is a list of the current line-up of built-in tools (some of which are enterprise grade and others that are not quite there yet):

Personal Archive

You may be thinking ‘Awesome, Exchange has a built in archive solution, that’s going to save me money. Well, not so fast. Note the word ‘personal’ at the front. To explain see if the following scenario sounds familiar:

You’re working away and you get a message that your mailbox quota has been reached. That’s funny, you think... because I just deleted stuff last week. But all that ‘stuff’ is sitting in your Deleted Items and so it is still part of your mailbox size. And your quota is a stingy 1 GB (let’s say) so you need to clean house a bit (Inbox, Sent Items, Deleted Items, etc...). Easy enough, so what do you do? You’ve been taught how to create a .pst file for your mail. This file is on your desktop system and you can create it through Outlook and move tons of email over to it yet still access it from Outlook easily when you need it. A great solution! That resolves the quota problem right?

Indeed it does resolve the quota problem. However, it creates a new problem for admins. Discovery becomes harder if not impossible when end-users have the ability to pull stuff to their desktops (which may not be backed up).

The solution? The ‘personal archive’ feature in Exchange 2013. This is like a second mailbox in a sense in that you can place the archive in the same database as the existing mailbox for a person or you can put it in a completely different database (on a completely different system or drive). The value here for admins who are crazy about performance and mailbox quotas (due to the expense of putting mailboxes on high-performance, expensive disks) is that you can put the archive on cheap disks and you can make them larger quota sizes (or unlimited if you like) so that your end users now have a place to put old email without moving them to .pst files. Archived messages can be searched for just like messages in your primary mailbox (both are indexed for search).

The benefit for the users is that now they can see the archive email even when connecting through Outlook Web App (OWA), something they couldn't do with a .pst file. Note: You cannot access your personal archive through ActiveSync (mobile device) connections.

So, can the personal archive feature assist with being compliant with regulations? It can. It helps ensure data is discoverable and not floating around on desktops. However, this feature is not quite enterprise grade. Part of the reason is that end-users still have the ability to delete their own email. Let's see how you can prevent that with our next feature.

In-Place Hold (also Legal/Litigation Hold)

One of the biggest nightmares for compliance with Exchange 2013 is that out-of-the-box users have the ability to delete their own mail permanently. Unless you place mailboxes on legal hold an end-user can delete incriminating evidence before it is backed up.

Legal Hold (aka Litigation Hold) is a feature in Exchange that allows an administrator to stop users from deleting email. They don't know their mailbox is on hold so they still delete items, but they are still discoverable.

Note: The Recoverable Items folder is used with In-Place Hold when items are deleted. They are removed from the users view but because they are in Recoverable Items they can be discovered through an In-Place eDiscovery search.

The update to Legal Hold in Exchange 2013 is called In-Place Hold. This allows you to focus not on the entire mailbox but on specific items and timeframes. You can specify what you want to hold by using keywords, senders and/or recipients, start/end dates, message types (email/calendar items, etc...)

This all sounds cool doesn't it? But there are two issues with it. First off, you don't typically put a mailbox on hold unless requested to by HR, Legal, whomever. And by that point the person (if they have half a brain) has already deleted stuff. (The odd thing about deleting inappropriate or incriminating stuff is that you can only

delete one end of it. The other person still has it!!! And if they are in the same company it's still in Exchange!)

Now, you could be very proactive and just put all mailboxes on legal hold but that would just bloat your database really fast. Possible, but not your best course of action.



My honest feeling on this is that your best course of action to ensure a solid enterprise grade archive that is untouchable by end-users but helpful should they need to restore accidentally deleted email, is to go with a third-party, cloud-based solution. My personal choice in this regard is Mimecast. All email coming in/out/internal gets archived. End-users know it, so that is already a deterrent from them using your email system to send stupid/incriminating email. It's all held in the cloud and easily discoverable. End-users can restore emails they delete but cannot delete emails from the archive. Perfect solution!

In-Place eDiscovery

At times your organization may be required to “discover” content within Exchange (or Exchange Online). Reasons may be due to corporate policy, regulatory compliance or a good ol’ fashion lawsuit. Someone says in appropriate emails have been sent and you’re asked to provide Legal with all correspondence between all persons involved. An overwhelming task without tools for that level of discovery (aka eDiscovery).

Now with Exchange 2010 we had a feature called Multi-Mailbox Search which wasn’t super awesome. I didn’t like it personally but it was free and built-in. The next level in Exchange 2013 is called In-Place eDiscovery, which is a much more polished version. Using In-Place eDiscovery you can estimate your search results,

preview the results, copy the results to a Discovery mailbox or even over to a .pst file for transport elsewhere.

Again, you can only discover what you have right? So the discovery tool is only part of this puzzle. It's good to see it improving. And with Exchange you can give folks in authority (like HR) their own access to perform eDiscovery so that you don't have to stress about it as the Exchange admin.

Messaging Records Management (MRM)

You can tell users they need to clean up their mailboxes but oftentimes they misunderstand. They may ignore you completely or delete stuff only to have it sitting in the deleted items folder (which is still a part of the mailbox and takes space). It was clear to Exchange developers that an automatic way to help out the end users was needed. Email lifecycle management is essential.

Messaging Records Management was introduced with Exchange 2007 (MRM 1.0) and it involved Managed Folders. This evolved into MRM 2.0 in Exchange 2010 and now 2013 and it's a combination of Retention Policies and Retention Tags to assist with retaining messages that are important to retain and removing messages that are non-essential.

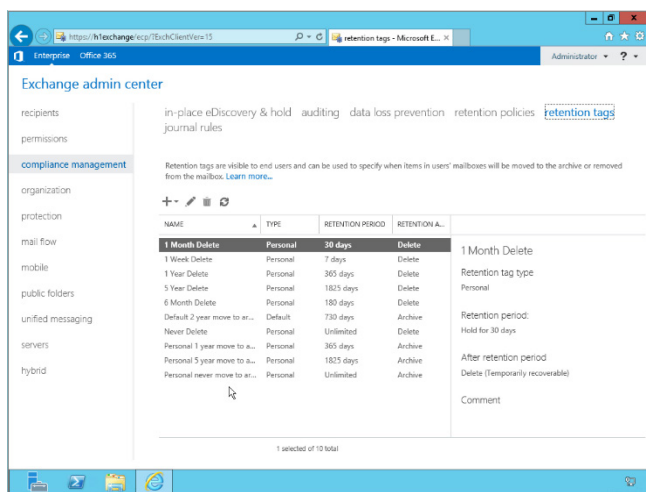
The value to a solution like MRM is that you can reclaim wasted storage space and assist end-users to be more productive by removing clutter. You also help in retaining the important messages for a necessary period of time.

So perhaps to help you start thinking about when you might use this imagine again the scenario above where you have end-users deleting items but putting them in their Deleted Items folder. What if there was a policy that said all email in the Deleted Items folder would be deleted after 90 days. That would resolve the bloat there. What if you had another policy that said all mail past the 180 day mark in the users Inbox would be moved to the personal archive (assuming you are using that feature). Can you see how this feature might help because it works automatically instead of leaving it up to your users.

So there are retention tags with retention settings that can apply to the whole mailbox or to specific default folders (like Inbox or Deleted Items) or to specific messages (depending on the tag time). These tags are added into retention policies, which are then applied to end users. Note: you can only have one policy applied per user. The Managed Folder Assistant (a service running on the Mailbox role) applies the policy to the end-users' mailbox.

There are three different types of retention tags:

- **Default Policy Tag:** Applies retention settings to untagged mailbox items
- **Retention Policy Tag:** Applies to default folders like Inbox, Deleted Items, etc...
- **Personal Tags:** Allow users to apply to folders created or individual items



Retention Tags in Exchange 2013

Journaling

Journaling is an interesting feature in that it allows you to set up a separate mailbox (that you call your “journal” mailbox) and you

allow email to be recorded into that mailbox for easy review or to use as a form of archive.

Microsoft makes it clear that journaling may not satisfy your legal requirements for regulatory compliance so don't think of it in the same category as an archive, however it may help. One way it might help is if you perhaps have a company or legal policy that requires a review of employee/client communications. With journaling you can easily review the communication by searching through the journal mailbox, which captures the traffic you have specified (depending on your configuration settings).

There are two flavors of journaling:

- Standard journaling, which is performed at the database level and has the journaling “agent” focus on that entire database (and you choose if you want incoming/outgoing or both journaled).
- Premium journaling allows the agent to focus on a granular level on individual mailboxes or distribution groups (and it requires an Enterprise CAL, or client access license).

Note: Office 365 doesn't allow access to the databases so you cannot perform standard journaling, however some of the plans you can choose will allow you to perform premium journaling. Make sure you know the features allowed in the plan you choose.

Transport Rules

Transport rules is one of my favorite Exchange features ever since they showed up in Exchange 2007. They're similar to junk email rules that you might set up for your Inbox, the difference being that Inbox rules work when a message is delivered whereas transport rules are handled in transit.

Basically there is a transport rules agent that processes the rules created so messages are analyzed as they move within your organization and you establish the criteria to be met. Now in legacy Exchange (2007 and 2010) the transport was handled by the

Hub Transport server role (and, if used, the Edge Transport server role in the perimeter). In Exchange 2013 the rules and the agent are handled by the Mailbox server role.

The way you create a transport rule is to determine the condition, the action and any exemptions. So here is the basic anatomy of a rule:



Anatomy of Transport Rules

It's NOT that complicated and that is one of the things I love most about transport rules.

Now as for the practical uses for these rules, they include the following:

- You may need to apply disclaimers to messages that leave your organization and rather than leave it up to end-users to configure this in their signature you can configure a Disclaimer transport rule.
- You can prevent inappropriate content from entering or exiting your organization based upon specific key words or file types.

- You can filter information that you have listed as being confidential or redirect inbound/outbound messages for inspection and approval before being sent.
- You can track or archive messages that are sent from specific individuals or teams of individuals to ensure compliance is met.

Data Loss Prevention

Data Loss Prevention (DLP) is a new feature in Exchange 2013 and what it does is piggyback off of transport rules to help protect your organization from giving out sensitive information.

Sometimes people are so comfortable with email and trust implicitly that is a safe means of communication that they put their personal sensitive information in it: banking information, SS#'s, credit card information, etc...

DLP policies allow us to configure special kinds of “transport rules” for Exchange to look for this sensitive data being sent. Policies can be outright enforced, or they can be enabled as tips so that users are warned by a Policy Tip that they are about to violate a policy (although this requires users to use Outlook 2013).

There are different ways to implement the DLP policies (like creating your own custom policy from scratch) but the easiest way is to pick a policy from one of the templates.

You'll note in the following figure that there are templates based on US laws and if you could scroll you would see that there are a mixture of other laws included from other countries.

Note also that the definitions for each template make it clear that the use of the policy doesn't ensure compliance with any regulation. The names are to help you make mental connections not to imply compliance.

DLP policy from template

Name:

SS# Confidentiality

Description:

*Choose a template:

<ul style="list-style-type: none">U.K. Access to Medical Reports ActU.K. Data Protection ActU.K. Financial DataU.K. Personal Information Online Code of Practice (PIOCP)U.K. Personally Identifiable Information (PII) DataU.K. Privacy and Electronic Communications RegulationsU.S. Federal Trade Commission (FTC) Consumer RulesU.S. Financial DataU.S. Gramm-Leach-Bliley Act (GLBA)U.S. Health Insurance Act (HIPAA)<li style="background-color: #cccccc;">U.S. Patriot ActU.S. Personally Identifiable Information (PII) DataU.S. State Breach Notification LawsU.S. State Social Security Number Confidentiality Laws	<p>U.S. Patriot Act 15.0.3.0</p> <p>Helps detect the presence of information commonly subject to U.S. Patriot Act, including information like credit card numbers or tax identification numbers. Use of this policy does not ensure compliance with any regulation. After your testing is complete, make the necessary configuration changes in Exchange so the transmission of information complies with your organization's policies. Examples include configuring TLS with known business partners or adding more restrictive transport rule actions, such as adding rights protection to messages that contain this type of data.</p>
--	--

DLP in Exchange 2013

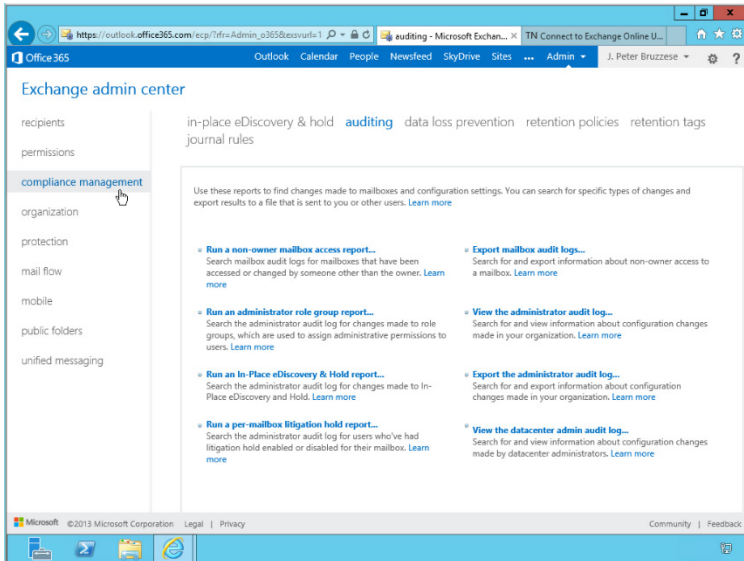
Mailbox/ Administrator Audit Logging

Auditing is very important in compliance. You have to be able to show that the mailbox owner is the one who sent the email in question and not someone with access permission. Even Administrators have to be auditable in that they can prove who they gave permissions to. Everything can be logged. Note... everything CAN be logged, but not everything is logged by default.

