SysAdmin Handbooks

The SysAdmin Handbook

The Best of Simple Talk

ISBN: 978-1-906434-42-7

The SysAdmin Handbook

The Best of Simple Talk SysAdmin

First Published by Red Gate Books, 2010

Red Gate Books
Newnham House
Cambridge Business Park
Cambridge
CB2 0WZ
United Kingdom

ISBN 978-1-906434-42-7

Editors: Tony Davis, Andrew Clarke and Michael Francis

Cover Art: Paul Vlaar

Typesetting & Design: Gower Associates and Matthew Tye

Table of Contents

Exchange 24

General Articles 394

Unified Messaging 479

PowerShell 493

About the Authors

Jaap Wesselius

Jaap Wesselius is a senior Exchange consultant for DM Consultants (http://www.dm-consultants.nl/), a Microsoft Gold Partner with a strong focus on messaging and collaboration solutions. Prior to working for DM Consultants, Jaap worked for eight years for Microsoft Services in The Netherlands, specializing in Exchange Server. Jaap is a BSc, MCSE, MCITP, and MCT, and was awarded the Microsoft MVP Award (Exchange Server) for his contributions to the Dutch messaging and collaboration community. Besides Exchange Server, Jaap is also very active in virtualization, and is a founder of the Dutch Hyper-V community. You can reach Jaap at J.Wesselius@dm-consultants.nl or Jaap@hyper-v.nu.

Ben Lye

Ben Lye is a senior systems administrator at a multinational software company. He has over ten years' experience of supporting and administering Windows and Exchange, and has been MCSE and MCP certified since 1999. Ben is passionate about automating and streamlining routine tasks, and enjoys creating and using tools which make day-to-day administration easier.

Michael B. Smith

Michael B. Smith is a well-known author and consultant in the Exchange Server and Active Directory arenas. His most recent book is *Monitoring Exchange Server 2007 with System Center Operations Manager 2007*, published in February 2009. You can contact Michael via e-mail at michael@TheEssentialExchange.com and read his blog at http://TheEssentialExchange.com.

Jonathan Medd

Jonathan Medd is a Senior Technical Consultant for a large local government organisation in the UK. He has been working with Windows Infrastructure products since 1997, in the last few years particularly around Active Directory and Exchange and, most recently, virtualisation with VMware products. He has held the MCSE certification since 1998 and VMware VCP certification since 2008. In 2007, he discovered Windows PowerShell and he spends a lot of time encouraging IT pros he meets to use PowerShell by talking with them, by giving presentations to User Groups, or via posts on his blog at http://jonathanmedd.net. He also co-hosts a regular PowerShell podcast which contains info on how to learn PowerShell and what's going on in the PowerShell world – you can find this at http://get-scripting.blogspot.com. You can follow him on Twitter at http://twitter.com/jonathanmedd.

Brien Posey

Brien Posey is a freelance technical writer, and a five-time Microsoft MVP. Over the last thirteen years, Brien has published literally thousands of technical articles and whitepapers, and written or contributed to dozens of books. Prior to becoming a freelance writer, Brien served as CIO for a national chain of hospitals and healthcare facilities. He has also served as a network administrator for the Department of Defense at Fort Knox, and for some of the nation's largest insurance companies.

Matt Simmons

Matt Simmons is an IT Administrator with several years' experience on small and medium networks. He is currently employed in the financial services industry, and has previously worked in logistics and Internet services. Matt maintains a systems administration blog, and spends his spare time reading and learning new things. He can be reached at STANDALONE.SYSADMIN@GMAIL.COM or via his blog at HTTP://WWW.STANDALONE-SYSADMIN.COM.

Neil Hobson

Neil is a Principal Consultant with Silversands, a UK-based Microsoft Gold Partner and is responsible for the design, implementation, and support of Microsoft infrastructure systems, most notably Microsoft Exchange systems. He has been in the IT industry since 1987 and has worked with Exchange since V4.0 days. He has been an Exchange MVP since 2003 and spends some of his spare time helping others in various Exchange mailing lists, and the public Exchange newsgroups; he also contributes to the MSExchange Blog.

Desmond Lee

Desmond Lee specializes in end-to-end enterprise infrastructure solutions built around proven business processes and people integration across various industries. He is a popular speaker at major international events, and contributes frequently to several highly rated publications and public forums/newsgroups. Desmond is a Microsoft Most Valuable Professional (Communications Server), Microsoft Certified Trainer and founder of the Swiss IT Pro User Group, an independent, non-profit organization for IT Pros by IT Pros championing Microsoft technologies. You can follow his IT adventures at WWW.LEEDESMOND.COM.

Hilal Lone

Hilal works as a Senior Integration Engineer (SME on Email Security) with Continuous Computing (WWW.CCPU.COM). Continuous Computing specializes in deploying uniquely architected solutions comprised of telecom platforms and Trillium software. Currently, Continuous Computing caters to the LTE, DPI, Femtocell, and Professional Services market. Hilal has been working in IT messaging, Security and Networking for more than six years; he specializes in end-to-end messaging and is deeply involved in Information Security and Networking. His qualifications include B.Sc., M.Sc. (Information Systems), MCSE, MCITP (Enterprise Messaging), CCNA, CCSP, and CEH. Apart from Exchange Server, he specializes in multi-platform firewalls and Cisco Networking equipment.

Nicolas Blank

Nicolas Blank is an Exchange MVP and consultant at Symbiotech (HTTP://WWW.SYMBIOTECH.CO.ZA), a consultancy that specializes in the Microsoft Infrastructure and related tools space with a strong focus on messaging, collaboration, migration, and security solutions. Nicolas currently builds solutions based on Active Directory, Exchange, Office Communication Server and a variety of third-party vendors. Nicolas consults, speaks, writes, and delivers seminars on various topics and blogs at HTTP://BLANKMANBLOG.SPACES.LIVE.COM.

William Lefkovics

William Lefkovics, B.Sc., MCSE, is the Technical Director at Mojave Media Group, LLC in Las Vegas, NV. He is the co-author of *Microsoft Exchange Server 2007: The Complete Reference* and contributes a monthly column on Outlook at Windows IT Pro Magazine.

Nathan Winters

Nathan Winters is a Unified Communications Consultant for Dimension Data; a Microsoft Gold Partner on five continents whose clients include over 70% of the Global Fortune 100. He has been working in IT for four years, and specializes in Exchange, Active Directory, and Virtualization. Recent work has included an Exchange 2007 design for several clients and an OCS 2007 voice design including Enterprise Voice. Midway through 2006, Nathan founded the Microsoft Messaging and Mobility User Group UK, which holds regular meetings in the UK to discuss topics related to Exchange. In April 2007, he was awarded an MVP (Exchange Server) for his work with MMMUG and regular contributions to the Mark Minasi Forum. He is a regular contributor to the MSExchange.org website and his other articles have been published by Penton Media (Exchange and Outlook Administrator newsletter), Microsoft (TechNet Industry Insiders) and on the MMMUG website. For more of his articles see the links below:
http://www.mmmug.co.uk/files/default.aspx
http://www.msexchange.org/nathan_winters
http://www.windowsitpro.com/Authors/AuthorID/1651/1651.html
You can contact Nathan at Nathan@clarinathan.co.uk or through his blog at http://www.mmmug.co.uk/blogs/nweb.

Nirmal Sharma

Nirmal is an MCSEx3, MCITP, and was awarded a Microsoft MVP award in Directory Services four times. He specializes in Directory Services, Microsoft Clustering, Hyper-V, SQL, and Exchange. He has been involved in Microsoft Technologies since 1994, and followed the progression of Microsoft Operating System and software. He specializes in Microsoft technologies. In his spare time, he likes to help others and share some of his knowledge by writing tips and articles. He can be reached at nirmal_sharma@mvps.org.

Dr Masha Petrova

Dr Masha Petrova is the founder and CEO of MVP Modeling Solutions, a company dedicated to the improvement of engineering R&D processes, and a creator of www.SuccessfulUnemploymentToolkit.com. She received her Ph.D. from the University of California at San Diego, where she conducted research in chemical kinetics. She has created and taught courses on engineering computer modeling and has trained professionals all over the globe. She is a co-founder of the Women in Combustion group, a featured speaker for the ACS Speaker Service and an official instructor for the ACS courses. She writes a monthly column for the Product Design and Development online magazine. She is also a member of the Executive Board of Directors for the Franklin Foundation, an organization dedicated to promoting education and innovation in America's technology and science fields.

Bill Holmberg

Bill Holmberg has been an IT professional since shortly after building Altair, Heathkit, and Apple computers from kits in the late 70s. Since then he has helped companies build IT infrastructures, help desks, WANs, and even an onsite support group rated "World Class" by Gartner. As a consultant, for over 12 years, to AlphaBetas (which he helped found, and for which he is the CIO) he has helped bring over 100 games and software titles to market, whether as the Project Manager, QA/QC analyst, or even as voice-over talent. Bill also helps present classes funded by the Department of Justice for Minneapolis' InfraGard (an FBI volunteer group) on Security and Forensics Education (SAFE) and has dealt with many network intrusions for private clients. He is currently re-engineering the 8-location WAN to best practices for a food processor and recycler in Minnesota, and implementing their DR plan. In his "spare" time, he has resurrected his Grandfather's namesake family business, Hart Skis.

Introduction

Over the past two years, Simple-Talk has published articles on a variety of SysAdmin topics, from Exchange to Virtualization, and including everything from Powershell to Unified Messaging.

We have brought the best of these articles together to form The SysAdmin Handbook. With over fifty articles packed into this book, it will be an essential reference for any Systems Administrator, whether you have years of experience or are just starting out.

Simple-Talk is an online technical journal and community hub for working SQL Server and .NET developers and administrators, as well as enthusiasts and those taking their first steps with these Microsoft technologies. Simple-Talk provides in-depth technical articles, opinions and commentary to more than 388,000 subscribers.

The Simple-Talk SysAdmin newsletter is a fortnightly publication. Subscribing to the newsletter means you will get the latest straightforward and practical technical articles straight to your inbox. If you do not already receive the newsletter, subscribe for free HERE.

We would like to thank all the authors who have contributed SysAdmin articles to Simple-Talk over the past two years, as well as those who have taken their time to read the articles, and participate in the community. Because of this, we have been able to develop a successful SysAdmin community on Simple-Talk, filled with high-quality and useful technical articles.

The Simple-Talk Publishing Team

Exchange

High Availability in Exchange 2007

10 June 2008

by Neil Hobson

Neil Hobson writes about the ways that MS Exchange 2007 can ensure that your organisations messaging remains available. He looks at the way that the Mailbox Server role can be made more available using features as Single Copy Clusters, Local Continuous Replication, and Clustered Continuous Replication. He also discusses ways of improving the resilience of the other server roles.

If you are thinking of deploying Exchange 2007, the chances are that you have a requirement to implement the product with high availability in mind. More and more organizations are placing much more importance on the messaging infrastructure and, as a result, high availability and site resilience have a much higher emphasis than they did several years ago. In this article I'm going to cover the high availability options within Exchange 2007, giving you an outline of each one and the options you have for implementing the product with high availability in mind.

I'm going to assume that you've already made the decision to implement Exchange 2007 in a highly available manner, so I won't be covering the discussion points around whether high availability is appropriate for you or your organization. That will make a useful discussion topic for a future article. Another subject that will also make a good discussion topic for a future article is the site resilience option available for Exchange 2007 and therefore I won't be covering that at this time.

Finally, bear in mind that I'm going to be covering the native options available in Exchange 2007 and therefore no 3rd-party high availability solutions will be detailed within this article.

As you likely know, Exchange 2007 is comprised of five different server roles and each one has high availability options, some of which differ quite significantly from others. We'll take each of the five roles in turn and detail the options you have for making these roles highly available.

Mailbox Server Role

First, let's take a closer look at the high availability options for the Mailbox server role. This is where the users' actual mailbox data is stored so it's one of the key roles to make highly available.

When you think of high availability for mailboxes in Exchange, you will most likely think of clustering since many previous versions of Exchange have included this technology. So it will come as no surprise that clustering is also at the heart of high Availability in Exchange 2007, although there are some significant changes that may cause you to re-evaluate clustering if you have previously dismissed it.

Single Copy Clusters

The first technology we'll look at is Single Copy Clusters (SCC). This technology will be familiar to those Exchange administrators who have implemented clustering in previous versions of Exchange since it's essentially the new name for 'traditional' Exchange clustering.

SCC in Exchange 2007 is essentially the same as previous versions of Exchange clustering. This means that it still uses the shared storage model, where the actual mailbox and public folder databases only exist once within the storage infrastructure (hence the term *single copy*). The individual cluster nodes of the SCC environment can all access the same shared data, but only one node at a time can actually use it. Figure 1 shows a simplified diagram of how this might look with a three-node SCC environment.