You have to enable mailbox audit logging and oddly enough this has to be done through PowerShell and the Exchange Management Shell.

With administrator audit logging it's interesting that everything is logged by cmdlets in PowerShell. So if the administrator creates a new mailbox you will not see the Exchange Admin Center method he used but the New-Mailbox cmdlet that ran in the background as a result. Everything done in the web-based console is actually performed through PowerShell commands behind the scenes, so

all commands are recorded with the exception of the Get- and Search-cmdlets because these do not affect changes.



Auditing Options in Exchange 2013

Information Rights Management

IRM is one of my least favorite features in this list. Information Rights Management (IRM) is a solution to assist with information/data leakage by establishing rights-protected content. Traditional methods that assist with this include encryption, company policies and such, but IRM works with a feature called Active Directory Rights Management Services (RMS) so that you can restrict recipients with regard to what they can do with an email they receive.

For example, do you need to prevent an email from being forwarded or printed? Do you want to prevent an email from being cut and pasted using the Snipping Tool? Do you want to give an email or attachment an expiration date? That's what IRM is for.

Now you might be thinking, but J. Peter these are awesome capabilities that IRM enables, how could you NOT like it?!

Well, look... here is the problem. First off, it's not the easiest thing to set up and configure. Second, IRM is a deterrent, not an absolute prevention feature. For starters, it only works for in-house recipients. So once that email leaves your organization you lose that control. And even though you cannot use the Snipping Tool or Print Screen to capture emails you can use a third-party tool. Or you can take a picture of your screen with those nifty camera phones the kids are all the rage about these days. Of course we ol' timers could just take a pencil and paper out and jot down the email content. So... it's a deterrent to a degree but I'm not convinced the effort involved justifies the rewards obtained (aka the ends do not justify the means in this case).

The Big Takeaways

The world scene is ever changing. With scandals mounting the decision to fight back with stricter regulations with harsher penalties seemed to be the only resort for lawmakers. To be in compliance requires a great deal of effort and with email being such a key player in corporate communications it falls to the Exchange admins to institute proper levels of archiving and ensure eDiscovery and so forth.

Exchange is increasing its compliance management features with each release (and in Office 365) but currently there are aspects to these features that are not enterprise-grade or they simply require too much effort to ensure compliance whereas a third-party solution that handles the archive side with robust eDiscovery tools may be worth the price tag.

Listen to your legal team, do your homework on third-party options, and make wise choices. Oh... and stay out of jail!

Chapter 6: High Availability and Site Resiliency



Are you daydreaming about clones again? Every time we configure the DAG you have that same daydream!

High Availability with Exchange. No subject has been as much fun to discuss over the past few Exchange editions (starting with Exchange 2007). I feel this is one of the most valuable features in Exchange and the Exchange developers only continue to make it better and better.

Obviously at this level, the primer level, we won't be going into the nitty gritty of it all. But by the end you'll understand that there are ways to ensure availability of your Exchange environment (and I don't mean your Mailbox servers only).

To truly grasp the present, we have to take a step into the past once again.

High Availability and Site Resiliency

In the old days we talked about ways to make servers more fault tolerant. In other words they can lose a hard drive but still keep working. To accomplish this required redundancy. Multiple power supplies, multiple drives configured in mirrored sets or RAID 5 arrays. More expensive hardware and software solutions were developed. Servers were clustered to either balance their load between them, improve availability or provide for higher performance by combining processing ability. To increase the availability of your servers and services was possible... but it was going to cost you lots!!!

Exchange, as a mission critical application, became a heavy focus with regard to availability promises and Service Level Agreements (SLAs). SLA's might be made between a buyer and seller, between a company and a set of contractors, even between the company and its own employees (individuals or teams). It's a pledge of sorts that they will maintain a certain level of quality work, often demonstrated by servers being up (uptime) and available to serve. Obviously uptime is worthless if the services aren't available, so just having a server continuing to be up does us no good. This is where the concept of 'business continuity' replaces the discussion of fault tolerance from times past. Hardware, software, application, network and ultimately, end-user access to all of this. You can have everything running, but if the user can't access it from where they are, it might as well not be running at all.

Note: SLA's often indicate the amount of uptime in percentages (like 99.9%) with some kind of restitution provided should this not be met.

It's not only important to have redundancy and availability of a disk (with databases), but of the whole server, of multiple servers, connections internally, connections to the Internet... you don't want a single point of failure to keep people from accessing their email. And so you do your best to eliminate those single points

of failure (and yes, that can get expensive). You may even have a secondary datacenter set up to allow for site resiliency. Or you might use a Branch office to be your secondary location.

Two terms you may hear include Recovery Time Objective (RTO) and Recovery Point Objective (RPO). This is usually part of a SLA and they stipulate the RTO (acceptable time without service being available) and RPO (how much data, past and present, must be restorable in the RTO).

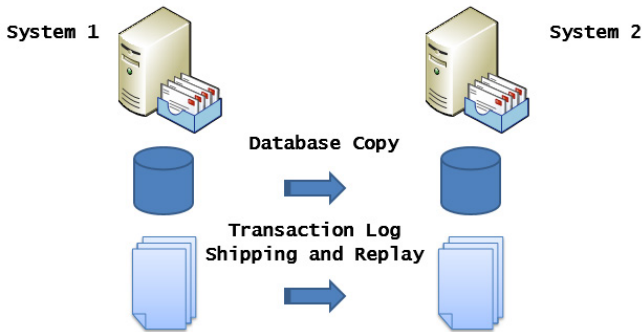
In the past these were negotiable terms and companies knew that the faster the recovery and the more data that need to be restorable, the more the expense. But in modern times you often hear people say the RTO is immediately, and the RPO is up to the nano-second. They want no downtime and they want all their data. Now that is a challenge... and Exchange has been designed well to try and handle that challenge.

Mailbox Server Availability and Resiliency

The secret to mailbox availability and resiliency lies right in the very database architecture we discussed in Chapter 3. Do you remember we talked about email coming into a server and going first into transaction logs and then into the .edb database (aka the active copy of the database)? Well, if they can go into one database, why couldn't those little 1 MB files be shipped over to another server with a duplicate of the database and be played (or re-played) into that database to provide for a secondary copy (aka a passive copy). And so that is what they did!

Continuous Replication

Starting with Exchange 2007 a feature called continuous replication was introduced where the database is initially copied and the log files are shipped over and replayed constantly to keep the passive copy of the database up-to-date.




Now you might be thinking, there has to be more to it than that right? Well, that's the underlying concept. Obviously there were some kinks to work out in the beginning.

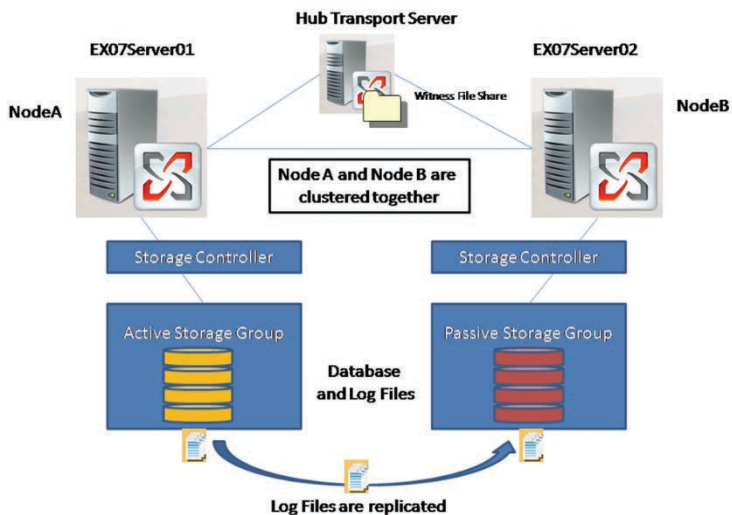
CR and Exchange 2007

With Exchange 2007 there were several continuous replication offerings (only one of which might truly be considered “high availability” because it used clustering services and was automatic). Let's review the 2007 offerings:

Local Continuous Replication (LCR). LCR is a single server solution that uses asynchronous log shipping and replay from one set of disks to another. The solution is one that requires a manual switch to move from the primary copy of the data to the secondary. It's called the “poor man's cluster” by some although there is no clustering technology used. The problem with this solution is that it only provided disk redundancy and resiliency. So if the server died, the mailbox server was off-line.



Cluster Continuous Replication (CCR). CCR is a clustered solution that only allows for 2 nodes in the cluster where one is the active node and the other is the passive node for automatic failover. This solution allows for two different systems and two different sets of storage offering a greater level of availability because single points of failure are eliminated. Asynchronous log shipping and replay is used in this solution to keep the database up to date between the active and passive copy of the data. This was the only “true” high availability offering because clustering services were used and so the process for failover was automatic. It offered both server and disk redundancy but more was needed for site resilience.

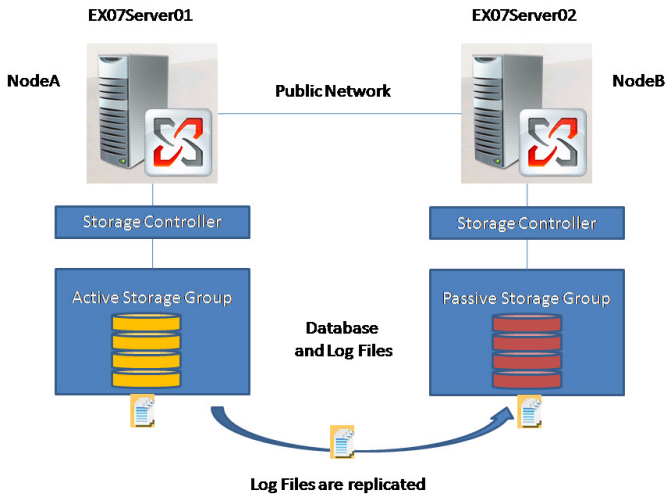


Note: CCR uses a feature called the transport dumpster to help ensure the passive is brought up to date with email that may not have been shipped over. The transport dumpster holds on to email for a brief period of time. It's been replaced by Safety Net in 2013.

Single Copy Clusters (SCC) SCC is NOT a continuous replication solution. It was available with Exchange 2007 RTM though so I thought I would mention it here. It offered a similar solution to what we had in Exchange 2003, multiple systems with a single storage group that is shared between the nodes of the cluster. Again, we have one active server with one or more passive servers waiting for a failover. This solution did provide for server redundancy and the disks were usually on some kind of an array with RAID providing the fault tolerance for the disks.

Standby Continuous Replication (SCR) SCR was introduced in Exchange 2007 SP1. It has the same overall features as LCR in that there are no clustering parts being used and the goal is to allow for availability from one site to another. It's a continuous

replication solution with a manual failover if something goes wrong.



Ok, so this is what we had with Exchange 2007, now let's move forward to see what they did with 2010.

Exchange 2010 Database Availability Groups (DAG)

So, you recall with Exchange 2010 that storage groups are gone and the focus is on the database right? Well, something else that is gone are all the high availability (and semi-high availability) offerings. What?! Yep, everything you just learned is not in 2010. LCR, CCR, SCR (even SSC) all gone. But... continuous replication has been retained and put to work in a new solution called the Database Availability Group (DAG). And it is awesome... seriously.

So, DAGs use continuous replication. There is an active copy of the database (that's the one everyone still connects to for their mail). And you can have a passive database copy too. In fact, up to 16 servers can be part of a DAG. The design options are

outside this book's scope, but suffice to say 16 servers means you shouldn't be losing your availability all that easily if you design and deploy this properly.

In addition to continuous replication a DAG uses clustering components like heartbeats and a witness server (with a witness directory) to connect members of a DAG and maintain quorum. Ok... I know, I know... I just used too many unfamiliar (unexplained) terms and concepts and you felt nauseous. Let's go through this slowly.

A heartbeat is a simple method that servers in a cluster use to "check-in" with one another to ensure they are still alive (so to speak). It's essential a way for the servers to say "I'm ok!" back and forth between them. Does the lack of a heartbeat indicate the server is down? Well, it could mean that. If you have two servers, one has the active database and the other has the passive and the passive no longer gets the heartbeat from the active... it thinks "I must step up here and become the active!". But what if the network cable just got cut (or something crazy like that)? If the passive becomes the active... while the active is still up and running... this is bad. Like crossing the streams bad (look it up). The terms we use in the clustering world are split-brain syndrome and world chaos. Neither sounds good right?

So how do we prevent this? Well, two servers aren't going to cut it. You need to at least have a third server to be the witness. And with DAG it doesn't have to be another Exchange server, it can be any server acting as a witness server with a witness share. This file share witness resource is really only needed when there is an even number of DAG members and so you need the witness as a referee to maintain quorum.

Techie Note: As you add servers to the DAG and go from odd to even (2, 4, 6, etc...) the quorum model is changed

automatically from Node Majority to a Node and File Share Majority model.

Now quorum is a term that is used in many situations and it applies to a voting process. It's a "consensus of voters". So in our scenario above with 2 servers, if there is a third file share witness to provide quorum and the line is cut between the active and passive, the passive will check with the witness to see if it should take over. If the witness can still communicate with the active, the passive stands down.

Now expand this concept and consider that we might be dealing with multiple sites and multiple passive copies of a database. A lot of communication has to continue working for this to succeed.

You might be thinking, do I really need all those passive copies? Well, if the active fails, the passive will automatically go into action and your end users will never know there was a hiccup. If two go down, a third would be nice. If the site goes down, it sure would be good to have a secondary location (Branch or datacenter of some sort) with another passive or two. You can see how this could get expensive too. Each server is hardware (even if virtualizing) and software and licensing and administration time and maintenance and so forth.


Another value to the passive copies is the ability to apply a lag on the log replay during the configuration. This will make it so that if a virus or corruption hits your active database you have time to revert back to a lagged copy that wasn't affected.

To create a DAG you simply create the Database Availability Group and give it a name. You then add the members to the group (these would be Mailbox servers and they can only be members of one DAG at a time). And then you can consider which of your active databases you would like to have passive copies of on the other members of the DAG.



Lagged copies or database copies can also exist in remote datacenters for the ability to provide a high availability solution to your end users.

Consider the DAG configuration below:



There are 3 servers and they are all in the same site in this case. Each has an active database (just one for simplicity). We've created a DAG and put all three servers in the DAG as members. We decided to make 2 passive copies of each database and one of the passives is a 2 hour lagged copy (just as an example). If System 1 goes down, System 2 will be ready to make the Passive DB1 into the active until you repair System 1.

Theresa Miller says "think of it like RAID, but instead of disks you are using databases".

DAC Mode

One thing we haven't talked about is failing over between sites. Is this possible? Absolutely. And if everything is ok with regard to the connections to the secondary site there is no different than if the server holding the passive copy was local. However what happens if you have two sites (a primary datacenter and a secondary datacenter). Imagine having 2 Exchange Mailbox

servers in each site with a witness in the primary site. What happens if the primary site has a blackout? Well, it's 3 to 2 and the 3 are down from the loss of power right? So at that point even though we have 2 fully functional servers in the secondary datacenter they will not mount their passive copies and immediately provide a failover. They have no way of communicating with the witness to provide quorum. So are you dead in the water?

Well, in this case if the blackout were to continue for period of time and it was decided you need your email up and running (perhaps you have workers in some locations with power, they simply need to access their email) you can provide a manual switchover (not a failover). To perform the switchover you have to manually shrink the cluster to regain quorum. So this is where DAG provides site resiliency not as a high availability process but more as a disaster recovery process. It's not automatic, but it doesn't take long to do the switchover.

Now, this is great news. You can manually step in and get email back up and running with your passive copies. But what happens if the blackout ends and your primary site comes back online? What if the Internet connection is down and there is no way for them to know the others are serving as the active now? That could be a problem. They wake up and think "hum... guess the other servers are down" not realizing they were the ones that went down and the others are active. That's... again... bad. Split brain at the database level.

To resolve this problem a feature called Datacenter Activation Coordination Mode (DAC mode) was created. It's a property setting that you configure (through the Exchange Management Shell) when you create your DAG (or when you extend it to another datacenter or AD site). If you lose power and have to manually switchover to your secondary site and the power is restored to the primary datacenter while the WAN link is down,

DAC mode prevents the servers in the primary datacenter from mounting their databases even though they think they have a quorum.

Exchange 2013 DAG

While there were huge changes in high availability offerings between Exchange 2007 and 2010, the changes between 2010 and 2013 are more subtle. We still use DAG with 2013 and everything you just learned still applies. But they have continued to improve the solution.

One improvement is the ability to place your witness server in a third location. This would allow for a site failover in the event one site loses power because the other site would still be able to communicate with the witness (in the third site). It's still not a perfect solution because there are many ways the quorum could be disturbed by WAN connection outages and so forth.

In addition, there have been some improvements with regard to automated maintenance for lagged copies (which is connected to another new feature in 2013 called the Safety Net, which replaces the transport dumpster feature). Some minor improvements like automatic database reseeding after a storage failure and a few other tweaks have also been added.

To learn more about Exchange Server 2013 Database Availability Groups (with lots of graphics) check out this great article by Exchange MVP Paul Cunningham on his blog, ExchangeServerPro.com:
<http://exchangeserverpro.com/exchange-server-2013-database-availability-groups/>

Non-Mailbox Server Availability and Resiliency

I know it feels like the only important server in the world is the Mailbox server, but that's not true. Obviously you want to make sure your databases are accessible by Exchange so you put a lot of effort into that plan. Perhaps you have a backup/recovery solution you use and DAG for availability and so forth. But all of that comes to a screeching halt if you don't have other, non-mailbox server considerations in order.

First off, with legacy Exchange (2007 and 2010) the other server roles had to be considered. Now the Hub Transport was easy. To provide redundancy and availability all you had to do was add another HT server into your environment. The Edge Transport was/is similar in that you add a second one, copy over the settings (manually or clone it) and you're moving forward. The Client Access needed redundant servers set up (remember you can put server roles together so you didn't necessarily need lots of new servers, just servers with multiple role deployments of Exchange). With CAS servers you needed to provide the load balancing portion. You would do this, in most cases, with a hardware load balancer and you would put your CAS servers in an array.

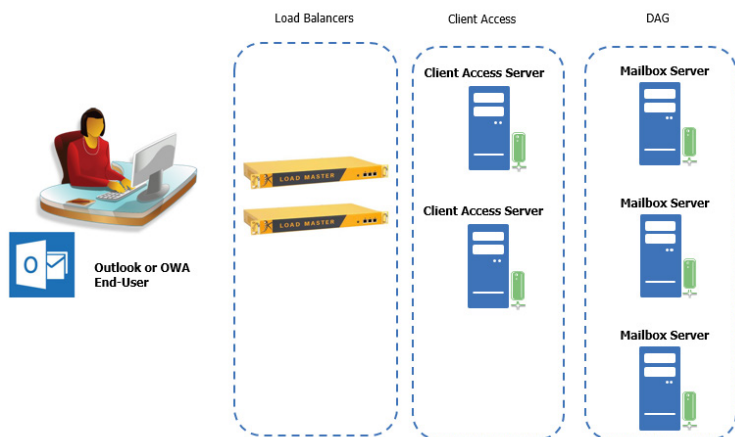
With the CAS role in Exchange 2013 remember the CAS handles ALL client connectivity (although it proxies or redirects back to the Mailbox server so it doesn't do any data rendering like its legacy version older sisters).

Because CAS no longer keeps any state or session data it doesn't need a Layer 7 load balancer, it can work with Layer 4. However, application aware load balancers (aka "smart" load balancers or Application Delivery Components) are the norm and are better in my opinion. Sure you can try DNS round robin in a lab (with no real load balancing) or you can play with Windows Network Load

Balancing (NLB), but for production environments you'll want a hardware or virtual load balancer.

Note: We don't discuss OSI model layers so when I use the terms Layer 4 and Layer 7 that may be a bit new. You can do some quick research on this but ultimately Layer 7 is smarter and higher up the chain (the application layer) as opposed to Layer 4 (the transport layer).

Here you can see an example of two load balancers (these are KEMP LoadMasters) so that we have redundancy of even our load balancers for higher availability. We have 2 CAS servers and 3 Mailbox servers configured as members of a DAG.



Aside from your Exchange servers you have to make sure you have redundant domain controllers, DNS servers, routers, switches, WAN connections, etc... And you may look to alternative, third-party solutions to provide some form of continuity. In the event your Exchange environment goes down (either on-premise or cloud-based with Office 365) it would be nice to have an alternative way to keep working (aka continuity).

The Big Takeaways


There is so much to learn from this chapter. It may require two or three reads to get it all straight. Possibly some additional research. But ultimately here is what you want to remember:

Exchange uses continuous replication (starting with Exchange 2007) to ship transaction logs over to a passive copy of the database (or databases with DAG) in order to provide a way to failover or switchover to another server if necessary. That server could be in the same site or it could be in another site, but the goal is to provide little or no downtime.

The solution has evolved over the years and in its current form it's called Database Availability Groups (DAG). You can have up to 16 members of a DAG. The configuration options and the design and deployment aspects of DAGs are limitless. It's worth researching DAG design options to see how these can truly help your organization obtain high availability and site resiliency.

In the end, however, DAG will keep your Exchange environment available for that point in time. But that won't help you if you need to reach back 5 years for discovery purposes (so don't forget the need for an archive) or if you need to restore something from a few months back that is no longer in deleted item recovery. In those cases an easily accessible archive or a backup solution will be needed.

Chapter 7: Unified Messaging



Next to Exchange high availability in terms of topics I love to talk about is Unified Messaging. People have so many misconceptions about the solution. Some are outright afraid of playing with it! I'm here to tell you it's not crazy hard to learn and it's not going to require a telephony overhaul... that's called LYNC! (one of the most difficult solutions I've ever played with so don't expect me to pen a Conversational Lync any time soon... I'll farm that out. Maybe a Baby Talk Lync if I were to write it. But I digress...

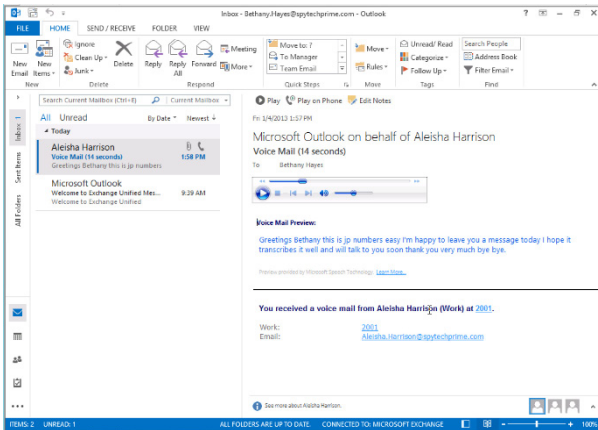
Unified Messaging (UM) includes services designed to provide a "universal" Inbox of email, voicemail and (if configured) incoming faxes. If implemented within your environment you'll have a variety of great new features added to your end-users'

world, services that they may already have in some form through another vendor but will now come through Exchange directly.

Unified Messaging 101

So, within your organization you probably have a PBX (or modern IP-PBX) that handles your incoming calls. Think of the PBX like a phone server. You may never touch it (or even see it) but if your company has 10, 100, 1000 people working in it you realize that there are telephone experts (in the field of telephony) who make that all “happen” for your company. Now, you probably have a voicemail part to your environment as well. So you have a company number, you get an extension... and if someone calls and goes to your extension and you don't pick up, they leave you a voicemail. Simple enough, right?

Ok, so again, the UM services provided in Exchange (starting with Exchange 2007) isn't looking for you to throw out your entire telephony infrastructure but rather for you to break off one little piece of it... the voicemail piece. With UM, if configured properly, you can have voicemails left for users and these will be placed in their Inbox as an MP3 (or some other audio format) and can even be transcribed in the email itself! How cool is that?!



Outlook 2013 with a Voicemail and Voicemail Preview

Note: With Exchange 2007 there was an incoming fax portion provided but with 2010/2013 you can configure this if you have a partner fax server solution with a URI provided by the fax solution provider. So, basically, Exchange 2007 could also receive fax calls, but 2010/2013 require you to use a third-party fax service to which it will send the incoming fax calls.

With Exchange 2007 there was a specific role called the Unified Messaging role. And this carried forward with Exchange 2010. You could install it with other server roles (not the Edge, but all the others). However, with Exchange 2013 going down to 2 roles, as you recall from Chapter 2, the UM services have been placed on the Mailbox role and they are installed automatically when you install the Mailbox role. Not to worry, if you don't use the UM services they won't interfere with the performance of your Mailbox server in any way so you don't have to worry about turning these services off.

Unified Messaging Features

There are a lot of cool features that come with the UM role. Let's take a look at a group:

- Outlook Voice Access (OVA): Users can call their Inbox and access their voicemail, email, calendar, and contacts (all read to them with text-to-speech) and they can update things (like their schedule) using voice.
- Voice Mail Preview: Uses speech-to-text to take a voicemail and put a text preview of it in your Inbox. It's not perfect but it uses a best-guess method for words it doesn't know.
- Incoming Fax: Can be configured after you establish a relationship with a fax vendor and then faxes will be sent to your Inbox as a .tif file.
- Call Answering Rules: Like Outlook rules for phone calls, these rules allow end-users to determine how they want calls to be handled.

- **Play on Phone:** Allows users to play their voicemails on a phone, rather than through the computer's speakers (for a bit more privacy)
- **Auto Attendant:** Lets Exchange answer the phone for a department or an entire company, with either default or company-specific prompts to help people navigate to the right person in your company, or be provided information you want to automate. You can have the user respond with voice or DTMF. DTMF stands for dual tone multi-frequency and yes, you can forget that immediately. Just remember, DTMF means using your keypad. If the auto attendant cannot understand you or if you simply prefer to use the keypad it's good to have that configured to use DTMF (or have an alternate DTMF auto attendant ready)
- **Language Packs:** Allows you to configure alternative languages for your UM services. Depending on the language (and if there is an available language pack) you can have auto attendants in the language of the caller configured, with voicemail preview transcription provided as well for some language.
- **Message Waiting Indicator:** Always good to know you have a message.
- **Missed Call/Voice Mail Notification Texts:** Again, good to have notification capabilities.