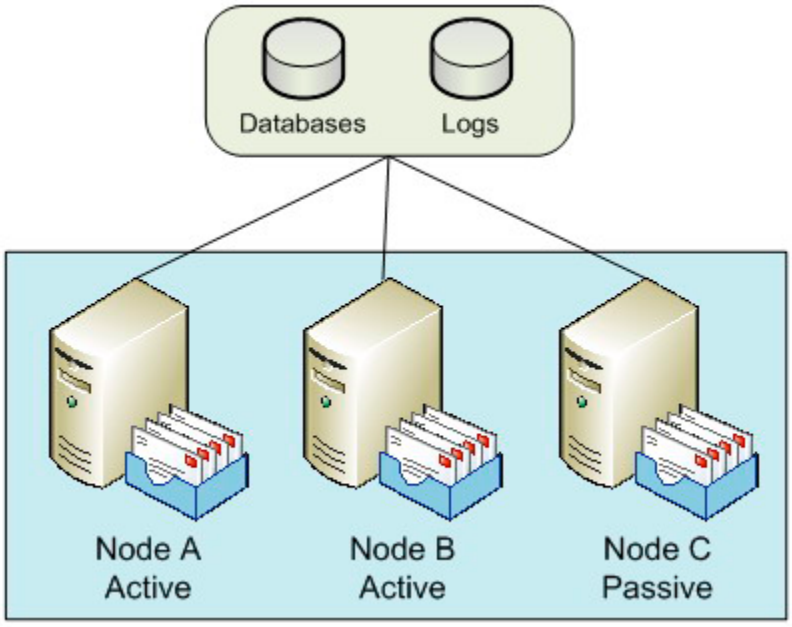

Figure 1: Single Copy Clusters.

In Figure 1, you can see three nodes within the SCC environment, two of these are curretly *active* nodes (i.e. they are actively servicing users), while the third is passive. Exchange 2007 clusters only support the active/passive model: it's possible to have up to eight nodes configured as long as one is passive

The passive node is available to take ownership of the data currently belonging to one of the active nodes, if and when an issue arises with one of those active nodes. With up to eight nodes available, some interesting combinations can be designed. For example, you might consider a six-active and two-passive node system, or perhaps a four-active and four-passive node system. The latter option actually falls within the Microsoft best practice recommendations, since Microsoft recommends that you have at least one passive node for every active node that you implement.

You might consider a configuration such as a two active and two passive node SCC environment to be wasteful in terms of hardware resources, since two nodes aren't actually offering a service to the users. However, you have to consider the reason that you may have implemented such a design: high availability, and therefore the preservation of your messaging service in the event of a node failure. For example, consider the case in Figure 1, where you have two active nodes and a single passive node. If Node A fails, the service will failover to use Node C and users are largely unaffected. However, you are now in the situation where you no longer have a passive node available until you resolve the problem with Node A. Therefore, if Node B were to fail before Node A is brought back into service, there is nowhere for the services on Node B to failover to and thus users will be affected. Therefore, having at least one passive node per active node is good practice.

You can clearly see that an SCC environment is excellent at providing high availability in cases where there is a server failure. Although SCC environments can and do provide good levels of uptime, the main drawback with SCC is the fact that there is only a single copy of the

storage present. Some organizations implement their Exchange databases on replicated Storage Area Networks (SANs) to overcome this. Here, synchronous data replication can be used to ensure that the Exchange databases are copied to a different SAN, typically located in a different data centre.

One key area to consider regarding clustering is that if you deploy the Mailbox server role on a clustered solution, either using SCC or Clustered Continuous Replication (CCR), which will be described later, no additional Exchange 2007 server roles can be combined with the Mailbox server role. Consequently, the minute you decide to deploy a SCC environment for your Mailbox servers, you will need additional servers to run other roles such as the Hub Transport and Client Access Server roles.

Local Continuous Replication

This is the first of the new *continuous replication* technologies available within Exchange 2007. The first and most obvious point to make about Local Continuous Replication (LCR) is the fact that it is a single-server solution and not a clustered solution. Therefore, LCR will not protect you against the failure of an entire server. Having said this, LCR does implement the new **log shipping and replay** functionality that Exchange 2007 provides. It does this by shipping the transaction logs generated by a storage group, known as the **active copy**, to another separate set of disks that are connected to the same server, referred to as the **passive copy**. Once the logs have been transferred to the alternate disks, they are replayed into a copy of the Exchange database that also resides on these disks. Thus, a separate copy of the database is maintained in near-real time fashion on the same server, and you therefore have data redundancy. Should there be a problem with the production database the administrator can switch over to using the backup copy of the database fairly quickly. An overview diagram of LCR is shown in Figure 2.

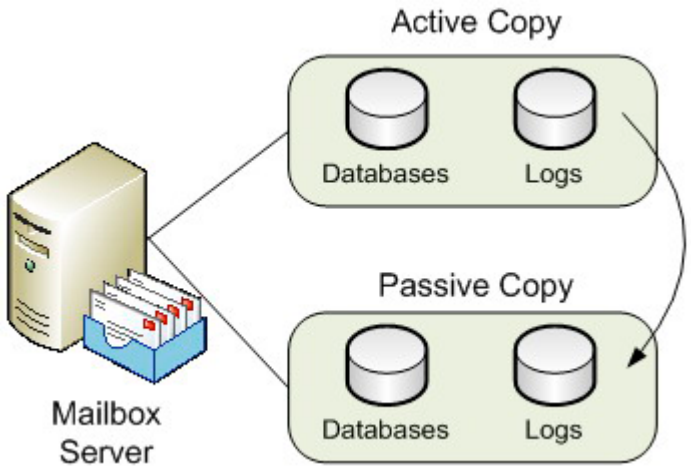

Figure 2: Local Continuous Replication.

It makes the most sense to deploy the second set of disks via a separate disk array controller to cover for the failure of the primary disk array controller. As I've already mentioned, LCR is not a clustered solution and therefore does not protect you should the entire server fail. However, it is a good solution to protect against mailbox data corruption and additionally offers the ability to offload Volume Shadow Copy Service (VSS) backups from the active copy of the databases. In this scenario, you configure the VSS backup to be performed against the passive storage group. This is better for disk performance on the active storage group as well as allowing online maintenance to be unaffected by the backup process.

Clustered Continuous Replication

Whilst SCC offers you protection against server failure and LCR offers you protection against data failure, Clustered Continuous Replication (CCR) offers you both server and data protection. As you can guess from the name, CCR is the second of the new continuous replication technologies available with Exchange 2007. A CCR environment is a two-node cluster only, consisting of an active and a passive

node. The key difference between a CCR environment and a SCC environment is that the CCR environment does not use shared storage. Rather, both nodes of the CCR environment have their own copies of the Exchange databases and transaction logs. The transaction logs from the active node are asynchronously copied to the passive node and replayed into the database. Should a problem occur with either the active node itself or the active node's databases, the Exchange server can be failed over to run from the previously passive node and its own copy of the databases. A sample CCR environment configuration is shown in Figure 3.

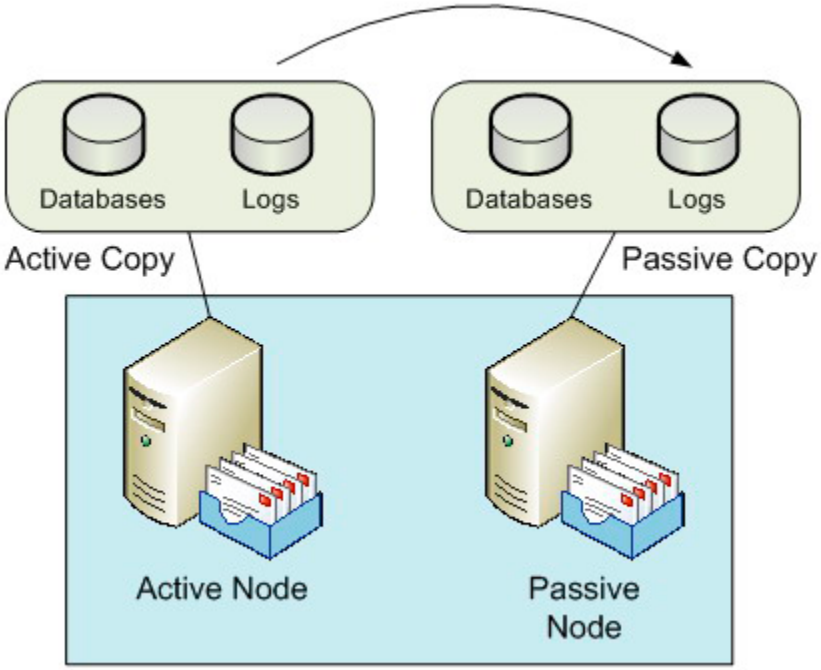

Figure 3: Clustered Continuous Replication.

As I stated earlier, a CCR environment consists of two nodes; it cannot be scaled out to eight nodes as with the SCC environment. However, the fact that CCR implements two separate copies of the databases and transaction logs is usually attractive to many organizations. Also, CCR can be implemented using Direct Attached Storage (DAS) rather than SAN storage. This, too, can appear attractive to many organizations although it is important to note that this may not necessarily be a straight decision to make as there are other factors to consider. Also, VSS backups can be taken from the passive node, which is obviously good for the performance of the active node, plus it has the added benefit of allowing online database maintenance to run outside of the backup window.

Some organizations may choose to implement the active CCR node in one data center with the passive CCR node in a different data centre. This is technically fine, although it should be noted that if you are deploying CCR on the Windows 2003 operating system, both CCR nodes must be in the same subnet. You will therefore need to check with the network team in your organization as to how the network is configured across the two data centers. This restriction does not apply to Windows 2008 since clusters can be implemented with nodes from different subnets, although the nodes must be in the same Windows Active Directory site. With nodes in different data centers, the CCR environment can be considered a **stretched cluster**. You will also need to factor in considerations such as the available network bandwidth between the two data centers as well as the latency of the network connection. Such factors can affect the performance of the system, especially when you consider the fact that the various server roles will communicate via Remote Procedure Calls (RPCs) and thus require sufficient bandwidth.

Consequently, many organizations choose to implement a CCR environment within their production data center and Standby Continuous Replication (SCR) to a disaster recovery data center. CCR then protects from server and data failure within the same data center, which is more likely than the complete loss or failure of a single data center. Having said that, you may still need to plan for that eventuality which is where SCR comes in. SCR is a site resilience solution available in Exchange 2007 Service Pack 1 that currently falls outside of the topics that I want to cover in this article.

Hub Transport Server Role

If you have chosen to implement one of the high availability solutions for the Mailbox server role, your decision will likely be meaningless if you do not consider high availability for the Hub Transport role. The reason for this is an architectural change to the way messages are routed in Exchange 2007. In Exchange 2007, all messages are routed via the Hub Transport server role, even messages between users on the same database on the same Mailbox server. Therefore, if you have deployed an Exchange 2007 CCR environment, for example, but at the same time deployed just a single Hub Transport server, you should be aware that no messages will flow around your infrastructure should the single Hub Transport server fail.

It's therefore vital to include the Hub Transport server role in your high availability planning. Fortunately, including high availability with the Hub Transport server role isn't as complicated as the Mailbox server role, since this feature is built into the Hub Transport role by default. Therefore, all you essentially need to do is deploy multiple Hub Transport servers; no clustering or load balancing is required. For example, if you've deployed a single CCR environment consisting of two clustered Mailbox server roles, you'd likely deploy two independent Hub Transport servers for a balanced design.

Load balancing and redundancy is automatic with the implementation of multiple Hub Transport servers, and therefore no additional configuration is required. If two Hub Transport servers are deployed within the same Active Directory site, the Mailbox server will use them both; if one Hub Transport server fails, the remaining Hub Transport server will continue to process the messages from the Mailbox server. Additionally, when considering communications across Active Directory sites, a Hub Transport server in one Active Directory site will automatically use multiple Hub Transport servers in a different Active Directory site.

As I've already said there is no need to deploy hardware load balancing solutions or technologies such as Windows Network Load Balancing (NLB) when considering the Hub Transport server role. However, you could find yourself in the position where you wish to implement the Hub Transport server role on the same servers as the Client Access Server role, which themselves may be configured with Windows NLB. If that's the case, it's important to know that you can use Windows NLB to load balance the client connectors if you are using Exchange 2007 Service Pack 1, but take care to ensure that your Windows NLB configuration excludes the default SMTP communications on port 25 that occur between Hub Transport servers. We'll look a bit more at Windows NLB in the next section.

Client Access Server Role

As with the Hub Transport server role, if you've elected to implement high availability for your Mailbox server role you should consider the same for the Client Access Server role. However, you may be thinking that you don't need to do this based on the fact that maybe you're not implementing client access methods such as Outlook Web Access, POP3, IMAP4 and so on. After all, these are the types of clients that communicate via a Client Access Server, right? Well, whilst that is absolutely correct, don't forget that the Client Access Server also runs the Availability and Autodiscover services, which are also important for clients running Outlook 2007. If you're spending the money to implement high availability for the Mailbox and Hub Transport server roles, implement high availability for the Client Access Server role too. The most common way to do this is to implement Windows NLB to load-balance the relevant communications ports. Of course, hardware load balancing can be used too. With Windows NLB, you essentially have a single IP address that is shared amongst the servers, which have their own IP addresses themselves. Figure 4 gives you an idea of how this might look, where a user is connecting to Client Access Servers.