Making UM Work

First of all, I never recommend the IT admins or Exchange admins, who already have enough to do and learn, dive into the world of telephony. That is a respectable field in and of itself and experts already exist for it... so rely on them. Your expert or team of experts should be heavily involved in your configuration plans for UM integration.


However, it doesn't hurt to dabble a bit in another field right? So let's do a little dabbling and you'll see it won't kill you. To start

with there is the Public Switched Telephone Network (PSTN). Think of the PSTN like the Internet of phones; it's what lets you pick up a phone in Alabama and place a call to Zimbabwe just by dialing a number. Some phone company provides the lines coming into your company, including connecting them to the PSTN. Now you have to know that if you have 1000 people in your company they don't literally run 1000 phone lines into it. Instead they provide what are called "trunk lines". Depending on the discussed needs of the company a ratio is decided upon to ensure there are enough lines available when people pick up their phones to work. Obviously a call center will need more lines through the trunk lines than a normal business.

The lines come in and are configured to work through PBXs or IP-PBXs. These are devices that ultimately allow you to have a phone extension and be able to call the cubicle next to yours or outside to a pizza place for lunch. Again, we're not looking to break this - we LIKE pizza.

Now, if you have a legacy PBX in your environment you have to first check to see if it will work with UM. If not you will need to replace that (and you might as well go 21st century with the IP-PBX). If your legacy PBX will work with Exchange, you'll need to purchase a VoIP Gateway, which pretty much translates the legacy PBX communication method into an IP based method so that traffic can go over your network wires to your server. If you are already using an IP-PBX than you should be able to get it to work with the UM services if the type of IP-PBX is supported.

So, imagine this better through the simplified drawing here with an AudioCodes VoIP Gateway.



The Super Cool UM Lab

Ok, so I don't want you to freak out and think you have to go out and buy all that stuff to get UM up and running in a lab. In fact, when I work with it, demo it, etc... I use a simple lab setup using an AudioCodes MP-114 VoIP Gateway (and it's awesome).



2 cheap phones plugged into the FXS ports. Two lines out for FXO if I want to test that and a line for network connectivity. Now, the actual configuration of the gateway takes some learning (again, the telephony team is a must, but this is still dabbling). Step-by-step configuration is provided to help you get UM up

and running though, so you don't have to over-stress about the complexity.

Oh... I should mention, FXS (Foreign Exchange Station) System or Subscriber is an interface that drives a telephone, delivers battery, etc.. FXO (Foreign Exchange Office) connects to phone lines. PBXs have both FXO and FXS interfaces... and the telephony speak can stop here.

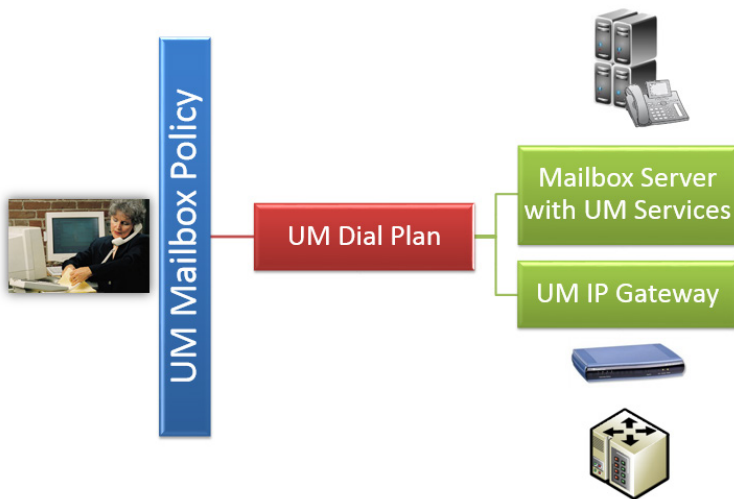
Configuring the UM Services

Remember, Exchange databases and such have been around for 20 years. Veteran Exchange admins will tell war stories about the ol' days in the trenches with it. But UM was just released in 2007 and many admins have still never played with it. So, it's one of those subjects that offers even footing... your own modern battle to boast about overcoming perhaps.

UM does have its complexities too, so don't let anyone fool you into thinking this is a solution you enable and it works. You will have to learn a bit about the configuration side to it, starting with dial plans. Dial plans, IP Gateways, Policies and Auto Attendants are all part of the fun of configuration that is Unified Messaging.

Remember, this is a primer, so we'll pitch this underhand and easy.

Let's start with the dial plan. A dial plan is a way to tell Exchange that there is a set of extensions that belong together and are, let's say, 4 digits. You can configure only one dial plan per user. So if the New York office has extensions already setup for their PBX of 4 digits starting with 2000, you might have one extension be 2001, and then 2002, 2003, etc... So your New York office will have a dial plan that specifies a 4-digit extension length. A user in New York and is UM enabled would have the New York dial plan configured for her user account, let's say she's assigned extension 2132. Now, each dial plan has to know where the real PBX or IP-PBX is at and so there is also a UM IP Gateway configured, which can use the IP address or FQDN of the VoIP Gateway or the IP-PBX for communications.

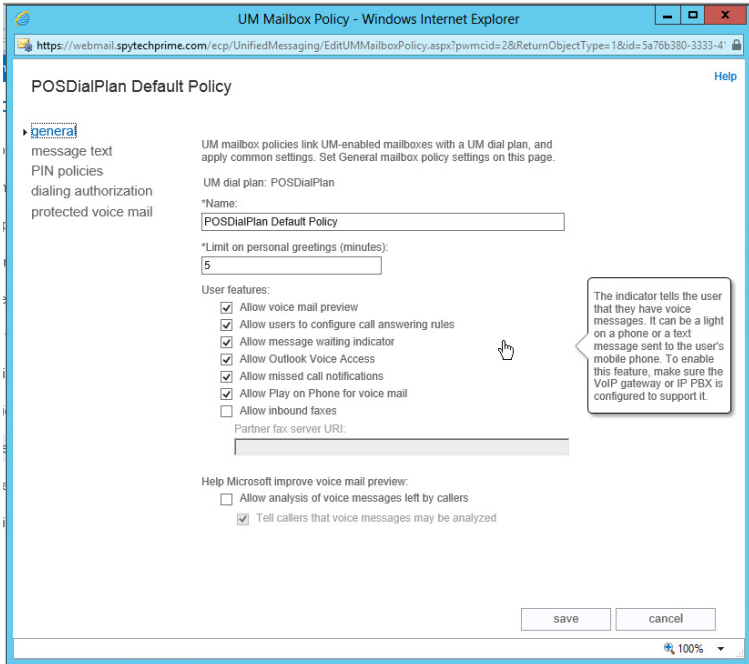


Ok, so let's take a step back. Once you have a good understanding of your existing telephony infrastructure and have your new VoIP Gateway ready (if necessary, unless using an IP-PBX) you go to your Exchange Admin Center and need to configure UM (which is already installed in Exchange 2013). So, what do you do next?

- First, create a UM dial plan to match the extensions set up in your PBX or IP-PBX
- Second, create and configure your UM IP Gateway(s) so Exchange knows where the PBXes are on the network
- Third, configure the UM Mailbox Policy so that user mailboxes get the right UM settings
- Fourth (optionally) create and configure the UM Auto Attendant(s) so that Exchange can answer the phone for your company or department

Once complete, you UM-enable end-users, associate a policy with their account and configure their PIN settings (unless set to be done automatically).

One thing to keep in mind is that the UM Mailbox Policy is the way you can enable/disable certain user features with UM. For example, you can turn Outlook Voice Access on/off, or Play on Phone, etc... It's here that you can configure the inbound fax server URI, a host of other configuration pieces.



UM Mailbox Policy in Exchange 2013



I'll be honest, it's a bit tough to be able to show you all the different screens and settings in all of this. It's too much to show you everything here, but if you are really interested in how it is all configured I recommend you check out my video training courses through Pluralsight. I demonstrate how to configure just about everything in Exchange, including Unified Messaging.

The Big Takeaways


Unified Messaging provides a “universal Inbox” for your end-users so that now they can receive voicemail right in their Inbox (and, if configured, incoming fax).

You can dabble in a lab environment with this but if you plan on deploying it for real, make sure you have telephony experts to assist with configuration of your PBX/IP-PBX connection to your Exchange environment.

The process for configuring UM on the server side is to create a dial plan and then UM IP Gateway. Configure the UM Mailbox Policy, perhaps establish auto attendants, enable end-users and start testing.

And remember, UM also allows for features like Outlook Voice Access (OVA) so you can check your email on your phone and alter your schedule if need be. Late for a meeting? Call OVA and let it be known. UM is worth testing and implementing. At least, in my opinion.

Chapter 8: Exchange Virtualization



Virtualization may be a new concept for you so we'll start this off by explaining what virtualization is and how it can benefit an Exchange (or any server) environment.

There is a good deal of debate about virtualizing Exchange (parts of it or all of it) but some of that is simply the inability of people to change and accept technology (yes, even if they work with technology).

“In my day, we didn’t have no virtual-zation mumbo jumbo! If you wanted to have multiple OS’s on one computer you used dual booting!!! Sure you could only run one OS at a time! But we liked it... we LOVED it!” a la SNL’s Dana Carvey.

Virtualization 101

Virtualization is a big buzz word these days. It’s more than a marketing term though. In short, if you think about your computer and all that hardware running, and then you think about your Operating System and how you have to install that on top of your hardware (on the bare metal so to speak) what are you really doing? You install the OS program on the hard drive right? Well when your computer boots up and looks at your hard drive it finds out where the OS is located and begins booting up your system. That OS is installed on the “bare metal”. But you know what? That makes life so much harder when you want to install an OS because you have to worry about the type of hardware, do you have all the drivers, what if the server outgrows the hardware, can you move it...? And that is where virtualization comes in.

Virtualized systems have a hypervisor. The hypervisor is a thin layer that sits between the hardware and the OS itself. So now you can install things much easier because it’s not directly on the bare metal.

Now initially we had what are called Type-2 hypervisors. These you installed right on your desktop and they allowed you to install another OS on top of your existing OS. So I remember running Windows Vista with VMware’s Workstation (a Type 2 hypervisor) running Server 2003 R2. I could perform testing and do demonstrations with the virtualized server I had. To the server itself it thought it was running right on bare metal. But in reality it was running on top of an OS with a hypervisor in-between. The negative here is that this made for such a huge performance hit. The better approach was to go with a Type 1 hypervisor.

A Type 1 hypervisor (like VMware’s ESXi or Microsoft’s Hyper-V) sits right on top of the bare metal. And the virtual machines

(or VMs) sit on top of that. Each virtual machine (VM) requires processor power and memory but most modern systems are so powerful that they are being underutilized these days. So it's pretty awesome that you can utilize them more fully by running different server types (ones that have solutions that cannot be installed on the same server so you need multiple servers anyway).

Now the hypervisor itself is not the cool part anymore. The cool part is the management solutions that are used on the back end to keep track of all your VMs and also assist with backing up the data and being prepared to help move VMs when necessary, or have them be redundant and fail over to another system if necessary (if one server crashes). These are features that modern virtualization management

Benefits to Virtualizing Exchange

We have to be honest here... Exchange runs best on the bare metal. Exchange MVP Clint Boessen summed it up best on his blog when he said "Exchange performs best when it can interact with the physical components of a server directly. If you disagree with this statement that's usually a symptom exhibited right after a VMware conference - hopefully it will go away."

So when we speak of "benefits" we are primarily speaking of benefits to your environment as a whole. For example, virtualizing Exchange servers brings server reduction (which is reduction in power needed to run those servers, reduction in heat generated and cooling necessary). It also brings server consolidation which is great for space savings and helps reduce underused servers in your environment. Most admins would agree that server management is easier too when your servers are virtualized. These are all good things. The negative is that there are complexities and costs that may be hidden at first. And admins need to have the skills to administer properly (which at this time, if they don't... they may fit the grumpy legacy admin at the outset of this chapter).

Exchange Virtualization

Exchange admins have been virtualizing Exchange from the moment the technology existed. At first they learned that it wasn't always best for their Mailbox servers (some still believe that). But the truth is, Microsoft didn't support it. We still did it ... but Microsoft didn't begin supporting virtualization of Exchange until August of 2008.

Now, you may be wondering which vendors are supported by Microsoft for virtualization of Exchange. Well, obviously Microsoft's own Hyper-V (that makes sense). And yes, VMware... Citrix... Red Hat... and a host of others who are part of the SVVP program (Server Virtualization Validation Program): <http://www.windowsservercatalog.com/svvp.aspx>

Which solution is the best one for virtualizing Exchange? I'll be honest, it doesn't matter. If it is supported, it's supported. I'm sure if you ran tests (and they have) and peered into the nano-seconds of performance response and such VMware would eke out a win over Hyper-V. But that's not really the deciding factor here. Which virtualization solution are you most comfortable working with? Which one do you already use? Keep using it, so long as it is supported.

One thing to keep in mind is that although Microsoft "supports" virtualization of Exchange, there are also lots of rules with regard to doing it. For example, Exchange 5.5 and 2000 are not supported at all in production in hardware virtualization environments. Exchange 2003 is supported (barely) and there are some key conditions you have to meet. And then things start to loosen up with Exchange 2007 forward.

There is a support document for legacy Exchange virtualization conditions you can find here: <http://tinyurl.com/4jtpg8>

There is a special Exchange 2013 Virtualization page you can find here: <http://tinyurl.com/d2wdx8r>

Best Practices for Exchange Virtualization

I liked this next quote from one of our reviewers, Phoummala Schmitt, in an article she wrote on the Petri site:



I've been virtualizing Exchange servers since the days of running Exchange 2003 on ESX 3.0, which at the time was not supported by Microsoft. I'm pretty confident that in those early days I was doing everything that the Microsoft support statement indicated not to do. Over the years, I've learn many lessons – some the hard way – on how to get the most out of running an Exchange virtual machine (VM).

Over the years speaking and writing about this subject you pick up a handful of best practices to help folks avoid making mistakes. Here are a few:

Don't Oversubscribe Memory: Some hypervisors have the ability to allow you to oversubscribe or dynamically adjust the amount of memory available to the guest VMs. While this may work for some workloads it doesn't work for Exchange because Exchange uses memory on an ongoing basis, so it won't release it and that causes problems. So dynamic memory allocation shouldn't be used with Exchange.

Avoid Virtualization Sprawl: Sometimes admins try to squish too many VMs on one server. With Exchange you are encouraged to provide the appropriate processor, RAM, storage and network connectivity that you would a physical server (a little extra wouldn't hurt). You don't want your Exchange server to be hurting for resources. A rule of thumb for virtual hosts is that they consume CPU overhead of 5-10% (but there is no absolute number here). Exchange does support a 2:1 processor to logical processor ratio (but even in this case it's better if you go 1:1).

Ensure Multiple Network Connections: VMs are sometimes configured on a server with all of them using the same network connection. That could be improved upon by a simple quad-port network card so you have port density for your virtualized servers.

Use Pass-through iSCSI: Although you have the option of going with .vhd files, you will experience better performance (especially with your Mailbox servers) if you go with pass-through storage. (Network-attached Storage or NAS is not supported). Also, dynamically expanding disks or disks that use some kind of differencing or delta mechanisms are not supported (disk size must be fixed). Oh... and snapshots are not supported (they aren't application aware yet).

Note: You can use the Exchange Server Role Requirements Calculator to size your storage properly. It's an Excel spreadsheet provided by the Exchange Team and it's an amazing tool. Here is the 2013 version (there are legacy versions too): <http://tinyurl.com/n5k73kq>

Live Migration and vMotion technologies (along with others of a similar nature) are supported for Exchange relocation for a planned migration of your VMs, however, failover activity at the hypervisor level just result in a cold boot when the VM is activated at the target. Hyper-V's Quick Migration is not supported.

Note: Microsoft really wants us to use DAG for our server failover and switchover needs with Exchange. Keep that in mind.

If you want to test your Exchange environment (virtualized or not) you can use tools like Jetstress and LoadGen to mimic real world stress on your environment to ensure you've designed it well. Jetstress tests the performance of the disk subsystem and LoadGen will simulate client connectivity (and perceived traffic loads).

Which Roles to Virtualize

Depending on which flavor of Exchange you are running, either 2007/2010 or 2013 this discussion changes a bit. The role that causes the most amount of heated debate is the Mailbox role.

So, a little story. I'm out speaking at the TEC Conference in Vegas in 2009. At some point I'm speaking about virtualizing the Mailbox roles and a man in the back of the room yells out "you can't virtualize the Mailbox role!" which was technically not accurate but he meant you shouldn't virtualize it. I thought "how rude... who does this guy think he is?" Well, he was an HP expert... and did I mention he was built like one of those Ultimate Fighter guys. I told him to pipe down... or we'd take this outside. (Ahem... not quite... I respectfully disagreed with him... in a very peaceful manner. He was quite scary).

Why the division? Two experts, two opinions. Well, his experience was driving his comments and at that time he knew that production mailboxes would do better on physical servers, which in some cases was true, but not all. Remember, support for virtualization of Exchange was only 1 year old and we only had Exchange 2007, so the difference of opinion was warranted. The critics felt/feel that the Mailbox role is so CPU and I/O intensive that virtualizing this role in a production environment is a mistake and performance will suffer.

As time has progressed and Exchange 2010 and now 2013 have arrived we see virtualization of server roles becoming more common, including the Mailbox role. Now, if you aren't convinced about that you might try virtualizing a member of your DAG for a passive copy, or for a secondary datacenter. So you can keep your production Mailbox server on physical hardware and go with virtualization for your DAG.

For Exchange 2007/2010 the Client Access and Hub Transport roles are supported for virtualization. What about the Client Access 2013 server? That is supported too. Just provide the proper resources for it to run properly. As for the Edge Transport, you CAN virtualize it but you would only do that if you planned on putting other servers in your perimeter on the

same box (otherwise why bother). Some worry about an escape attack if the hacker goes through the VM and there have been a few security exploits that may allow that to happen.

A virtual machine (VM) escape is an exploit where the attacker can run code to break through the virtualized server and interact with the hypervisor. This would be a very scary situation because the hacker could access other VMs and data.

As for the UM role being supported for virtualization it wasn't with Exchange 2007. Starting with Exchange 2010 SP1 it was supported but with odd requirements (tons of RAM and only stand-alone on the system, no multi-role install). With Exchange 2013 it is fully supported and baked into the Mailbox role.



I think that that the decision to virtualize Exchange is more complicated than whether or not it works and is supported. For example, the design of the VMware infrastructure is a factor too. I have been in situations where a large enough Exchange farm would warrant its own VMware farm. This is based upon the organizations design principles for VMware. In that case it is important to determine if the extra layer of complexity and support required for VMware is worth the benefits of putting Exchange on VMware. So when making these kinds of design decisions as to whether or not to use bare metal hardware/VMware (Hyper-V) or both be sure to consider this factor too.

Should You Virtualize Exchange?

There's no single right or wrong answer to that. Take stock of your environment and needs and consider ways virtualization can assist. Perhaps start small and only virtualize servers that aren't critical and see how they perform. Personally... I don't think I've installed Exchange on bare metal in over 5 years. Maybe longer.

The Big Takeaways

First up, never get into a heated debate with an Exchange geek who happens to also look like he could be an ultimate fighter.

Next, it's time to embrace the future because virtualization is a technology that is here to stay. It's one of the pillars to public and private cloud technologies and so we need to perhaps evolve a bit and compromise with regard to Exchange being installed in a virtualized environment. The biggest key is to ensure you provide Exchange with the same resources (CPU/RAM/etc...) that you would give the server if it were installed on bare metal.


Know what is and isn't supported based on your flavor of Exchange. Know the best practices for Exchange and virtualization: ie. use pass-through iSCSI, not fixed .vhds... no memory overcommit or dynamic memory... and so on.

Yes... it's true, running Exchange on bare metal is best. But can we gain other benefits by virtualizing Exchange while making the performance hit seem negligible from the end-user perspective? Absolutely.



With the right planning and good foundation for your infrastructure and storage, virtualizing Exchange can be a great option. I currently maintain a completely virtualized Exchange environment for a global company, across 2 datacenters using a stretched DAG. We have no physical Exchange servers in our environment that supports users in over 50 countries.

Chapter 9: Exchange Security



“This isn't what I meant by 'make sure our Exchange server is secure'”

I'll be honest, this one is going to be hard. Hard for me to explain and hard for you to understand. You see, security for Exchange data involves a lot of different pieces. One minute we'll be talking about certificates for your server, another we'll be discussing anti-spam/anti-virus or, if we want to jump to end-users we can discuss email encryption. It's a real hodgepodge and none of it is super easy to grasp. That's why I let it hit the TOC so late in the book.

Not to worry, by this point you are no longer newbies. Hey, if you made it through Chapters 1 through 8 you deserve to be here.

Anti-Spam and Anti-Malware Features

Microsoft knows how much junk is sent into Exchange every day. They could leave the junk mail protection to third-party companies (and for a while they did) but they realized some folks weren't doing anything to protect themselves and so they had to bake in some form of protection.

Anti-spam features have been evolving for a while in Exchange but the new anti-malware feature just arrived with Exchange 2013. Even though you have these features built-in, many will look for more help, either with an Edge server in the perimeter (partnered up with an anti-virus solution) or something in the cloud. Microsoft recommends their own Exchange Online Protection tool. I personally like Mimecast's protection because it comes with a bevy of additional features. You'll need to do your homework and decide for yourself.

Both the anti-spam and anti-malware features are super easy to explain. One spam or malware (viruses, etc...) become identified in the world at large and your anti-spam/malware tools are updated to recognize it (along with a billion other pieces of junk) they stand at the watch and divert (aka quarantine), delete or reject stuff that matches.

For example, Exchange has an anti-spam feature called Content Filtering. This was a feature formerly known in Exchange 2003 as Intelligent Message Filter and it can examine messages based on keywords, message size and so forth. It then gives the message a spam confidence level (SCL) from 0 to 9. The number is a gauge to indicate if a message appears likely to be spam (9) or not so likely (0) and everything in-between. Based on the SCL number you can take actions (Delete – which won't even notify the sender, Reject – where the sender is told, Quarantine – where it will be sent to an email address for analysis). The higher the SCL, the stronger the reaction.

There are a variety of anti-spam features in Exchange. All of these features try and do one thing... protect your organization from harmful junk. Sometimes it's just spam, sometimes more, but it's good to have help. Sometimes it isn't enough and you

need to look at doubling up on your efforts to keep spam/malware out.

Speaking of malware, Exchange 2013 also has a built-in malware filter. It's not super robust just yet (a version 1.0 product, although the Office 365 flavor is more of a 1.5) but it will detect malware and delete it and/or send alert text if you configure it to do so.

Role-Based Access Control (RBAC)

With Exchange 2010/2013 we evolved permissions in Exchange away from access control lists (ACLs) and moved toward roles.

The concept is simple in theory and the underlying permissions themselves are based upon something solid, PowerShell. Ultimately the way permissions are determined is through underlying cmdlets and parameters attached to the roles that are assigned to role groups. So the default roles have existing cmdlets attached that can be altered to make for enhanced roles or lesser roles depending on whether you add or remove cmdlets and/or parameters.

There are currently 12 built-in Role Groups. If you want to give someone permissions on a broad scale, just assign them into one of the default groups. For example, to allow a person to perform Discovery searches and place a person's mailbox on legal hold, you add them to the Discovery Management Role Group. If you want to give them control over the entire organization you add them to the Organization Management Role Group. So, there are 12 of these different groups you can utilize, including the following:

- Organization Management
- View-Only Organization Management
- Recipient Management
- UM Management
- Discovery Management
- Records Management
- Server Management

- Help Desk
- Hygiene Management
- Compliance Management
- Public Folder Management
- Delegated Setup

Each Role Group has Roles assigned to them to break it down further. There are 67 different Roles. For example, the Discovery Management Role Group mentioned a moment ago has 2 Roles assigned, the Mailbox Search role and the Legal Hold role. Those roles have entries that are based upon PowerShell. There are cmdlets and parameters that are assigned to each role to allow a person who is assigned to a Role or Role Group the ability to use the EAC to perform the tasks that utilize PowerShell behind the scenes through cmdlets and parameters seeing as how all Exchange management eventually ends up with a PowerShell command being run.

By default, the Exchange Administrator is made a member of the Organization Management role group and that has nearly all the roles assigned. In smaller organizations you might have one or two IT administrators handling the Exchange environment and so they might both be in the Organization Management role group and will be capable of performing all tasks. However, if your organization is mid-to-large in size you might begin delegating others to the various role groups so that they can take some of the load off your plate. However, in some cases the built-in role groups may not suffice. You may need to create specific role groups using roles of your choosing, and that is also possible.

You'll find RBAC to be quite simple if you stay within the default role groups and roles and administrate the process through the EAC, however, it's obvious that this can become much more involved when you start looking into more granular control through the EMS.

Now... here is where it becomes interesting. While it is relatively easy to find the Role Groups and definitions for each, and it is even easy to locate the Roles with explanations for these as well,

it becomes a bit of a challenge to locate the entries that go along with those roles.

Now if you are reading this and thinking “why would one need to get that involved in the process?” I’ll tell you. While the default Role Groups and Roles are great and certainly more extensive than anything we’ve had in legacy Exchange permission options they are designed to be flexible and allow for the ultimate in granular permission settings. You can create your own Roles and Role Groups (based off of those Roles). Often times the way this is done is by using an existing Role as your parent and having a child Role strip out whatever permissions you need to change. The one caveat here is that you cannot have a child role have more permissions than the parent.

So... if you don’t know what permissions you are starting with (ie. the cmdlets and parameters themselves) and only have some foggy explanation about the Role and what it does, you may have difficulty creating the new roles. You need that information and it has to be easier to get at than it currently is... which is through various PowerShell commands seeking out the management role entries of the roles.