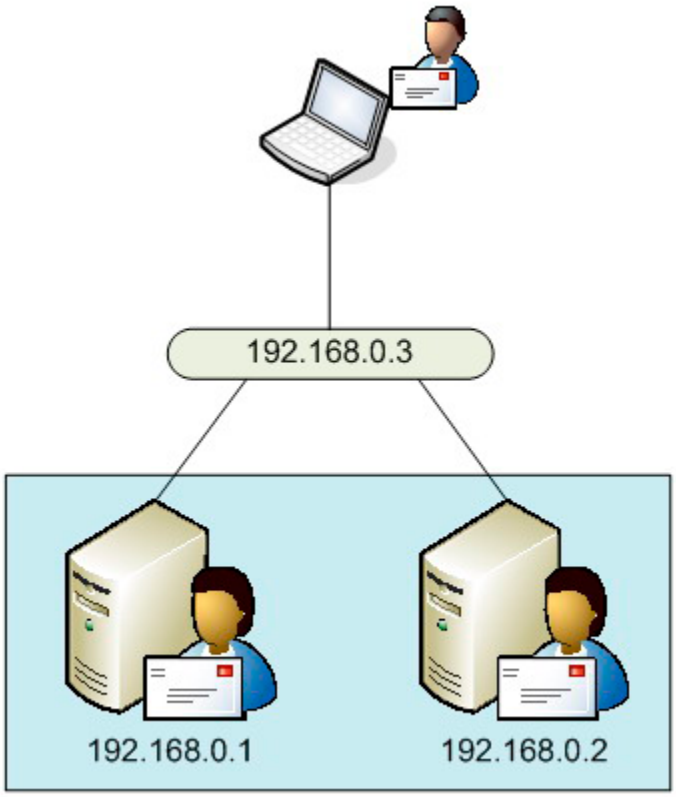

Figure 4: Windows NLB.

As I mentioned in the previous section, on Hub Transport servers, some organizations choose to combine the Client Access Server and Hub Transport server roles onto the same server. This configuration wasn't officially supported by Microsoft in the Release To Manufacturing (RTM) version of Exchange 2007 but is supported now that Exchange 2007 Service Pack 1 has been released. However, don't forget what I previously mentioned within this article around the load balancing of client connectors, and the need to exclude the Hub Transport server SMTP communications from the load balancing configuration. Equally it should be noted that you don't actually have to combine the Hub Transport and Client Access Server roles on the same server. In fact, in larger deployments the Client Access Server and Hub Transport server roles are separated onto different servers. This means that deployments consisting of a single CCR environment can actually see six servers deployed, namely two for the CCR environment, two for the Hub Transport server role and two for the Client Access Server role. This can be further added to by the deployment of the Unified Messaging and Edge Transport server roles which will be covered in the next two sections.

Unified Messaging Server Role

High availability for the Unified Messaging server role is similar to that of the Hub Transport server role in that you need to deploy multiple Unified Messaging servers to achieve redundancy. However, there is a little more configuration to do to achieve this. The key to achieving this high availability is to ensure that the Unified Messaging servers are configured in a single **dial plan**. The dial plan is effectively the link between an Active Directory user's telephone extension number and their Exchange 2007 mailbox that has been enabled for Unified Messaging. Your Voice over IP (VoIP) gateways find your Unified Messaging servers via the dial plans and will attempt to connect to one of those Unified Messaging servers. Ultimately, if no response is received, the VoIP gateways will try the next available Unified Messaging server. I use the term gateways (plural) since you'll likely implement multiple VoIP gateways in addition to multiple Unified Messaging servers to complete your redundant configuration.

Of course, in the example situation where two Unified Messaging servers are available, the VoIP gateways can be configured to balance incoming calls across both Unified Messaging servers. Note that technologies such as Windows NLB and round-robin DNS aren't used to provide this high availability and redundancy. Rather, the VoIP gateways can be configured with either the IP addresses or the Fully Qualified Domain Names (FQDNs) of both the Unified Messaging servers.

Edge Transport Server Role

The Edge Transport server is unique amongst the five Exchange 2007 server roles in that it is the only server role that has to be deployed independently, and therefore cannot coexist with any other server role. Achieving high availability for the Edge Transport role once again involves the implementation of multiple servers. Since your Edge Transport servers effectively sit in between your Hub Transport servers on your internal network and the Internet, you simply need to ensure that your Hub Transport servers and the SMTP servers of everyone else on the Internet know how to contact them.

For inbound messages, the typical way to achieve this is to implement multiple Mail Exchanger (MX) records in DNS. You can either implement MX records with identical weightings, in which case both Edge Transport servers will be used equally, or you can implement the MX records in priority order. In the latter case, one Edge Transport server will be used in preference to the other, with the secondary Edge Transport server only being used if the primary Edge Transport server fails.

For outbound messages, the preferred option is to create Edge Subscriptions for each Edge Transport server. An Edge Subscription essentially links the Edge Transport server to the Hub Transport servers in a specified Active Directory site, thereby allowing, via Send Connectors, email messages to flow from the Hub Transport Servers to the Edge Transport servers.

Summary

In this article we've looked at the high availability options for each of the Exchange 2007 server roles. As you've seen, the most interesting options are available for the Mailbox server role but at the same time it's vital to implement high availability for the other server roles if you've elected to choose a clustering solution for your Mailbox servers.

Message Hygiene in Exchange Server 2007

03 July 2008

by William Lefkovics

Around four out of every five email messages are spam. Now that the nuisance threatens to engulf what has become the normal route of business communications worldwide, all the vendors of Mail Servers are taking more active steps to confront the problem. William Lefkovics explains how Microsoft Exchange 2007 is meeting the challenge of the rising tide of spam.

Message Hygiene in Exchange Server 2007: A Defense-in-Depth Approach

Unsolicited Commercial E-mail (UCE), or spam, is still a huge problem. Different research firms place the total percentage of spam to be anywhere from 70 to 90% of the total volume of e-mail around the world. Different enterprises may see variances in either direction from this range. It is still too much.

Spam Botnets

Most spam today is distributed by networks of unsuspecting workstations infected with malware. These bots are small applications capable of, among other things, forming and sending SMTP messages en mass by remote command. Spammers typically pay for use of botnets (short for "robot networks") rather than maintain their own. In the United States, the Federal Bureau of Investigation (FBI) has been working through "Operation Bot Roast" which identified 1 million infected hosts in the US alone and has resulted in the arrest and conviction (or guilty pleas) of several botnet operators. (http://www.us-cert.gov/press_room/botroast_200711.html)

Microsoft, in response to this ongoing battle, improved upon their set of native tools in Exchange Server 2003 Service Pack 2. They provided a more complete system for administrators to eliminate known spam early in the SMTP conversation and give them more control in assessing spam levels on more questionable messages. Exchange Server 2007 expands and improves on this defense-in-depth approach giving administrators effective sequential tools without additional software cost.

There is no silver bullet that will prevent spam from reaching inboxes, but there are several steps Exchange administrators can take to reduce inbound spam significantly. With Exchange Server 2007, Microsoft introduced server roles. There are five roles in total - Mailbox, Client Access, Hub Transport, Unified Messaging, and Edge Transport. The Edge Transport role is intended to reside in a perimeter network and perform the majority of the inbound message hygiene functionality for the Exchange organization. The Edge Transport server is typically the gateway SMTP server for the enterprise. Assuming so, message hygiene begins here.

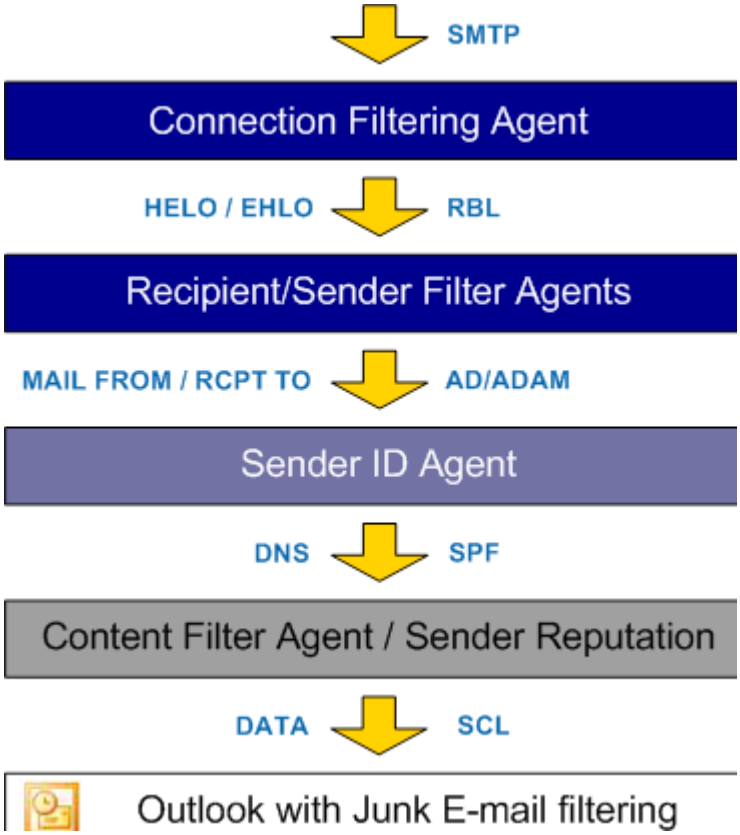

Figure 1.

Figure 1 shows the inbound flow of messages passing through the various anti-spam agent layers. Each layer implements a different test in the effort to validate messages. If 80% of messages arriving at your gateway are spam, then it certainly makes sense to drop those e-mails as early as possible in the SMTP conversation. The anti-spam tools available on an Exchange 2007 Edge server are shown in Figure 2.

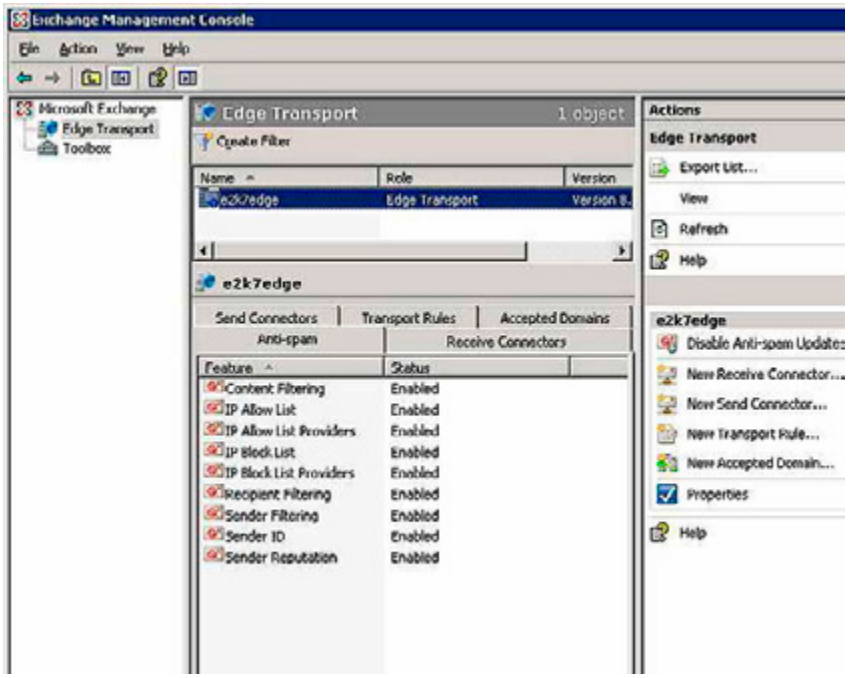

Figure 2.

Figure 2 shows the anti-spam tools available on an Exchange 2007 Edge server. Double clicking or selecting the desired anti-spam feature in the middle pane, and clicking Properties in the Action pane, will open the configuration settings for that feature. The four IP lists, allow and block, make up the Connection Filtering Agent.

Connection Filtering Agent

The first line of defense, and probably the most prolific one in terms of the potential volume of effective message filtering, is connection filtering.

Feel the Power!
The Exchange Management Shell (EMS) is an extension of Windows PowerShell for managing Exchange 2007. The EMC is built upon the EMS. For all the actions taken in the EMC, there is an equivalent EMS command. There is a complete set of cmdlets to manage connection filtering from the EMS. For example, *Get-IPBlockListConfig* on the Edge Server will return the IP Block list configuration.

For all the Exchange Server cmdlets run *Get-ExCommand* in the EMS. Microsoft maintains a library of these cmdlets as well. The anti-spam cmdlets can be found at http://technet.microsoft.com/en-us/library/ bb124601(EXCHG.80).aspx.

The Connection Filtering Agent is concerned with the source IP address of the host that is connecting on TCP Port 25. The agent references the IP address of the host making the connection request against an IP Allow List and an IP Block List. In addition it can query an IP Block List Provider or an IP Allow List Provider. If the IP address is listed on a block list, either internal or at a block list provider, and not on an allow list, then the connection is dropped. The IP Block list is maintained by the administrator. IP Block List Providers, also called Real-time Block Lists (RBLs) or DNS Block Lists (DNSBLs), maintain databases of known spam hosts.