There are several tools that can help but I like the free CodePlex tool RBAC Manager: <http://rbac.codeplex.com>



There is so much more that can be said about RBAC but not without going much deeper. The best chapter to read, in my opinion, on the subject is Chapter 12 of the Sybex book *Mastering Exchange Server 2013*. Really well done. Aside from that you can jump on Pluralsight and watch some of my video lessons about it which are equally riveting!

Certificates

Talk about a discussion that could be its own chapter. Certificates are part of real-world Exchange, not just lab-world

Exchange. The book “Mastering Exchange Server 2013” said “We think that most people, they don’t understand what certificates really are or how they work. Certificates and PKIs are ... “stark naked voodoo” mainly because they’ve traditionally been complicated to deploy and play with”. I thought that was a funny line but ultimately the author was trying to say most folks place concept of certificates in the same category as the dark arts.

We’ll steer clear of the difficult parts (no talk of X.509 cert standards or anything weird). Let’s just try and explain what a certificate does. So your clients are connecting to your servers and they may be using protocols like HTTP, SMTP, POP and IMAP to do it. It’s important to secure communications from server to server and from client to server. Secure Socket Layers (SSL) is used for securing communications (one of the methods). By default, client communications use SSL for encryption for Outlook Web App, ActiveSync and Outlook Anywhere. And SSL requires a digital certificate.

A certificate is like a verification card. A way to authenticate that the holder is truly who they claim to be. One Exchange MVP, Lasse Pettersson, likes to compare a certificate to a passport and a Certificate Authority (CA) to a global passport agency that is trusted by everyone. There are three different types that can be used with Exchange including:

- Self-signed: These are automatically created and used by Exchange the moment you install it. They allow Exchange to work out of the box and they are fine for lab environments, but they aren’t meant for production Exchange. Imagine a person approaching you and saying “I’m trustworthy, you can use my services. Here is my card... that I wrote and signed myself.” That may not impress you as much as if it was signed by someone you knew and respected. So, the self-signed are temporary until you obtain the appropriate SSL cert.
- Windows PKI-generated: You can set up your own Windows Server with Certificate Services and obtain a PKI cert through your own organization. So, if you are comfortable running your own in-house certificate

authority (CA) this is a viable option. However, many find that the low cost of option 3 and the ease of deployment make it the better choice.

- Trusted third-party certs: These are purchased from a trusted certificate authority (CA) for reasonable prices. These certs, when provided by a known CA, are automatically trusted by client computers and mobile devices. If your organization is allowing external access to Outlook Web App, ActiveSync to mobile devices or Outlook Anywhere 3rd party certificates are the best option.

Note: With CAS and Mailbox servers residing on separate servers you only have to worry about changing the self-signed certs on the CAS servers because the Mailbox server doesn't accept direct connection from clients. However, with multi-role servers you have to change the certificate.

Now, one of the most important tasks with a certificate is getting all your names registered. Some of these are known by clients (like mail.companyname.com or something for each of your offered services: OWA, OA, ActiveSync) and others may be used behind the scenes like the server FQDN or Autodiscover services (like autodiscover.companyname.com). You may have legacy Exchange servers in your environment and may need a legacy.companyname.com name registered.

Because we are looking at registering so many names you want to obtain a Subject Alternative Name (SAN) aka Unified Communications (UC) certificate. These SAN/UC certificates allow you to pay for one certificate (rather than multiple certs) and add all your server names and external URLs to it.

Another option with certificates is for you to purchase a wildcard certificate. Something like *.yourcompanydomain.com so that it covers all the subdomain naming you can come up with. Although wildcard certs work, many are not comfortable with the security implications of having a certificate that can be used open-ended. So SAN/UC certs are preferred because they are created

specifically for the names you provide and are thus considered more secure.

Once you have all your planning done you are going to follow the steps provided by Microsoft to generate a certificate request specifically for Exchange 2013 and then use that request to obtain your cert through your provider. I use GoDaddy, others use DigiCert or some other provider. The choice is yours really. Once you get the certificate you will import it on your server(s) and assign it to services.

See... not as scary as they make it out to be right? No more difficult to grasp than high availability or Unified Messaging right?

Transport Layer Security (TLS)

Email flows internally in a secured environment. TLS is used with internal communications and it's the latest version of the SSL protocol.

At times you also have partners that you connect Exchange to using send connectors. You can configure mutual TLS authentication to provide session-based encryption and authentication. With mutual TLS each server validates the other server's certificate as opposed to TLS where no authentication is performed or sometimes one side authenticates.

When sending email to another organization that doesn't have TLS the email message will not be encrypted. TLS will only work if both the sender and receiver have it enabled with the mail systems.

Client-Side Protection Concepts

Security is more than server-side, its client-side too. Some of the best security you can provide for your organization may be the result, not of technology, but of training.

Some additional client-side options include things like dual factor authentication. This would be something that protects user access to the domain, which would enhance security.

The use of an email encryption solution may be worth considering. S/MIME is a client-side technology that provides signed or encrypted messages, however because of the nature of S/MIME messages they may not fall within your company's policy lines because they cannot be scanned, cannot have disclaimers applied, cannot be inspected, etc.. so that may not work for your organization.

We discussed the use of IRM earlier on when discussing Regulatory Compliance. It's not a perfect technology but it does offer a few deterrents that add to your security for clients.




One interesting set of training is from KnowBe4. They partnered up with Kevin Mitnick, notorious social engineer, and created training that should scare end-users and train them properly not to click links to bank accounts and things of that nature. They also have mock email tests that allow administrator to see if a person has benefited from the training or not.

The Big Takeaways

This chapter is a cornucopia of different security considerations. Anti-spam/anti-malware options, permissions, certificates and so forth. And we just scratched the surface.

Obviously there is more you can do to secure Exchange. Hardening Exchange (as they call it) may include firewall and port usage hardening. Certainly making sure all your servers are up to date and patched is key. If you are using Hyper-V for your virtualization you want to consider using Server Core for your parent. So there are lots of additional security options to consider. But hopefully you have a better idea now of what is involved to protect your mission critical email solution.

Chapter 10: Office 365 (Exchange Online)



*"Come on J. ... let's pick one,
it's time to embrace the future!"*

What is Office 365? Well... it's a confusing name for a great solution. Its predecessor had an even worse name: Business Productivity Online Suite (or BPOS for short).

The reason Office 365 is confusing is because many folks think it is referring to the next flavor of Office, and to a degree they are correct (I'll explain that). But the primary offering is actually Microsoft's hosted versions of Exchange (Exchange Online), SharePoint and Lync.

Let's break down what Office 365 is all about.

Clearing Up the O365 Confusion

As mentioned, Office 365 is partially all about the hosted services you can obtain by choosing a package that fits your needs. At the same time it's also about subscription Office (if you pick a plan that includes the Office suite).

There are three “service family” plans: Small Business (up to 25 users) Midsize (up to 300 users) and Enterprise (over 250 users). Even if you have a small business of 10 people you can choose an Enterprise plan if it has the features you need/want.

Every plan you choose has a base of services (they all include the Office 365 Platform, all include Exchange Online, all include SharePoint Online, almost all include Lync Online, all include Office Web Apps) and then vary with add-on services like Project Online, Yammer Online or the Office applications subscriptions.

Logically, the plan you choose will have a price tag attached and this will often drive the decision on which plan is best for you. You want to be careful that the plan you choose includes features you want. For example, if you get a small business plan you may not have some of the regulatory compliance features you would like to have (like premium journaling). You can always upgrade your plan if you need to but it would be better to know up-front what your plan supports. These plans are not just based on number of seats, they have enabled/disabled features to consider and some include Office while others do not.

So, you might be thinking “Ok, so if I go with Office 365 I get Exchange Online right? Exchange 2013?” The answer is yes and no. Initially, when they upgraded the BPOS platform to Exchange 2013 you would have gotten that flavor of Exchange. But one of the coolest things about Office 365 and Exchange Online is that they are always making improvements to it. And you don't have to wait for a cumulative update or a service pack to see the improvement or new feature. It's online first! So you don't get Exchange 2013... you get the latest flavor of Exchange available, Exchange 365. Note: Eventually many of the online features will be provided to the on-prem edition through updates and service packs.

So the Office 365 flavor can be more capable than the on-prem version in some cases. Case in point is the anti-spam features. When Exchange 2013 RTM'd you could only administer anti-spam through the Exchange Management Shell. No EAC option. But with Exchange Online (through Office 365) I can see now that some of the anti-spam features are in the GUI. So I get the latest interface in the cloud version of Exchange.

The same is true of your Office applications. You can still buy Office 2013 and install it directly on a desktop for a user. But if you buy the subscription (with your Office 365 plan) your user's Office products will update to the latest features and such immediately.

Again, Office 365 gets all the enhancements first, and in some cases may be the only platform to get enhancements. There is no guarantee that a feature will come down the pipe to your on-prem version.

Hosted or Cloud-Based Exchange

I may have clarified what Office 365 is but not what hosted Exchange is, or Exchange Online specifically.

Hosted Exchange isn't a new concept. Providers years back said 'hey, we can set up Exchange for you and give your company accounts with their domain name (just point the MX records to us) and you can have Exchange without the stress!' It's a great idea really and one that smaller businesses (and mid-level too) have appreciated. But those earlier, multi-tenant deployments came with a lot of limitations. As an Exchange admin I couldn't get under the hood and make any real changes.

Modern hosted Exchange providers have been evolving so that they provide higher end services at a reasonable price in order to try and compete with Microsoft's Office 365. Another vendor trying to compete with it is Google Apps, which offers hosted email and services as well, but I think the pendulum has swung back in Microsoft's favor on that. Google Apps is yesterday, Office 365 is today.

In addition to hosted Exchange you can go with a dedicated virtual server that has a full version of Exchange on it. So in that case you have more control over the Exchange environment but don't have to worry about the hardware it is running on.

Depending on the organization you work with (healthcare, finance, government) hosted email or Exchange may not be an option for you. You may need on-premise Exchange. But, if that isn't a concern and you are looking to go toward a hosted or cloud-based solution you need to do your homework and choose one that works best for you needs.

Hybrid On-Premise/Office 365

What some companies are doing is mixing the two options together. Because Microsoft built Exchange and offers O365 they have made it easier for the two to work together. Some call it the best of both worlds. The organization can keep mailboxes in-house that are of a more sensitive nature while allowing Office 365 to handle non-critical mailboxes (like temporary workers perhaps). Or they can use the archive features of Office 365 combined with on-premise mailboxes.

With the hybrid model users can find each other across platforms through a common global address list (GAL) and can share calendar information (aka free/busy data). Exchange admins can use the same Exchange Admin Center tool to administrate both, which makes it convenient as well.

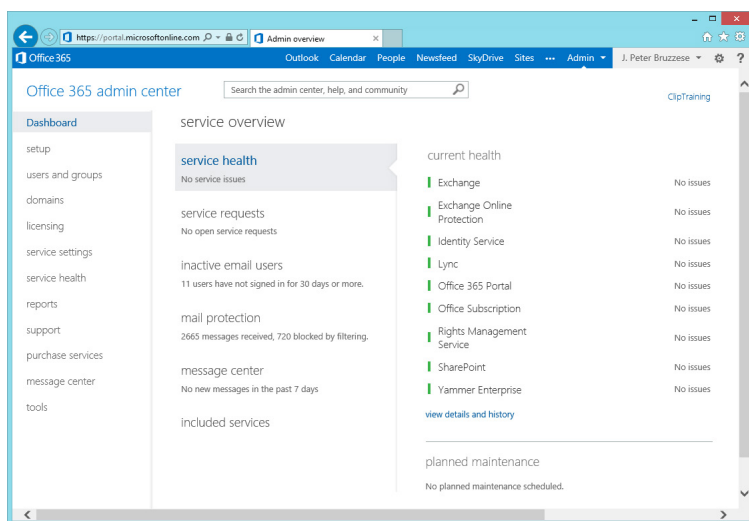
A Tour of Office 365

I personally love working with Office 365. When I log into the Office 365 admin center I'm greeted with an overview of services (as you can see from the figure). I can see immediately if there are any issues with my services and see if there are health issues.

I can easily add new users to my portal, pull up reports and more. It's very easy. But that isn't my favorite part. My favorite part is that even though this is a hosted, cloud-based deployment of Exchange for my organization, I can still select the Admin link in the top right corner and choose the Exchange administration

option. This brings to an (almost) fully functional Exchange Admin Center. I say ‘almost’ because you cannot configure server hardware (databases and such).

It’s really awesome to be able to control and configure Exchange Online using the same tools I’m used to using on-premise. Often times with hosted solutions it doesn’t work that way. You get some kind of proprietary tool set (web-based) that gives you very limited options. But with Office 365 you get a very robust administration experience. As close to on-premise as you can hope for with a hosted solution in my opinion.



Office 365 Admin Center

Note: Office 365 can be managed through the EAC but it can also be managed through a remote PowerShell session.

The Big Takeaways

Office 365 is Microsoft's hosted suite of communication and collaboration solutions including Exchange Online, SharePoint Online, Lync Online and several other options depending on the plan you choose.

There are a variety of plans to choose from with different features and price tags attached. You need to make sure the plan you choose is best for your needs.

Some plans come with a subscription to Office so that users can install the latest version of Office applications. One of the values to Office 365 is that all of the solutions (the server-side ones and end-user ones) are kept up to date and are the latest iterations of those solutions available. So even if you have an on-premise Exchange 2013 server with CU3 or SP1 installed, the online O365 version of Exchange (aka Exchange Online) will still be more current. Same with your Office apps.