A good strategy is to select one or two of the well known DNSBLs to complement your manual efforts to maintain corporate IP Block and IP Allow lists. Over the last five years, Spamhaus has been a solid DNSBL for me. I have also used Spamcop over that same time period. Spamcop has shown some inconsistency but has performed well over the last year or two. If you choose just one, I recommend *zen.spamhause.org*. In my company, and at clients, this step alone blocks over 55% of inbound spam messages. These are messages that my Exchange system does not have to process any farther. An independent resource for DNSBL performance can be found at http://stats.dnsbl.com/.

Connection filtering is exceptional! That is, it allows for an Exception List. You can configure Exchange to allow messages addressed to specific SMTP addresses to pass through even if the source host IP address is on a block list. It might be a good idea to add *postmaster@yourdomain.com* to the Exception List.

Configuring Exchange 2007 to query a DNSBL is simple. From the Exchange Management Console (EMC), open the configuration settings for IP Block List Providers and select the Providers tab. The Add button opens an input box to enter the name of the DNSBL.

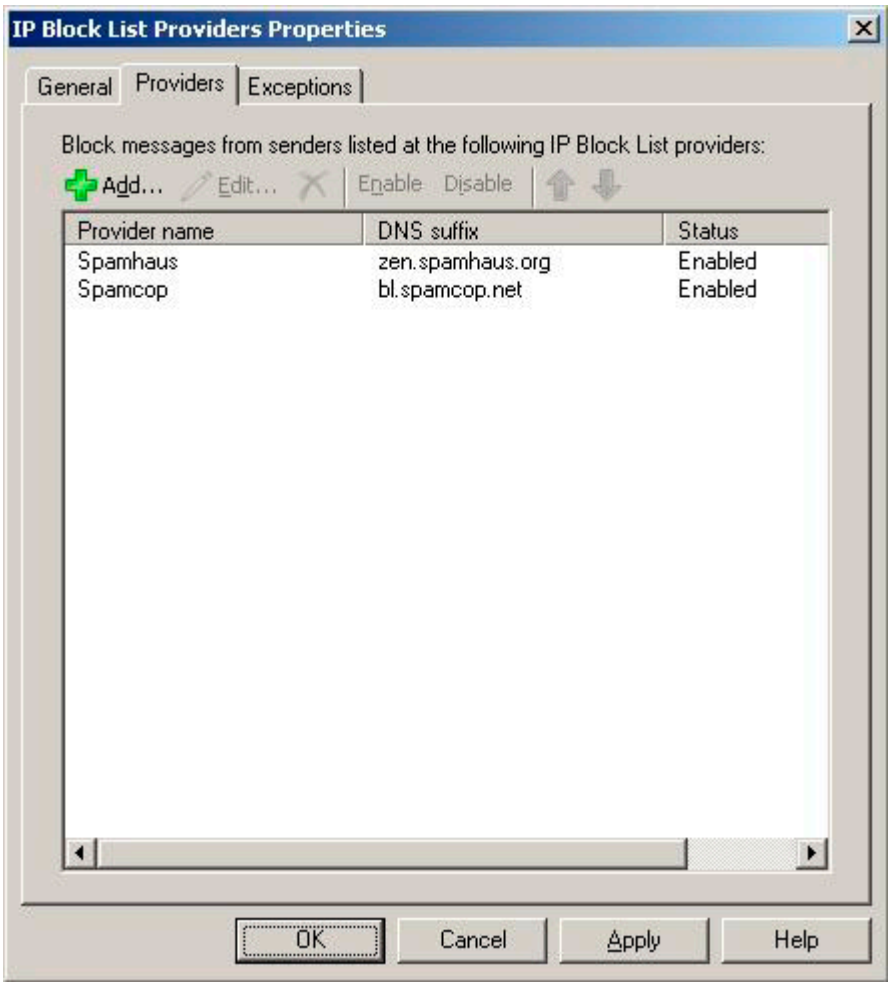

Figure 3.

Figure 3 shows the IP Block List Providers Properties window with two DNSBLs entered.

Sender Filtering Agent

If the message survives Connection Filtering, the administrator can have Exchange reference the sender SMTP address against a list of addresses to block. The sender address is available after the MAIL command in the SMTP conversation. Specific SMTP addresses can be blocked, or entire domains. Adding **.domain.com* to the Sender Filtering list will also block all subdomains for *domain.com*. While the default action is to block messages, Sender Filtering can also stamp the result in the message header and allow it to pass through. Content Filtering will consider the stamp status in its rating of the message (more about that later).

The EMC has a simple interface to enter blocked senders in the Sender Filtering properties. Addresses can also be added at the EMS command line with a simple cmdlet such as:

```
>Set-SenderFilterConfig -BlockedSenders badaddress@domain.com
```

Recipient Filtering Agent

The next command in the SMTP transaction is the RCPT command, where the intended recipient of the message is exposed. Here we can apply a filter based on the recipient information. It is possible to block incoming messages for specific addresses, but more importantly, it is also possible to have Exchange query ADAM (or Active Directory if the Anti-spam agents are run on the Hub Transport server) to validate the recipient address.

No Edge Server? No problem!
In the absence of an Edge Transport server, as may be the case in smaller companies, the Hub Transport role can host the Anti-Spam agents. First, on the Hub Transport server, you must run the PowerShell script *Install-AntispamAgents.ps1*, found in the *\scripts* folder on the Exchange 2007 DVD or extracted download.

After the script has completed, the Exchange Transport Service must be restarted. There will now be an Anti-Spam tab for the Hub Transport Server node under the Organisation Configuration container in the Exchange Management Console. Exchange 2007 sp1 moves a couple of settings to the server container.

Selecting the checkbox by "*Block messages sent to recipients not listed in the Global Address List*" can reduce the workload of other anti-spam agents and the server in general. If Exchange Server accepts the message without validating the recipient address, and the address does not exist in the organization, then it will have to submit a Non-Delivery Report (NDR) back to the sending server.

Having the server check the directory for addresses does make it easier for spammers to perform address harvesting. The connecting server can test for a series of generic addresses in the same session. Exchange 2007 allows for a short delay when servers provide the RCPT address before replying with 550 5.1.1 to indicate user unknown. A delay of a few seconds can reduce the value of the connection to spammers while not reducing performance significantly on the Exchange side. This delay is called *SMTP Tarpitting*. It is still effective in some instances but, with the advent of spam bots, connecting hosts show little or no adverse reaction - they just seem to go on unabated. Individual enterprises should monitor Recipient Filtering to gauge its effectiveness.

Finally, The Recipient Filtering layer is also an easy place to prevent users from receiving any external e-mail, by including their address in the Recipient Filter block list. They will still be able to send to the internet. Again, the interface in the EMC is a basic address entry form with an equivalent cmdlet such as:

```
>Set-RecipientFilterConfig -BlockedRecipients address@domain.com
```

SenderID Filter Agent

The anonymous nature of the SMTP protocol allows senders to use incorrect or nonexistent domains in the source address of a message. SenderID is based on the Sender Policy Framework (SPF) and uses the same DNS record format. SPF, originally called "*Sender Permitted From*" is outlined at *www.openspf.org* and RFC 4408. With the SenderID Filter Agent enabled, Exchange Server executes a DNS lookup for an SPF record for the sending SMTP domain. If a published SPF record is available, Exchange will be able to determine if the IP address of the SMTP source is authorized to send e-mail for the sending SMTP domain. This anti-spoofing measure is described in RFC 4406 entitled *Sender ID: Authenticating E-Mail*. As with other message hygiene layers, the administrator can conserve resources by configuring specific exceptions. Messages from specific domains, or those addressed to certain recipients, can be configured to bypass SenderID queries.

SenderID tries to determine the correct sender address to validate against, which is not always obvious. The SenderID Filter Agent will parse the message headers applying a specific algorithm to arrive at the Purported Responsible Address (PRA). This algorithm is defined in RFC 4407 called *Purported Responsible Address in E-Mail Messages.* This also differentiates SenderID from basic SPF.

The SPF record query returns a status. SenderID only deletes or rejects a message when the status returned from the lookup is set to FAIL. This is configurable in the SenderID Properties window in the EMC. No action is taken for other status levels which include PASS, NEUTRAL, SOFT FAIL, TEMP ERROR, and PERM ERROR.

Confused with SPF record formatting?

There are a few resources online to assist in creating a properly formatted SPF record and validating an existing record. Microsoft maintains such a tool, affectionately called Sender ID Framework SPF Record Wizard found at http://www.microsoft.com/mscorp/safety/content/ *technologies/senderid/wizard/.*

Content Filtering Agent

In Exchange Server 2003, Microsoft made available a separate download called the Intelligent Message Filter (IMF) Version 1. The IMF was updated to version 2 and included in Exchange 2003 Service Pack 2. In Exchange Server 2007, the IMF is now referred to as the Content Filter Agent, which could be considered IMF v3.

Content Filtering is engaged after the previous layers have assessed the messages and after the DATA command is fulfilled in the SMTP conversation. Content Filtering is a little bit of a "black box." However, with sufficiently large samples, people have worked backwards to compile reasonable algorithms that may be in play. Microsoft has a great deal of experience with spam. Hotmail.com, msn.com, and Microsoft.com have provided almost unlimited samples from which to create an effective content filter mechanism. Microsoft calls it SmartScreen technology and it forms the basis for content filtering in Outlook as well as Exchange Server. Exchange assesses messages and tries to quantify the likeliness that the message is spam. This measurement is called the Spam Confidence Level (SCL) and is stored as an attribute of the message. Content Filtering assigns a value between 0-9 and records this in an X-header. The higher the SCL, the greater the chances the message is spam. There is also an SCL value of -1 reserved for internal and authenticated messaging.

In the Content Filtering properties window, we can add custom terms to ensure that certain messages are either blocked or allowed to pass. A pharmaceutical distributor may want to whitelist the term Viagra, for example. In addition, recipient addresses can be listed as exceptions in the filtering process, such as *postmaster@domain.com*. New to Exchange 2007 is the ability to whitelist certain senders or domains bypassing content filtering. This is not available in the EMC, but the cmdlets are intuitive. To whitelist an SMTP address, the following can be run from the EMS:

```
>Set-ContentFilterConfig -BypassedSenders address@domain.com
```

Administrators can assign one of three actions to a message, based on its SCL value: delete, reject or quarantine.

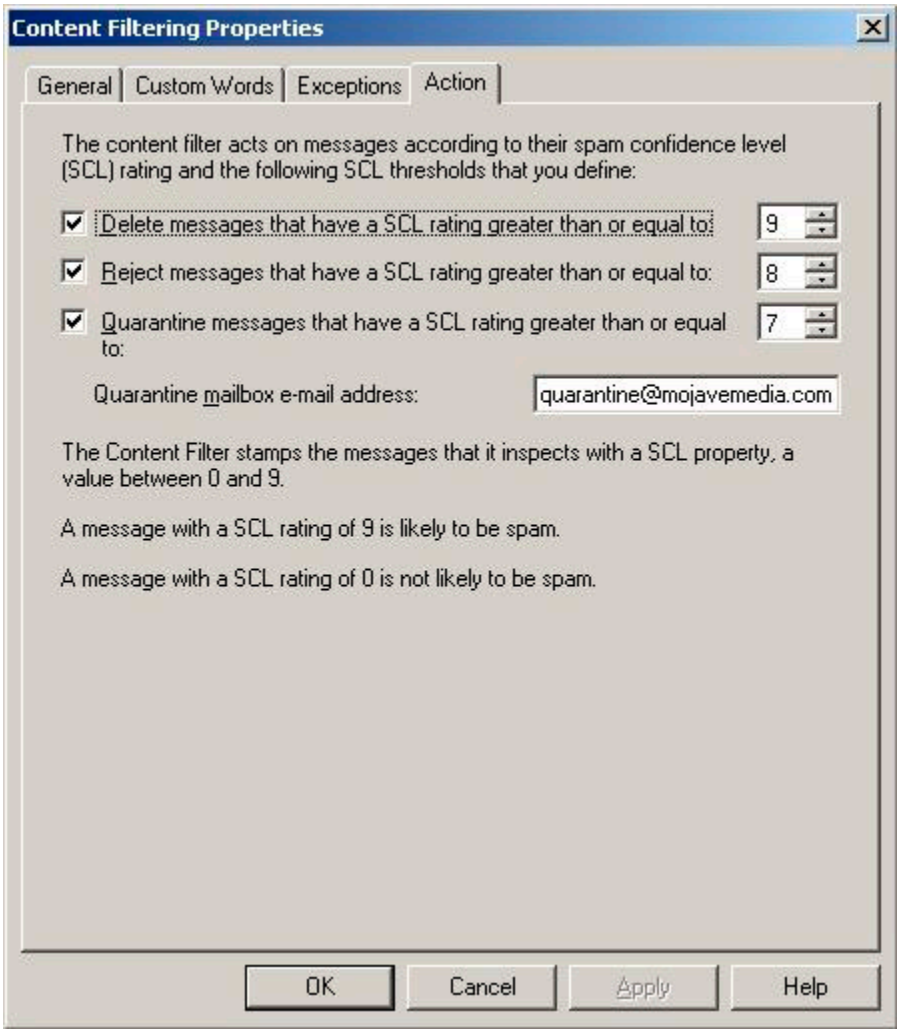

Figure 4.

Figure 4 shows the options in the EMC to assign actions, based on specific minimum SCL ratings. This also shows the current settings on my Edge Server. You will note in Figure 4 that messages with an SCL value of 7 will be quarantined to a separate mailbox. This means that SCL values of 6 or less will be have the rating appended to the header and the message passed through to the Hub Transport and on to the recipient. These settings are not required and different values may be more appropriate in different companies. Administrators should test different threshold values to see what is most effective for their enterprise.

There is also a Junk Mail Threshold for SCL messages. This is configured through the EMS only and simply determines the SCL value at which messages are moved to the Junk Mail folder of the recipient's mailbox.

The final thing I will mention is Sender Reputation. Sender Reputation is a new addition to the message hygiene efforts in Exchange Server. Exchange analyzes characteristics of SMTP sessions of unknown senders. After a sender has sent 20 messages, Exchange maintains a Sender Reputation Level (SRL) analogous to the SCL for content filtering. Sender Reputation, as implemented in Exchange 2007, is a Microsoft creation that is in its infancy and presents potentially compelling options for the future of anti-spam technologies.

Summary

Spam is a moving target. Currently there is no single tool to deploy in the enterprise that will solve this ongoing problem easily. With Exchange 2007, Microsoft provides, out of the box, a multi-faceted defense against UCE. The key to effective spam filtering, using the native Exchange Server 2007 tools, is to determine unwanted messages as early as possible in the message flow. When used together, the tools provide a competitive anti-spam solution at no additional software cost.

Using Exchange 2007 for Resource Booking

16 July 2008

by Nathan Winters

The process of booking various resources to go with a meeting room just got a whole lot easier with Exchange 2007. Nathan Winters explains Exchange's resource-scheduling functionality and offers practical advice on how you might go about configuring Exchange to allow users who are scheduling meetings to specify that such things as video conferencing, projection equipment, and refreshments are required for the meeting.

Introduction

Although previous versions of Exchange were regularly used to host mailboxes that represented resources, Exchange 2007 was the first version in which these mailboxes were actually different from normal mailboxes. This means that there is now no need to rely on add-in software such as the "Auto Accept Agent'. It also means that the design process has been rather more thought through, thereby ensuring a useful solution both at the server and client end.

In Exchange 2007, it is possible to create resource mailboxes that can represent either Rooms or Equipment. These exist as disabled accounts in Active Directory and can be managed in various ways. Exchange 2007 also contains the type of *"Auto Accept Agent"* functionality that was missing from Exchange 2003. Having a mailbox type of "Resource" is useful because it makes administration separate. It allows these mailboxes to have special functionality but yet to be secure as disabled users. It is also helpful to be able to distinguish between rooms and equipment as it is easier for users, and gives more flexibility when adding custom resource properties, as we shall see later.

With the "Availability Service" it is now possible to get real-time information about resource availability without the need to wait for public folder replication to occur (so long as Outlook 2007 is being used). To make best use of this solution you must ensure that Outlook 2007 is being used by those who will be interacting with the resources and managing their use. This is important because much of the user interface, such as the Scheduling Assistant, is held in Outlook 2007. I would however be careful in rolling out Outlook 2007 to the entire user base before thorough testing is carried out. Unless you are careful, you will have dissatisfied users; this is because Outlook 2007 does not always perform as expected, especially with third party add-ons.

Environment

My test environment consists of servers in two sites. Each site has a Mailbox server and a combined Hub-Transport and Client-Access server. Outlook 2007 and OWA 2007 are used as clients. Various users have been set-up in each site such as users, admin users (secretaries), resources (projectors), and rooms. I have configured such a test environment because it is useful to be able to show different sites with a different naming convention. Equally I was interested to see how time zones are handled and discovered that, in my view, this could be done better. For example, as I'll mention later, Outlook is not as flexible as I would like in helping people to arrange meetings in different time zones that require a room in each zone. You could easily test the vast majority of this functionality on a single Exchange server with a single Outlook client.

Creating or Migrating resource mailboxes

The resource mailboxes can be created from the GUI or from the command line as shown below:

```
New-Mailbox -database "Siteamb1\SG1\DB1" -Name SiteA-Room1 -
OrganizationalUnit "SiteA Rooms" -DisplayName "SiteA-Room1" -
UserPrincipalName SiteA-Room1@gaots.co.uk -Room
```

Note - If you are migrating the resource mailboxes from an Exchange 2003 system, then use one of the following commands to convert the migrated mailbox to a resource, either "room" or "equipment" as you wish.

```
Set-Mailbox -Identity Name -Type Room
Set-Mailbox -Identity Name -Type Equipment
```

Before moving on, the resource mailboxes must be set-up to "auto accept" meeting requests; otherwise they will always remain as "tentative" until an operator intervenes. This is done as follows:

```
Get-Mailbox | where {$_.ResourceType -eq "Equipment"} | Set-MailboxCalendarSettings
-AutomateProcessing:AutoAccept

Get-Mailbox | where {$_.ResourceType -eq "Room"} | Set-
MailboxCalendarSettings -AutomateProcessing:AutoAccept
```

Once suitably configured as above, each resource mailbox will send an auto-accept message to meeting organizers (Figure 1). The message text can be customized to suit requirements.

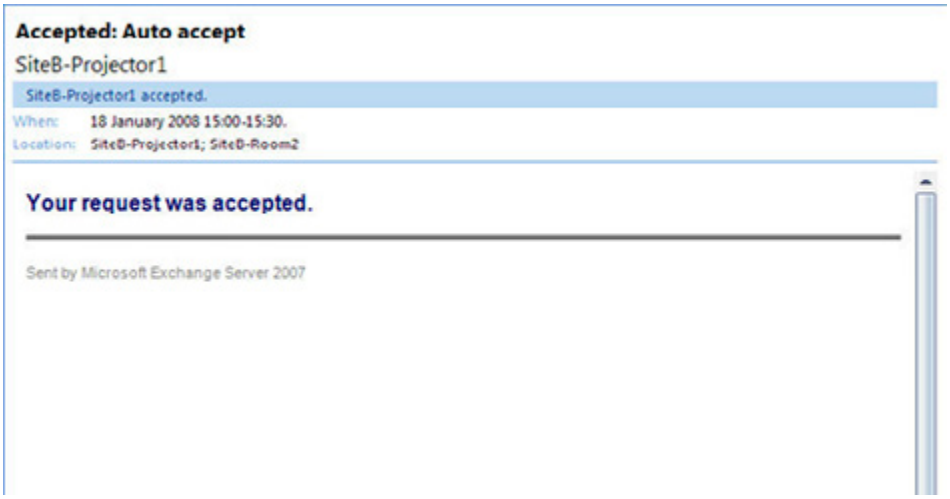

Figure 1 – Showing auto acceptance. This message can be customized.

Feature Investigation

I have now discussed how to create resource mailboxes and how to configure their basic option of responding by accepting requests. As you saw, the initial setup is really fairly simple; it is just important to remember that, when migrating from 2003, you must convert the mailbox to the new resource type. Now, let's move on and look at some of the key resource-centric features that Exchange 2007 offers.

Security and Booking Policies

Both security configuration and booking policies allow the administrator to grant control of resource mailboxes to certain users: Those users (using Outlook Web Access) can then define how they are used. For example, it is possible to specify the maximum meeting duration, to define and allow meetings to be scheduled only during working hours and to define how to handle conflicts.

The first step is to setup the resource-admin users as "delegates" of the specific resources in each site. The command below will get mailboxes with the type Equipment or Room and with "**sitea**" in the name. It will then give Delegate access to the account "**siteaadmin**." I have repeated the steps for **siteb**:

```
Get-Mailbox -RecipientTypeDetails EquipmentMailbox, RoomMailbox | Where {$_.Name
-match "^sitea"} | Set-MailboxCalendarSettings -ResourceDelegates siteaadmin

Get-Mailbox -RecipientTypeDetails EquipmentMailbox, RoomMailbox | Where {$_.Name -match "^siteb"}
| Set-MailboxCalendarSettings -ResourceDelegates sitebadmin
```

The "delegate user" for each resource mailbox may be sent a copy of the meeting requests which go to that resource depending on policy, and they are then able to accept or reject these requests.

Next set-up the admin accounts in **sitea** and **siteb** to have full access to their respective resources. The command below is for **siteb:**

```
Get-Mailbox -RecipientTypeDetails EquipmentMailbox, RoomMailbox | Where {$_.Name -match "^siteb"}
| Add-MailboxPermission -AccessRights FullAccess -User sitebadmin
```

With Full Admin permissions on the resource mailbox, the admin user can then open the mailbox as an additional mailbox in Outlook and have full control of all functions. Most important of these is the OWA Options/Resource Settings section of a resource mailbox: This is used to configure some of the Resource options such as; "Resource Scheduling Options," "Resource Scheduling Permissions," "Resource Privacy Options," and "Response Message." Every option that is available via the shell is also available in OWA as shown in Figure 2.

To access the relevant mailbox in OWA you will need to access a specific URL:

http://servername/owa/ResourceMailbox@domain.com

... and then log in with the account which has Full Access rights. Click on the Options button in the top right hand corner and then, in the left hand pane, click on the Resource Settings section.

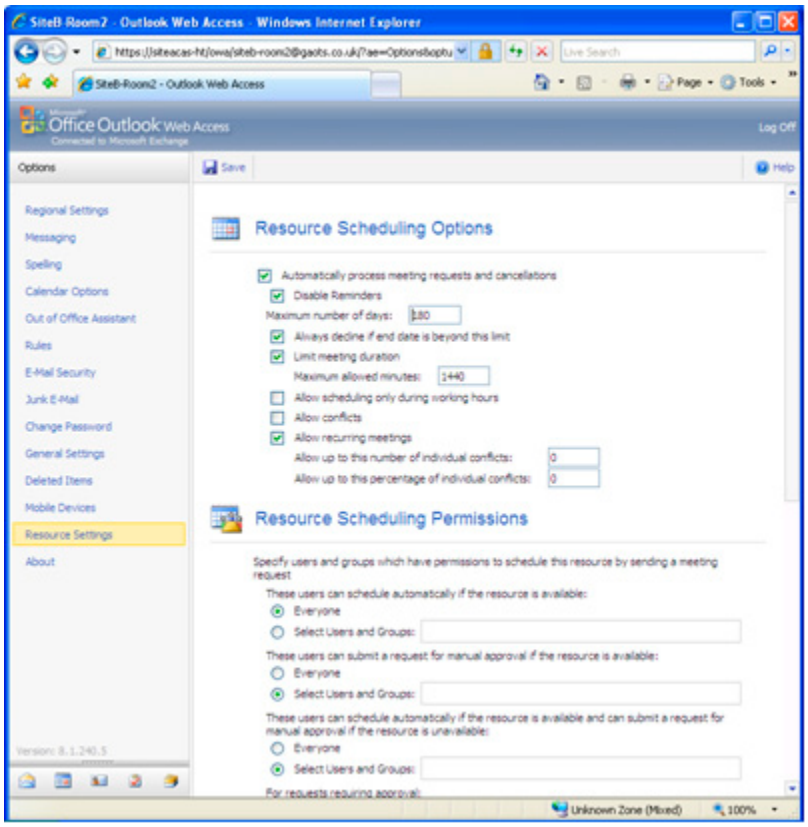

Figure 2 – The Resource Settings page in OWA options.

Once you have reached the Resource Settings section of the mailbox options in OWA, you can see the policies that are available. There are three policies, BookInPolicy, RequestInPolicy, and RequestOutOfPolicy.

- **BookInPolicy** is a list of users who can schedule the resource automatically. By default it is set to Everyone.
- **RequestInPolicy** is a list of users who can request to schedule to resource and have to have it authorized by the resource administrator.
- **RequestOutOfPolicy** is a list of users who can submit a request for a resource even if it is already booked. (This might for example be used for the CEO!)

This page enables you to allow conflicts, set maximum meeting length, allow recurring meetings, customize the text of notification mails and a host of other settings.

A policy example

So what if you want to set a policy that makes it mandatory for a delegate to authorise the request when anyone books catering?

This can be done by first setting up the delegate for the catering resource mailbox as above. Then log into OWA (also as above) and set the "These users can schedule automatically if the resource is available:" policy to "Select Users and Groups." On the "These users can submit a request for manual approval if the resource is available:" policy to "Everyone." Save the settings.