With Office 365 you can perform hybrid configurations of Exchange, meaning you can have a portion of your Exchange environment be on-premise and another portion be in the cloud, with O365.

There are alternative hosting options too. Alternative providers with different types of cloud-based Exchange or non-Exchange email offerings. You may find, with some research, that these better fit your needs. Perhaps a price point you prefer based on the services offered. Perhaps solutions that add greater value for a competitive price.

Parlez-vous Exchange?

Do you speak Exchange? Yes... yes you do speak Exchange. If you have read through the past 10 chapters then you can sincerely say you have a good grasp of conversational Exchange, including its online younger brother Office 365.

Through this book we have addressed the primary terminology and concepts behind Exchange (past and present) but there is so much more to learn. There are books that are 1000 pages on Exchange and even that isn't enough. We haven't even discussed configuration (step-by-step) or design/deployment, migration strategies, monitoring options, PowerShell (a monster subject in and of itself)... but that's ok. This book was meant to establish a base level of communication. "Conversational Exchange", not Fluency or Native Exchange just yet.

Perhaps as you read the chapters you did some research on the subject matter and that helped you round out the concepts. Or perhaps you have watched some of the videos available (I have a ton of them with Pluralsight on Exchange 2010 and 2013). Visual learners will no doubt appreciate seeing things done with Exchange and taking their knowledge to the next level.

For more insight you can sign up for videos at Pluralsight.com

You can visit my blog: ExclusivelyExchange.com or follow me on Twitter [@JPBruzzese](https://twitter.com/JPBruzzese)


Or maybe you are a hands-on kind of learner. You need to do it, install it yourself, immerse yourself in it, fail a few times with something and plough through so that you own it. And if you are that person you know there is a great deal ahead of you. But it's possible one day you may be an Exchange admin, perhaps even Microsoft Certified (the 2013 tests are quite difficult). And hopefully you'll look back and remember where you got your start. Right here with Conversational Exchange. And should I need a job at that time... (kidding)

We did our best to make this book easy to read. Nevertheless I know it wasn't "easy" per se. I had Exchange MVPs and Exchange gurus reading it, but I also had newbies, friends and family (including my own wife, who suffers enough listening to me talk about it day-to-day). They all contributed to the effort to make this book readable and I'd like to thank them for their help. And I'd like to congratulate you for making it through.

J. Peter Bruzzese



Appendix: Basic Exchange Prerequisite Knowledge



“Sorry sir... you can't go in here until you've mastered the basics of networking, DNS and AD.”

Folks, the last thing I wanted in the very first chapter of this book was to confuse you or scare you off. So I saved some of this information for the end of the book. It's important that you grasp a few basic underlying networking concepts before you dive right into the world of Exchange. If you peruse this information and feel you know it already, skip it! If you read the book and you don't want to learn more, skip it! If you don't think this information is part of what you need to learn Exchange, skip it! But... if you don't understand the basics of networking, or TCP/IP, or Active Directory... you might just want to keep reading.

How a Network Works

If you understand the way a small network works you will also understand, to some degree, how large networks operate. You may have a basic grasp of networking from your home network, where you know you pay for a connection to allow your home systems and devices to connect to the Internet. However, you can set up a home router that allows your devices to communicate with each other while not being completely exposed to the entire world. Your home network might have a WiFi enabled router with some systems plugged into it directly (or perhaps you have a Sonos bridge or a Hue bridge connected directly in) and then you have devices connecting through your in-house WiFi.

Let's dive just a bit deeper into the physical side to a network.

The Physical Pieces to a Network

Most home networks are designed not just to connect computers to each other or printers; rather, they are designed to link to the Internet connection coming into the home.

The incoming connection might be a DSL line or cable modem or satellite, depending on your local providers. Hopefully no one reading this book is still dealing with dial up. Now the Internet providers usually set up the connection to one computer in your home. Their little box has an Ethernet connection that uses a cable to connect to your computer's network port.

This cable is called a Category 5 Ethernet cable. Why Category 5? Well, as you might expect there were earlier categories, 1 through 4, which are not used anymore. The future categories are: you guessed it, 6 and 7; these are new to the Ethernet cable scene. Cat 6 is used for Gigabit Ethernet and is backward compatible with 5.

Some of the terms you might see with Ethernet cables include: 10BaseT, 100BaseT and 1000BaseT. These indicate the amount of data the cable can transmit per second, either 10 Megabits (not

bytes but bits), 100 Megabits or 1000 Megabits (often referred to as Gigabit speed).

Why is it called Ethernet? Ethernet was developed in the early 1970's at Xerox PARC by Robert Metcalfe and others. The reason it was called Ether-net was based on the concept of luminous "ether" which was once thought to carry electromagnetic waves through space. At that time, many networking systems were proprietary (that is, unique to a given environment) and the idea was to indicate that the Ethernet wasn't just for one type of system but for all systems.

The cables you might use in a home network are easy to distinguish from your phone cables but they do have some things in common. For example, the connectors look similar. If you look at the end of a phone wire you see a little head with a clip. If you look closer you'll see that there are copper looking pins inside. That is an RJ-11 connector. Now if you look at the end of an Ethernet cable you'll see that it is slightly bigger and has more pins, eight to be exact. That is called an RJ-45 connector.



RJ11

RJ45

Most wired networks are going to use the Category 5 Ethernet cables with RJ-45 connectors on the end. One end plugs into the back of your computer, either into the motherboard itself or into a network card. The other end plugs directly into the cable modem or FiOS that is provided by your Internet Service Provider (ISP). Your provider (or ISP) may be the same company that provides your cable television and/or home phone. However, to increase the use of that Internet connection

toward other computers within your home you will need a special device called a “router”.

A router is like a post office. Communication between the Internet, your home network and computers within your home network is handled through ‘packets’. These are like pieces of mail that travel from one home to the next. If you want to mail something officially to your neighbor you would take your mail to the post office and put it in the Local box. If you want to mail it to another state or country you would put it in the Out-of-Town box. The post office would handle it from that point.

Your router will send packets from one computer to another computer; and from the Internet to your computers. One thing to note is that when your router is simply connecting computers in your home it is actually acting as a switch, not a true router. That little point isn’t meant to confuse you but to help you when you decide to purchase a router at some point because you may want a router that also has a 4-port switch.

What about wireless? Well, most routers have the ability to be wireless access points. This allows your systems with built-in wireless connectivity like: most modern laptops, iPads, other tablet systems, and desktops with wireless cards installed or USB wireless connection, to access the router, each other and the Internet.

Whether wired or wireless, how do these devices actually talk to each other? How does the router know to which computer to send information?

MAC Addresses and TCP/IP

All devices that are on a network, on the Internet have a built-in MAC address (Media Access Control). These are typically assigned by the manufacturer of the device and are assigned using hexadecimal numbers, for example: 00-21-6A-3E-D5-5E.

Note: You can easily find a PC computer’s MAC address by opening a Command Prompt (click the Start orb and type cmd)

and then within the command prompt type in either **getmac** or **ipconfig /all**

We now have a clear way of seeing that every device is unique and that is great because it helps prevent confusion. Nonetheless, the numbers are not that easy to work with; and there is no order to them amongst devices. Let's say in your house you have a couple of different computers and some Wi-Fi enabled devices, like a tablet PC or mobile device or eReader (Kindle); you wouldn't want to write down and remember the Mac address of each device to communicate, would you?


You might be thinking, doesn't my computer also have a name? Can we use these names instead of the Mac numbers? You can use the name on your local area network (LAN), which is the network in your home. (Note: A LAN is different from a wide area network or WAN, which you don't have to worry about for your home). Then again the names don't really help your network keep track of all the devices and Mac addresses for easy communication. Instead all the devices on your home network and on the Internet use a TCP/IP Address. The computers locate each other using TCP/IP addresses and only then does it acquire the true MAC address of a computer to link for communication.

TCP/IP stands for transmission control protocol/Internet protocol and it is actually more than two protocols but a whole suite of them. Now, you may be wondering what we mean by protocols. Well, protocols are sometimes used to mean a language or a set of standards. Having standards causes the different manufacturers to follow a protocol when developing things that will work together on a network. Think of trying to follow a recipe when everyone has a different size teaspoon and cup measurement. It would never work. The standard measurement allows everyone to cook the same meal in much the same way TCP/IP has standards or protocols.

Going back to our post office analogy, in much the same way you might follow certain standards when boxing packages you send to someone; TCP/IP has a standard set for packets that go out on the wire and framed or boxed properly. You also need to make

sure that the packages are addressed properly. The same is true with TCP/IP. To send a simple document from one computer to another, even on the same network, the document needs to be broken up into packages and then sent over the wire to the other computer. The two computers might both be plugged into, or connected wirelessly, to the router. Yet, like a post office needs an address to locate the recipient, the router uses the TCP/IP address to locate or route its packages to its recipient.

Note the following graphic. It's meant to illustrate how TCP/IP uses IP addressing to deliver email in much the same way we have an address system that works for real mail.



What have we evaluated so far? A local area network (LAN) uses Ethernet cabling with RJ-45 connectors to connect computers to routers or the wireless router wirelessly; and the router helps to make sure packets of data get from one system to another. TCP/IP is the set of standards (or protocols) that make it possible for this communication to take place.

TCP/IP addressing also makes it much easier to bridge the gap between the MAC address and the computer or device. What is great about TCP/IP is that even though it is all numbers, it is a lot easier to work with and organize than MAC addresses.

How the Internet Works

You probably have some kind of Internet Service Provider like Comcast, AT&T, Brighthouse, CenturyLink or one of the main ones available. Thus, you have a connection coming into your home. That connection might allow you to plug in one computer or you might connect it to a router. Internally you may have a local network connected off that router with IP addresses that you have chosen. However, the Internet uses IP addresses that are given out specifically for use on the global network.

The router has an internal IP address, which is also called the default gateway once you configure your in-house computers to access the Internet. The router also has an external IP address, which connects it to the ISPs network. The router basically transfers data back and forth between the ISPs network and your internal home network.

Now when you open up your Internet Browser, (maybe you like Internet Explorer, maybe Firefox, maybe Google Chrome, maybe Safari or some other option) you type in the URL to the website you are looking to access. URL stands for Uniform Resource Locator which is a fancy way of saying website address.

The URL is made up of the protocol you want to use followed by a colon and two slashes, like: **http://**

Then you add the path to be able to locate the web site. That web site is being hosted on a server, or group of servers, and to access it you need to know the IP address of the server or server group. But, how can you know the IP addresses of every web site in the world? Well, we don't have to know every IP address, we just type in `www.microsoft.com` or whoever we are looking to reach. That path helps us to find the site through the use of Domain Name Service (DNS) servers.

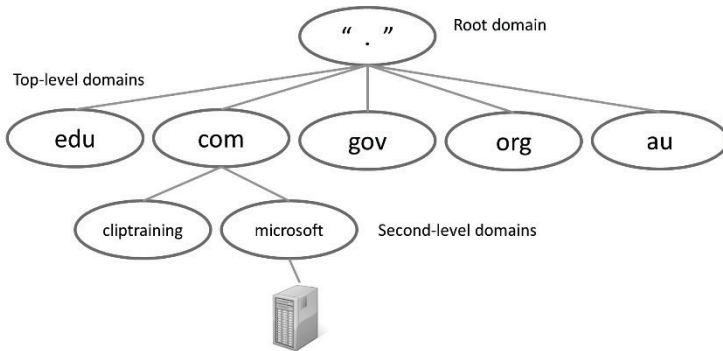
Domain Name Services (DNS)

It's simple. You need to call a plumber but don't remember the number, look it up in the phone directory. The directory is in

alphabetical order so you can find who or what you need by subject.

DNS Services are servers that are on the Internet to help us find the IP addresses of web sites or send emails to mail servers and so on. These servers are organized by domain, just like an alphabetical phone directory.

The root for the whole DNS system is a period (.) which is odd because we never type that in. If we typed a period (.) it would be at the end of the URL, because it is assumed we would leave it out. Instead we end our URLs with .com, .gov, .net, .org and so on. For countries, there may be ones like .uk, .cn and so on. That is why not all URLs we type in are .com, but can include other ending points.




The servers are registered under the domain that is chosen by the person who is setting up their domain name. For example, you might go to a popular domain name registration site like GoDaddy.com; and there is usually a domain finding dialog box for you to check to see if your domain name you want to register is already taken. Keep in mind you don't ever have to register a domain name. However, if you want to have your own web page for personal or business use you will need to register the domain name first.

When a domain is registered, the Domain Name Server (DNS) keeps a record of where the location of the web server is for that domain; and for other servers, if you have them, like email

servers and so on. When you type in something like `www.microsoft.com` your computer has no idea where the web servers are for Microsoft. Instead it very quickly checks in with a DNS server. That DNS server looks to see which category is needed. It sees `.com` and says okay; let's check to see if Microsoft is registered. It finds Microsoft and finds the DNS servers that are configured to tell you where the `www` services are hosted. An IP address is provided back to the computer to tell it how to find the servers for the web site. Now your browser knows the IP address on the Internet to connect to that `www` service and enables viewing of the web page. Remember, it already knows to use the HTTP protocol because that is in the beginning of the URL. It also knows to use Port 80 to make this connection to the channel on the web server hosting the pages. And after that expeditious process; viola... the page appears before you.