At this point, when anyone tries to submit a request, the delegate user will get a mail like the one in **Figure 3.**

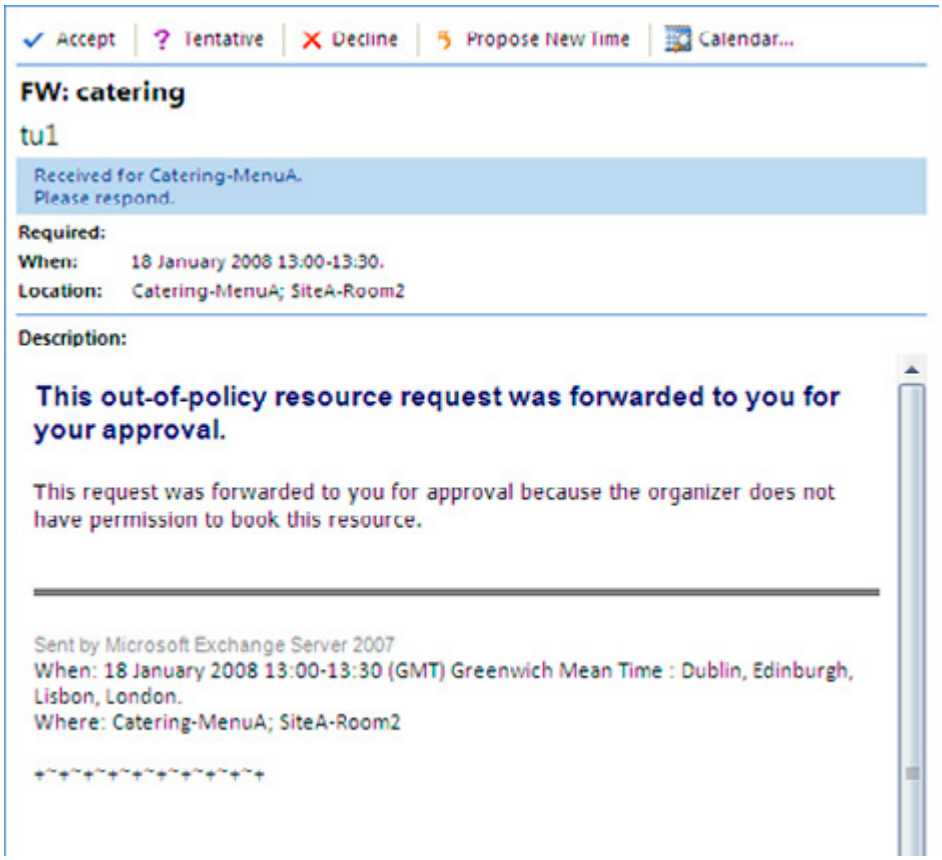

Figure 3 – The Request info forwarded to delegates.

The delegate user can then accept the request (and go ahead and book the catering with the caterer!).

Meanwhile the requesting user will get a mail as in Figure 4, showing the pending request. The meeting is held as tentative to prevent the resource being booked by someone else.

Figure 4 – The Tentative mail received by the user whilst waiting approval.

Once the delegate user approves the request, the meeting will be booked and acceptances sent out.

More info about the policies and the shell commands available can be found here:

http://technet.microsoft.com/en-us/library/aa996340.aspx

http://technet.microsoft.com/en-us/library/bb124987(EXCHG.80).aspx

Viewing Meeting rooms

So what does this look like from the client perspective? In Outlook 2007 it is possible to see all rooms, and their availability, in one window by using the Scheduling Assistant in Outlook 2007 as shown in **Figure 5.**

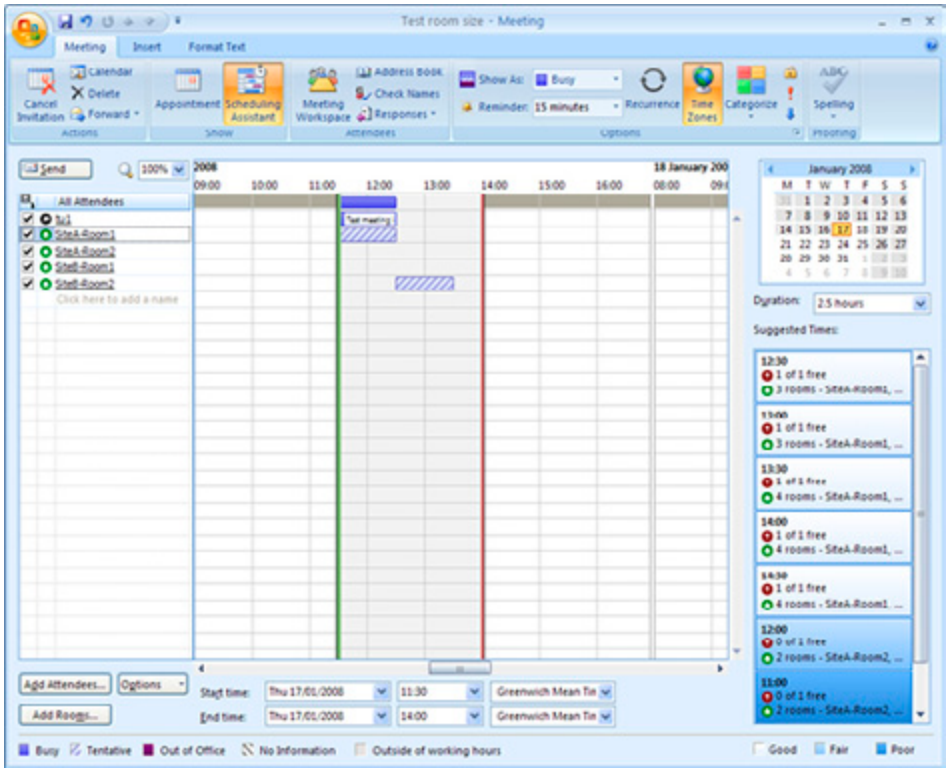

Figure 5 – The Scheduling Assistant.

All rooms can be added from the "Add Rooms" button. In the main window you will see which rooms are available: this is supplemented by the "Suggested Times" area on the right hand side.

Although it isn't obvious, it is possible to set the meeting time zone when creating a meeting using the Time Zones button highlighted in Figure 5. However, it is not possible to view and additional time zone in the scheduling assistant as you can in the main Outlook calendar.

Scheduling resources

The Scheduling tab is where most bookings will take place. The basic scheduling of resources is simple. First open a meeting request and enter a subject and if necessary, the attendees. Next move to the Scheduling tab. Use the new "Add Rooms" button to open the window shown in Figure 6. This shows the "All Rooms" address list which gives a view of the capabilities of each room. As you can see, "SiteB-Room2" is listed with a capacity of 2 which is a default property that can be customized on each room, using the Exchange Management Console as shown in Figure 7.

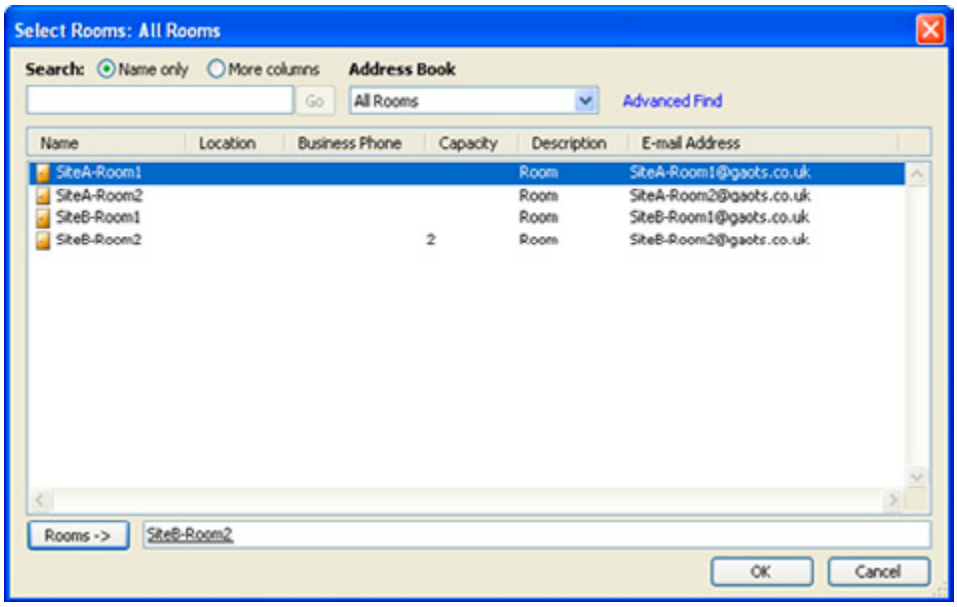

Figure 6 – The All Rooms address list.

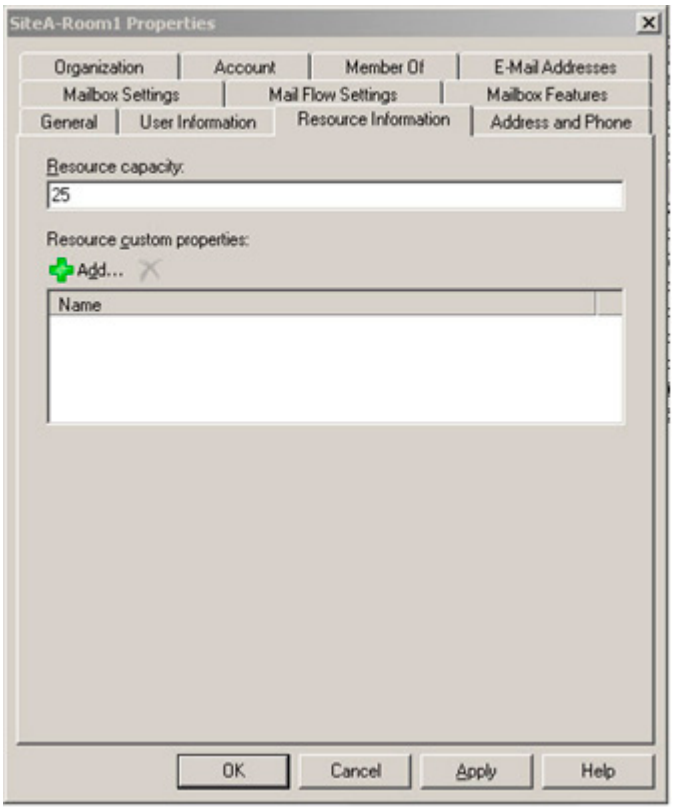

Figure 7 – Setting the capacity of the Room.

If you are searching for an available room, then select those with the correct properties (such as capacity), and add them.

When back on the Scheduling tab, you can then use the "attendees" tab to add other attendees or resources such as projectors. **Figure 8** below shows the booking facilities.

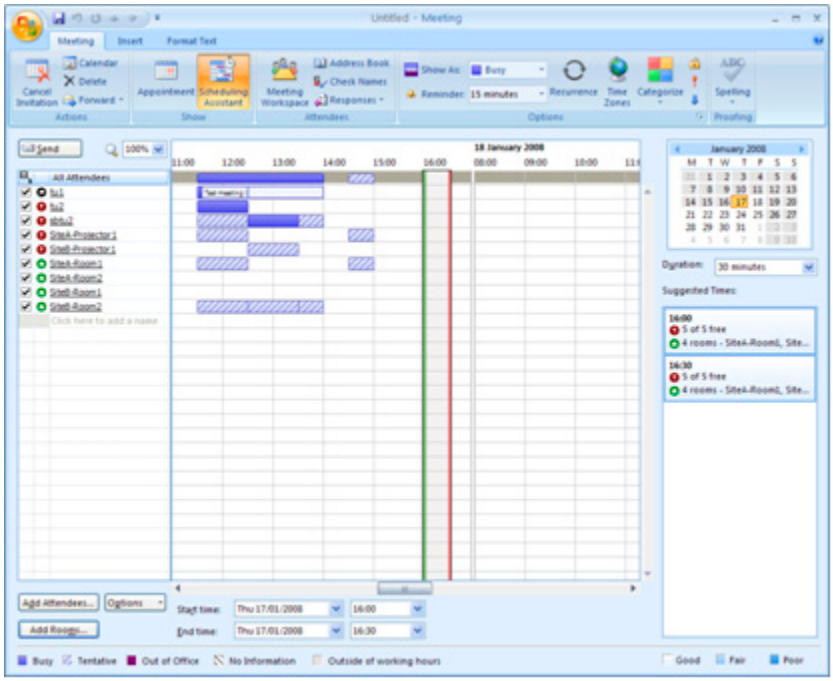

Figure 8 – Scheduling tab showing users, resources, and rooms.

Custom Resource Properties

As I mentioned earlier it is possible to add custom properties to resource mailboxes. If, for example, you have a catering menu and you want to use a resource mailbox to identify the menu code, then firstly create the custom resource types as follows.

First read the current resource configuration and store it in a temporary variable called **$ResourceConfiguration** by running the following command:

```
$ResourceConfiguration = Get-ResourceConfig
```

Next create your custom properties, in this case sandwich types - Fish, Vegetarian, Meat:

```
$ResourceConfiguration.ResourcePropertySchema+=("Equipment/Fish")
$ResourceConfiguration.ResourcePropertySchema+=("Equipment/Vegetarian")
$ResourceConfiguration.ResourcePropertySchema+=("Equipment/Meat")
```

Finally update the resource configuration of your organization by using the modified resource property schema, using the following command:

```
Set-ResourceConfig -Instance $ResourceConfiguration
```

Once you have created the custom resource properties, you will then add them to the relevant resource mailbox: this can be done in Exchange Management Console. The example shown in **Figure 9** shows two catering menus:

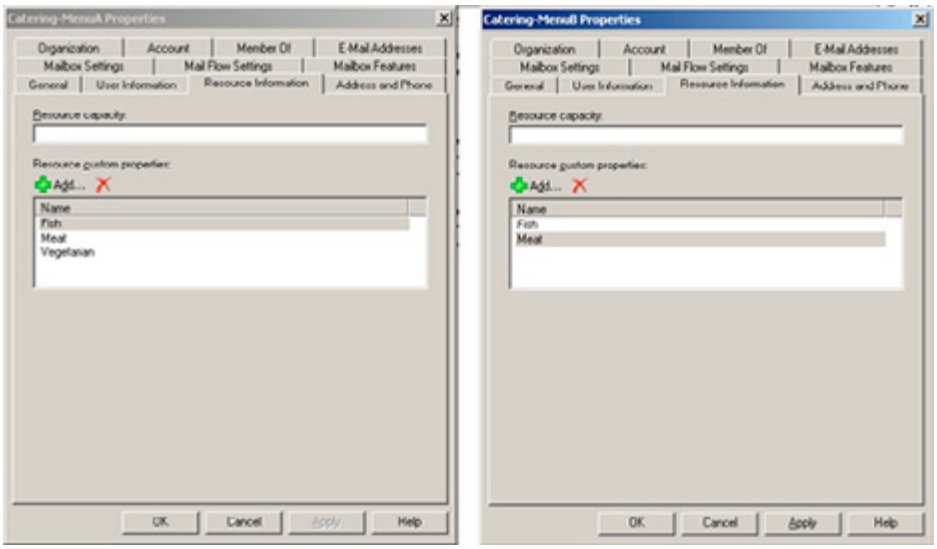

Figure 9 – Showing the method of adding custom resources.

Having done that, one way to view these extended attributes is by using the PowerShell command below:

Get-Mailbox -RecipientTypeDetails RoomMailbox |fl Name,ResourceCustom

However, with a little more investigation, I have found another more user-friendly way!

As it would appear that only the "All Rooms" address list has the right GUI part to show the "Description" field, I have edited the filter to include the EquipmentMailbox type as well.

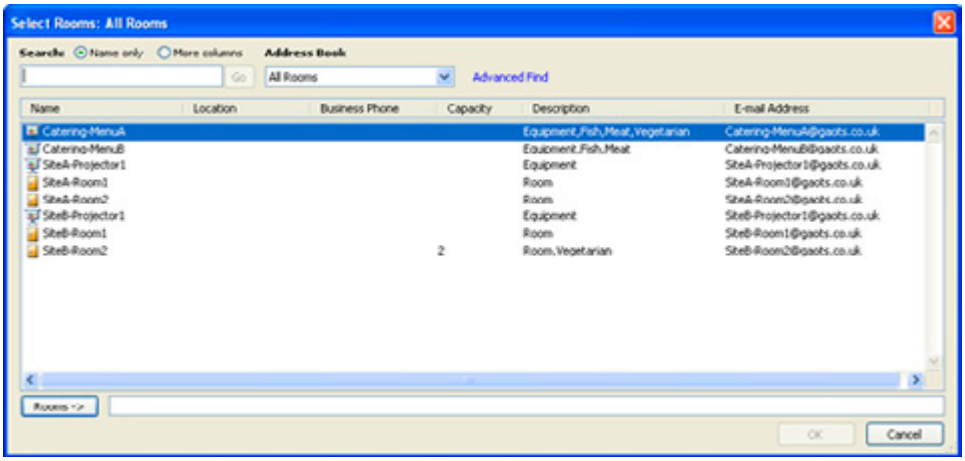

Figure 10 – The combined Rooms and Equipment resources showing custom properties.

I used this command to edit the filter.

```
Set-AddressList -Identity "All Rooms" -RecipientFilter {(Alias -ne $null -and
(RecipientDisplayType -eq 'ConferenceRoomMailbox' -or RecipientDisplayType -eq
'SyncedConferenceRoomMailbox' -or RecipientDisplayType -eq 'EquipmentMailbox'))}
```

The results of the filter change are shown in **Figure 10.**

Summary

This article describes the Exchange Resource Scheduling functionality. I realize I haven't covered every single feature, but I have aimed at doing enough to give you the skills needed to continue developing your resource booking solution using Exchange 2007.

In general, with a little investigation and some color-coding in Outlook, you can create a decent resource booking system. Of course it is not as polished as some of the 3rd-party systems, but then a lot of them are not that great!

Some improvements are necessary in the user-interface. For example, when you are adding custom resource properties you can't actually see them in Outlook unless you are modifying the Rooms address list. It should be possible to use the Exchange Web Services to create a custom front end for working with the "resource mailboxes" that you create. Let's hope someone does it soon.

Controlling Email Messages using Exchange's Transport Rules

22 July 2008

by William Lefkovics

Some tasks that should have been easy in previous versions of Exchange just weren't. Now, with Transport Rules in Exchange 2007, administrators have a much improved set of tools to manage message flow.

Microsoft Exchange Server 2007: Controlling Email Messages using Exchange's Transport Rules

As a long time Exchange administrator, one of the great frustrations I have found over the years has been how difficult certain seemingly trivial administrative tasks are to implement in Exchange Server. It is almost embarrassing to tell management that it is not that simple to add a disclaimer to outbound SMTP messages or insert text into the subject of certain messages or control message flow for compliance requirements.

With Exchange Server 2007 Microsoft addresses many of these deficiencies by completely changing the underlying architecture. SMTP has been brought back into Exchange instead of extending the SMTP service in Internet Information Server (IIS). Arduous transport event sinks have been replaced by Transport Agents on the Hub Transport and Edge Transport roles. Transport rules are implemented in the Exchange 2007 architecture through Transport Agents. There is a Transport Rules Agent on all Hub Transport Servers and an Edge Transport Agent on all Edge Servers. Administrators now have a basic UI to control messages using Transport Rules - one of the killer features of the latest Exchange Server.

Focus of Transport Server Roles

Edge Transport Servers perform gateway services for an Exchange organization. They are somewhat independent sentries in the perimeter of corporate networks. As such, their function in terms of transport rules focuses on message hygiene and security. Edge servers are not domain members and have no direct access to Active Directory.

Hub Transport Servers, on the other hand, are integrated in the Windows domain infrastructure and have access to Active Directory. They handle messages that remain internal to the organization in addition to content arriving from or departing to an Edge Transport server. The focus of transport rules implemented on the Hub Transport role is geared toward policy enforcement and compliance.

Because of the different focus of transport rules between the Edge and Hub Transports, the actual set of rules varies between them.

Scope of Transport Server Roles

Edge Transport servers work alone. Transport rules on the Edge are stored in a somewhat portable subset of Active Directory called Active Directory Application Mode (ADAM). A special updating mechanism called EdgeSync is used to keep ADAM fairly current for Edge with user information from AD. If there are multiple Edge servers in place, they do not share their instance of ADAM. Any updates must be performed separately to each Edge server. Different Edge Transport servers may control different connections to the Exchange organization or they may be clones of one to serve as redundant gateways for a single connection. Either way, Edge servers are not aware of each other, they are not members of the internal Exchange organization or Windows domain, and they operate independently.

By contrast, every single e-mail message sent in an Exchange 2007 organization must pass through at least one Hub Transport server. Even if the sender and recipient reside in the same database, the message leaves the store and passes through the transport pipeline on a hub transport server before returning to the mailbox server. Transport Rules are stored in Active Directory. This means (and this is important) that every Hub Transport server accesses the same set of transport rules. Messages sent through Exchange 2007 can not bypass transport rule processing! Historically, this has been a significant obstacle for administering an Exchange messaging infrastructure that meets regulatory compliance initiatives.

Write your own Transport Agent! What do you need? All that is required to compose and implement a custom Transport Agent is a server with either the Exchange 2007 Edge Transport Role or the Hub Transport Role installed. The Microsoft .Net Framework 2.0 Software Development Kit (SDK) and a suitable IDE, such as Visual Studio 2005 or 2008 (to compile the agent to a DLL) are also needed. Microsoft provides a set of classes through the .Net Framework extensions to programmatically access and even change SMTP message properties at various events through the transport pipeline. A little C# or Visual Basic programming skills does help.

For more information on developing Transport Agents see the Microsoft Exchange Server 2007 SP1 SDK or MSDN.

What types of messages do Transport Rules work against?

Almost every type of message that travels through the hub goes through transport rule processing. Standard e-mail messages with plain text, HTML, or RTF are all accessible by the transport rule agent. Transport rules do work for digitally signed messages and encrypted or opaque messages as well, but only aspects that it can access. A rule can still read the message header even if the message body has been encrypted.

Exchange 2007 Service Pack 1 added transport rule support for IPM.Note formatted messages, as you might see from applications that generate e-mail messages, as well as unified messaging e-mails, including voice mail messages, fax messages, and missed call notification messages.

Anatomy of a Transport Rule

Help Microsoft improve Transport Rules!

Ben Neuwirth at Microsoft recently posted a blog entry publishing a script that can be run against your transport servers to return a statistical analysis outlining which predicates and actions you use most. The script does not collect any personal data and you can review it before emailing it to Ben. The entry is found HERE.

Transport Rules are not all that different from Outlook Rules in their logic. Each rule is composed of at least one condition, an Action or Actions and optionally, one or more Exceptions. If no conditions are selected, then the rule will apply to ALL messages. Where there are multiple conditions, they all must be met; however, no matter how many exceptions there are, it only takes one to prevent the rule from firing.

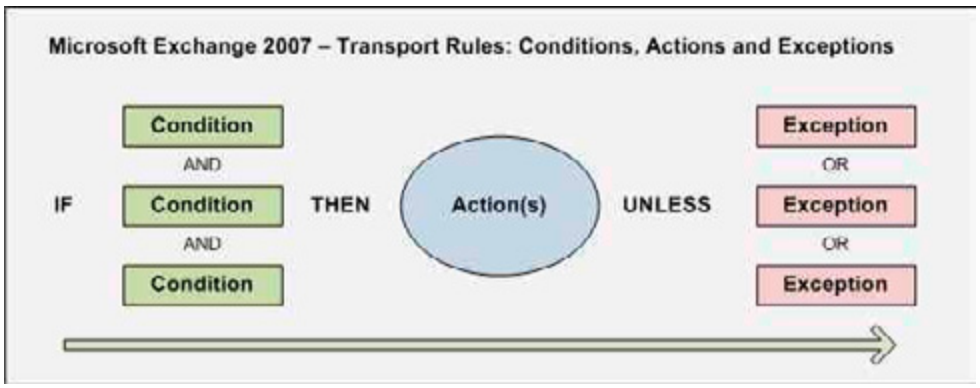

Figure 1.

Figure 1 shows this logical flow through a transport rule. Transport rules are quite flexible with a solid set of options.