When you want to access a more secure site you type in **https://** which requests a secure page type through Secure Socket Layers. Without going into great detail, SSL sites are more secure for your banking, purchasing and other secure transaction needs. The SSL connections are usually shown in your browser with a lock graphic of some sort to let you know it is safe to proceed. By default, SSL uses port 443 rather than port 80 for this connection.



DNS provides this hierarchical, organized set of registered IP addresses and domain names for everyone on the Internet and it is all behind the scenes. This is also true for other services like email. If you type in `bgates@microsoft.com` the DNS servers are able to locate the email server for the email to be sent to the correct server. The DNS server uses the same pattern we explained above except instead of the IP address hosting `www` services it provides the IP address for the server hosting the email services. It does this because the DNS Server has MX records configured. So when you type `www` the DNS server responds

with the IP address of the web server. If you type in an email address, the DNS server provides the IP addresses configured as MX records for your organization. So can see that DNS is essential to email and Exchange.

Exchange Server requires DNS

The Network around You

What an amazing thing if you work for a company that has a network with cables, routers, switches and more. Do you even realize what a tremendous learning experience is right in front of you? Too often though people go to work, sit down, log in, work, and log out at the end of the day. Never wondering what makes it all happen.

If you want to start learning more take a look BEHIND your computer. See the cables? Where do they go? Do you have a false floor that allows cables to be hidden? Or are they all just out in the open? How many servers does your company have? Why not ask your IT admin or network team about that. You probably log in and connect to servers for file saving purposes right? And your computer connects to a printer, to email services and more. Are your email services on-premise (located somewhere in the building) or are they hosted or cloud-based?

Active Directory

Active Directory (AD) is an identity management system and directory service. When you log into your work domain you need a username and password. The identity management system confirms that you are who you say you are and provides you with the ability to log in and access resources on the network (files, printers, etc...). At the same time your name is in Active Directory, which can be used to provide your address, phone numbers, position in the company and a host of other important details that can be searchable as a result of the directory service. There are many different directory services, but Active Directory is the one that Microsoft has created.

Servers Make the World Go Round

In a network you have client desktops, laptops devices and you have servers that provide services. Servers sound scary when the reality is that they are literally there to “serve” you, so don’t be too worried about them.

It’s good to have an idea of what different types of servers there are and what they do. I’m going to list out a few (not all) and you’ll note that I have a Microsoft slant here, not necessarily because I’m partial but because these are the ones I work with.

Server Type	Description
Active Directory	Provides a way for workstations (desktop systems) to log in. It maintains username/passwords for the people in your organization and provides your system a security access token when you log in. It also maintains directory information about persons (if you input that into the system, like address/phone/etc.) and offers a variety of tools for management of your network
DNS/DHCP	These are services provided within a network and could be included with other server services (like your AD server). As you recall, DNS provides name services (although in this case we mean internally on your network) and DHCP provides IP address leases for your client systems.
File (services)	These servers are designed to allow client connectivity so that persons can save their files to the network server. This is good for two reasons: It’s easier to back up the one file server rather than 100 clients, and it’s

	easier for collaboration when documents are on a network share.
--	---

Server Type	Description
Print (services)	Allows you to connect one or more printers up to a server and allow users to print through it. All the processing work is done on the server and the documents go into a print queue and are printed in the order received unless you tweak priority settings.
SQL	Provides database services, which are necessary when working with other server-types like SharePoint.
Exchange	An Exchange Server is used to provide email services. Users get a “mailbox” that allows them to send/receive email, keep track of their calendar and contacts, even receive voicemail (if configured properly).
SharePoint	SharePoint allows for easier collaboration within an organization. You access it through your browser and you can have web pages that have document libraries (with workflows, versioning and so forth), personal web pages, lists, and much more.
Lync	Provides a software based communications server system, that provides IM (instant messaging), Presence, VoIP (voice over IP) and conferencing capabilities (audio/voice and web conferencing).
Hyper-V	Virtualization services (discussed in a moment).

IIS (Web Server)	Internet Information Services allow that server to host web pages and web-based content.
-----------------------------	--

There are so many other server types for monitoring and managing (System Center tools), etc... and other options beyond Microsoft too.

The Big Takeaways

Ultimately, the big takeaway here is that there is a lot more to learn about the underlying network infrastructure than we let on at the beginning of the book. Truth is, to truly get into the world of Exchange you have to get a solid grasp of networking, DNS and Active Directory, Server installation and configuration and so on. So rather than scare you with all that to start with, we jumped into Exchange history.

Vendor Sponsor: Mimecast's Unified Email Management



“Do you really think you can fix this?”

“Email complexity? Oh yeah... we can get rid of that monster.”

Most information you read about when it comes to a third-party solution is written by the third-party. They tell you “we’re awesome! And here is a document that proves it! <cough><cough> written by us <said in a whisper>”. Even if it is true it certainly does cause an eyebrow to rise and the cynical side to us comes out.

That’s why I told my friends at Mimecast I wanted them to let me write this up in my way. I want you to see their solution through my eyes. I won’t be able to give you every last bell and whistle but I will certainly be able to tell you how it will add value to either your on-prem or Office 365 Exchange.

Mimecast was founded in 2003 by Peter Bauer and Neil Murray. These were regular people, IT admins, MCSE's, that saw a problem and went to work fixing it. The problem they saw was that email was becoming more and more complex to handle. They went to work on a solution that was in the cloud and provided email management.

Security

Email management can mean so many things, so what is it REALLY that Mimecast provides? Well, for starters, anti-spam and anti-malware. Keep the junk from ever reaching your on-premise Exchange or Office 365 servers. Mimecast's solution sits between your organization and the Internet and provides complete protection from spam, viruses, malware, phishing and data leaks.

Archive

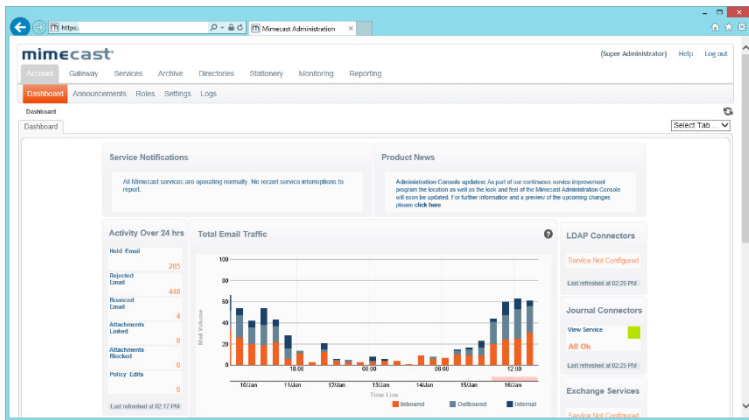
In addition, Mimecast provides an enterprise-grade archive solution with a powerful, high-performance eDiscovery piece. This reduces your on-premise storage costs because the archive ensures you have an accessible copy of that data at all times.

Let me explain this a bit further because I don't think everyone understands the value of this solution. If you recall the chapter on regulatory compliance we talked about having a personal archive, which is great for eliminating PST files but not great for enterprise archive and regulatory compliance protection. Why? Because end-users can delete whatever they want. And for that to stop you have to enable a form of legal hold (litigation hold or In-Place Legal Hold in Exchange 2013). This creates more storage bloat but does stop end-users from deleting things permanently.

With the Mimecast solution you have email archived before it even reaches your on-prem/O365 servers. Users can delete whatever they want, who cares? You have an archive. Now the cool thing is that this is an accessible archive, not backup tapes that sit in a vault. End-users are given tools that integrate with Outlook so that they can peruse their archive and find emails they

may have deleted accidentally and restore them (no IT intervention required... just a little training). BUT... if they want to delete an email that may be incriminating... nope, not possible.

I like to call this “preventative litigation”. Think about it. If you know, as an end-user, that everything you send and receive is being archived, is non-deletable, is easily located with eDiscovery... how stupid would you have to be to send something inappropriate? Hence, preventative litigation.



Continuity

I remember at 5 years old being in the movie theatre for the first Superman with Christopher Reeves. Do you remember the part where Lois Lane falls out of the helicopter and Superman catches her saying “Don’t worry maam, I’ve got you.” And she says “You’ve got me?! Whose got you!!!!???” Classic line. Good question though.

So, you have all these different types of Service Level Agreements out there (we talked about this in the book). SLA’s promise many things and one of them is availability of your services. But what happens if/when service goes down? It happens. It happens with on-premise Exchange and it happens with hosted solutions and even Office 365. Sure, the SLA typically offers some kind of restitution but what if you don’t want restitution, you want availability of service?

Here is where Mimecast is a brilliant solution. They keep users working during outages on-prem or in the cloud. They back it with 100% service availability SLA.

So, let's say service goes down. With Mimecast your end-users have no idea there is a problem. They can continue to send and receive email as if there was no failure because of the Outlook integration piece. So they just keep working. Once your servers come back online, Mimecast will sync up with them and the world keeps turning.

So even though you may have plenty of on-premise availability features in place (lots of redundant servers and so forth) you cannot be your own continuity solution. Even Microsoft, with O365, cannot be its OWN continuity solution. You need a third-party friend to help out with that. That's where Mimecast comes in.

File Archiving

In addition to the email archive capabilities Mimecast can also provide secure, reliable, scalable archive for all types of files. Data is all over the place these days. Network shares, end-user's drives, in SharePoint, in Exchange, in cloud services like Box or Dropbox. Mimecast brings all your data together into an easy to search archive. Imagine a single, manageable archive containing all of your corporate data. Brilliant.

Final Thoughts

I do a lot of product reviews. For the past 5 years (give or take) I've been writing them for the site MSEXchange.org (one of the best resources for Exchange info next to the Exchange Team blog itself). Of all the products I've reviewed... all of them... Mimecast is my favorite. I am only allowed to give 5 stars to products (top of the line) and I wanted to give Mimecast 6 stars. They wouldn't let me.

I integrated Mimecast with my personal company (ClipTraining) and our Office 365 solution. It took no time at all, very little stress on our side, and the benefits have been immediately

obvious. We have control over our business email like nothing we have had before. And it took very little effort to make it happen.

So, that's my personal opinion on Mimecast's Unified Email Management solution. I'd recommend you check them out. The added value you will receive for such a reasonable price point is unbelievable.

- J. Peter Bruzzese


Index

@ Symbol	7
Anti-Malware	111
Anti-Spam	111
Borenstein, Nathaniel	7
Certificates	113
Circular Logging	49
Creating Recipients (Mailboxes)	52
Cumulative Update	42
Deleted Item/Mailbox Retention	49,50
Extensible Storage Engine (ESE)	12
High Availability	
Continuous Replication	78
Local Continuous Replication (LCR)	79
Cluster Continuous Replication (CCR)	80
Database Availability Group (DAG)	82-85
DAC Mode	85
History of Exchange Server	
4.0, 5.0, 5.5	11
2000	13
2003	14
2007	15-17
2010	17-18
2013	18-21
Hybrid Exchange	122
Jetstress	106
Just a Bunch of Disks (JBOD)	41
LoadGen	106
MAPI	10
MIME	7
MS Mail	11
Office 365	120
Perimeter (DMZ)	28
Proxy	24
Regulatory Compliance	
Autodiscover	52
Recipient Types	55,56
Public Folder Mailboxes	57-59
Personal Archive	65
In-Place Hold (Legal/Litigation Hold)	66
In-Place eDiscovery	67
Messaging Records Management (MRM)	68,69

Journaling	70
Transport Rules	70-72
Data Loss Prevention	72
Audit Logging	73
Information Rights Management	74
Role Based Access Control (RBAC)	112
S/MIME	117
SAN/UC certificates	116
Server Roles Defined (2007/2010)	
Mailbox	26
Client Access	26
Hub Transport	27
Unified Messaging	27
Edge Transport	28
Server Roles Defined (2013)	31
Server Role Requirements Calculator	106
Service Level Agreements (SLA's)	77
Single Instance Storage	11
Smart Host	28
SMTP	7
Spam Confidence Level (SCL)	111
Storage groups	39
Transaction Logs	44
Transport Layer Security (TLS)	117
Transport Pipeline	34,35
Unified Messaging	
PBX (IP-PBX)	92
Outlook Voice Access (OVA)	93
Voice Mail Preview	93
Auto Attendant	94
DTMF	94
PSTN	95
Dial Plan	97
UM Mailbox Policy	99
Virtualization	102
Wildcard certificates	116
XENIX	11

Easily “converse” about Exchange Server in any setting.

Exchange Server has been evolving for 20 years and it has grown well beyond simple email services. The goal of this book is to help folks who work in fields around Exchange (sales/PR/marketing) as well as folks who need an introduction to Exchange features to be able to converse about it intelligently. You may never install Exchange, configure it, etc... but you can sound like you have!



About J. Peter Bruzzese

J. Peter is an Office 365 MVP, holds a variety of certifications (MCSE/MCITP: Messaging, A+/Network+, etc.), is an internationally published author and conference speaker, InfoWorld journalist, ClipTraining Co-Founder and CIO and more. Follow him on Twitter @JPBruzzese



ConversationalGeek®

Visit conversationalgeek.com for more books on topics geeks love.