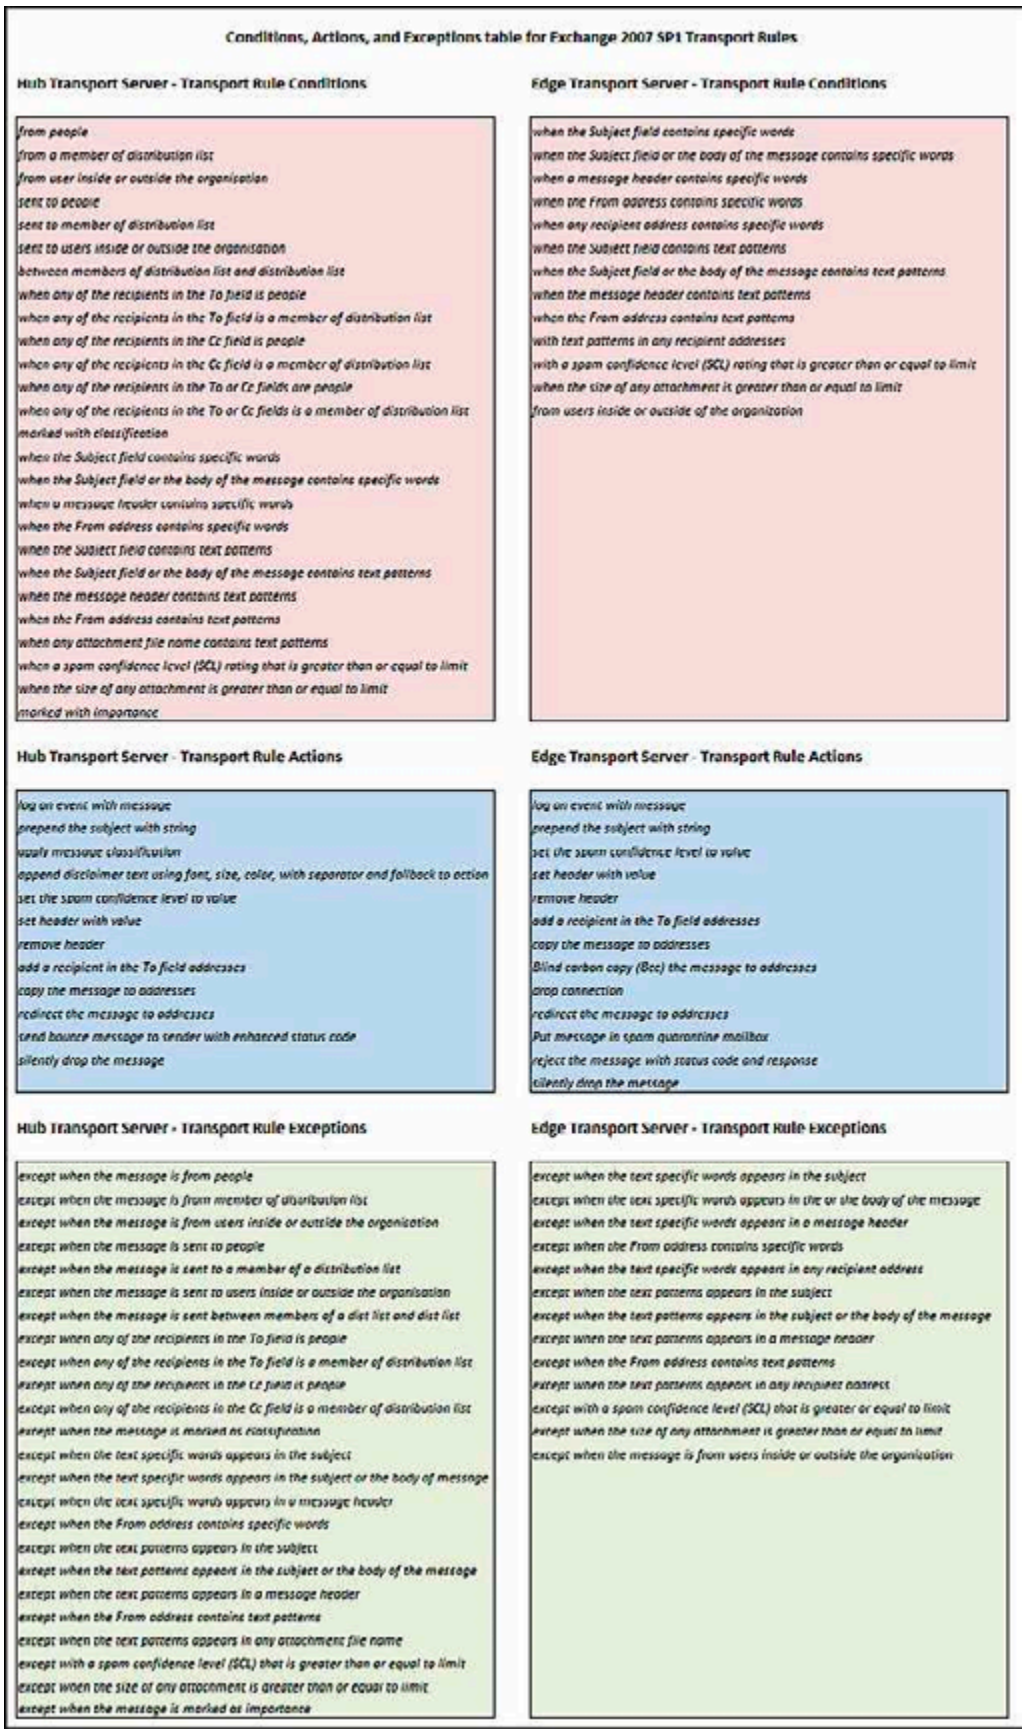

Conditions, Actions, and Exceptions table for Exchange 2007 SP1 Transport Rules

Hub Transport Server - Transport Rule Conditions

- from people
- from a member of distribution list
- from user inside or outside the organisation
- sent to people
- sent to member of distribution list
- sent to users inside or outside the organisation
- between members of distribution list and distribution list
- when any of the recipients in the To field is people
- when any of the recipients in the To field is a member of distribution list
- when any of the recipients in the Cc field is people
- when any of the recipients in the Cc field is a member of distribution list
- when any of the recipients in the To or Cc fields are people
- when any of the recipients in the To or Cc fields is a member of distribution list
- marked with classification
- when the Subject field contains specific words
- when the Subject field or the body of the message contains specific words
- when a message header contains specific words
- when the From address contains specific words
- when the subject field contains text patterns
- when the Subject field or the body of the message contains text patterns
- when the message header contains text patterns
- when the From address contains text patterns
- when any attachment file name contains text patterns
- when a spam confidence level (SCL) rating that is greater than or equal to limit
- when the size of any attachment is greater than or equal to limit
- marked with importance

Edge Transport Server - Transport Rule Conditions

- when the Subject field contains specific words
- when the Subject field or the body of the message contains specific words
- when a message header contains specific words
- when the From address contains specific words
- when any recipient address contains specific words
- when the Subject field contains text patterns
- when the Subject field or the body of the message contains text patterns
- when the message header contains text patterns
- when the From address contains text patterns
- with text patterns in any recipient addresses
- with a spam confidence level (SCL) rating that is greater than or equal to limit
- when the size of any attachment it greater than or equal to limit
- from users inside or outside of the organization

Hub Transport Server - Transport Rule Actions

- log an event with message
- prepend the subject with string
- apply message classification
- append disclaimer text using font, size, color, with separator and fallback to action
- set the spam confidence level to value
- set header with value
- remove header
- add a recipient in the To field addresses
- copy the message to addresses
- redirect the message to addresses
- send bounce message to sender with enhanced status code
- silently drop the message

Edge Transport Server - Transport Rule Actions

- log an event with message
- prepend the subject with string
- set the spam confidence level to value
- set header with value
- remove header
- add a recipient in the To field addresses
- copy the message to addresses
- Blind carbon copy (Bcc) the message to addresses
- drop connection
- redirect the message to addresses
- Put message in spam quarantine mailbox
- reject the message with status code and response
- silently drop the message

Hub Transport Server - Transport Rule Exceptions

- except when the message is from people
- except when the message is from member of distribution list
- except when the message is from users inside or outside the organization
- except when the message is sent to people
- except when the message is sent to a member of a distribution list
- except when the message is sent to users inside or outside the organisation
- except when the message is sent between members of a dist list and dist list
- except when any of the recipients in the To field is people
- except when any of the recipients in the To field is a member of distribution list
- except when any of the recipients in the Cc field is people
- except when any of the recipients in the Cc field is a member of distribution list
- except when the message is marked as classification
- except when the text specific words appears in the subject
- except when the text specific words appears in the subject or the body of message
- except when the text specific words appears in a message header
- except when the From address contains specific words
- except when the text patterns appears in the subject
- except when the text patterns appears in the subject or the body of the message
- except when the text patterns appears in a message header
- except when the From address contains text patterns
- except when the text patterns appears in any attachment file name
- except with a spam confidence level (SCL) that is greater than or equal to limit
- except when the size of any attachment is greater than or equal to limit
- except when the message is marked as importance

Edge Transport Server - Transport Rule Exceptions

- except when the text specific words appears in the subject
- except when the text specific words appears in the or the body of the message
- except when the text specific words appears in a message header
- except when the From address contains specific words
- except when the text specific words appears in any recipient address
- except when the text patterns appears in the subject
- except when the text patterns appears in the subject or the body of the message
- except when the text patterns appears in a message header
- except when the From address contains text patterns
- except when the text patterns appears in any recipient address
- except with a spam confidence level (SCL) that is greater or equal to limit
- except when the size of any attachment is greater than or equal to limit
- except when the message is from users inside or outside the organization

Figure 2.

Figure 2 shows the various Conditions, Actions, and Exceptions for the Edge and Hub Transport Servers. When these predicates and actions are selected, there are variables to include, such as text, addresses, and other properties.

Unfortunately, you can not add your own actions or predicates to Microsoft's transport rules interface. You can develop your own custom transport agent to fulfil such a need, of course.

Regulatory Compliance

Many companies these days are required to assert greater control over their messaging solutions. Regulatory agencies in various countries demand certain e-mail communications be archived and discoverable. In addition, corporate policy may be designed to minimize liability exposure by providing employees with a working environment safe from harassment or loss of productivity through electronic communications. Transport Rules in Exchange 2007 provide rudimentary solutions to assist administrators in deploying effectively compliant messaging systems. For example, archiving sensitive information may be required in some jurisdictions. Who wants to be investigated by the UK Information Commissioner's Office or the US Securities and Exchange Commission and not be able to provide the content they require?

Ethical Walls

Since every message must pass through a hub transport server, the hub becomes the point of message control for policy enforcement and compliance. Transport rules can fire on messages where senders or recipients belong to specific groups. This makes it easy to allow or prevent mail flow based on universal distribution group membership. A transport rule can prevent confidential email sent, either intentionally or accidentally, from the Finance department to a factory worker simply by restricting delivery of e-mail between those groups. This virtual blockade of e-mail communication between groups is referred to as an Ethical Firewall or Ethical Wall. A policy may be put in place to have CFO emails, which are often of a sensitive nature, blocked from being sent to factory workers. In the rare case where the CFO needs to send something, then perhaps the HR department can send that e-mail instead. Every company is different and Transport Rules provide some flexibility for securing the flow of e-mail for diverse scenarios. Ethical walls using Transport Rules reduce the potential for confidential information from getting into the wrong inbox.

Message Classifications

> ***Note***
>
> *Message Classifications are not available to Outlook clients by default. Outlook 2007 clients require Message Classifications to be exported from AD and copied to the local registry on the workstation for Outlook to access. This manual process allows administrative control on who can apply classifications to messages.*

With Exchange 2007 there is a special type of server-side message categorization called Message Classification, usable with OWA 2007 and Outlook 2007 clients. These are custom labels stored in Active Directory that can be applied to e-mail messages. Transport rules can either act upon messages with specific classifications or can assign a message classification to messages based on specific properties. Exchange 2007 actually has a few sample message classifications by default. These are not accessible through the EMC; however, they are fairly easily managed using the EMS.

To create a list of Message Classifications using the EMS type the following cmdlet:

```
[PS]C:\>Get-MessageClassification | ft
```

This will generate a simple table as shown in Figure 3 where you can also identify a new Message Classification we added for Simple Talk articles.

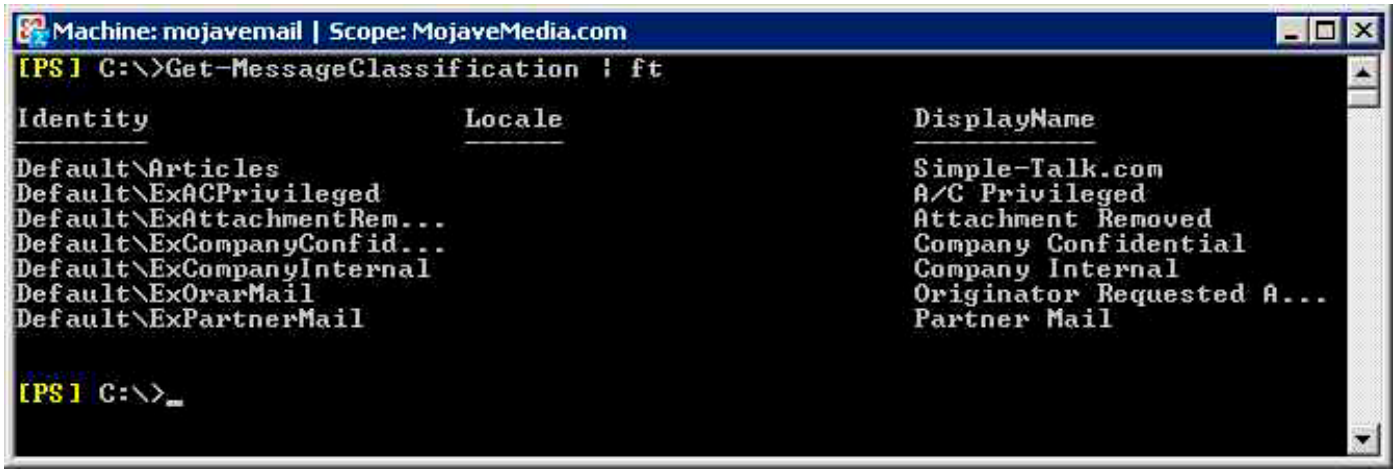

Figure 3.

Working with Message Classifications goes beyond the scope of this article. More information can be found at the source of course: http://technet.microsoft.com/en-us/library/aa998271(EXCHG.80).aspx.

Creating a new Transport Rule

What permissions are needed to create Transport rules?

Well just to see the transport rules, the administrator must be delegated at least the Exchange View-Only Administrator role. To create or modify existing transport rules, the administrator must have the Exchange Organization Administrator role.

As you probably already know, the Exchange Management Console (EMC) was built upon the Exchange Management Shell (EMS). Each action performed from the EMC holds an equivalent EMS cmdlet. Transport rules can be managed from either the EMC or the EMS. We will look at both options.

Whether you are creating a new transport rule on a Hub server or an Edge server, the process is very similar. We will walk through an example using the Hub Transport server where messages are copied to another mailbox based on keywords in the subject. On an internal Exchange 2007 server, the EMC has a few containers outlining menu options based on scope. Hub Transport rules are stored in Active Directory in the Exchange Configuration container, so they are replicated throughout the entire AD forest. Logically, Transport rules are thus managed using the Hub Transport option under the Organization container in the EMC as shown in the left pane in Figure 4.

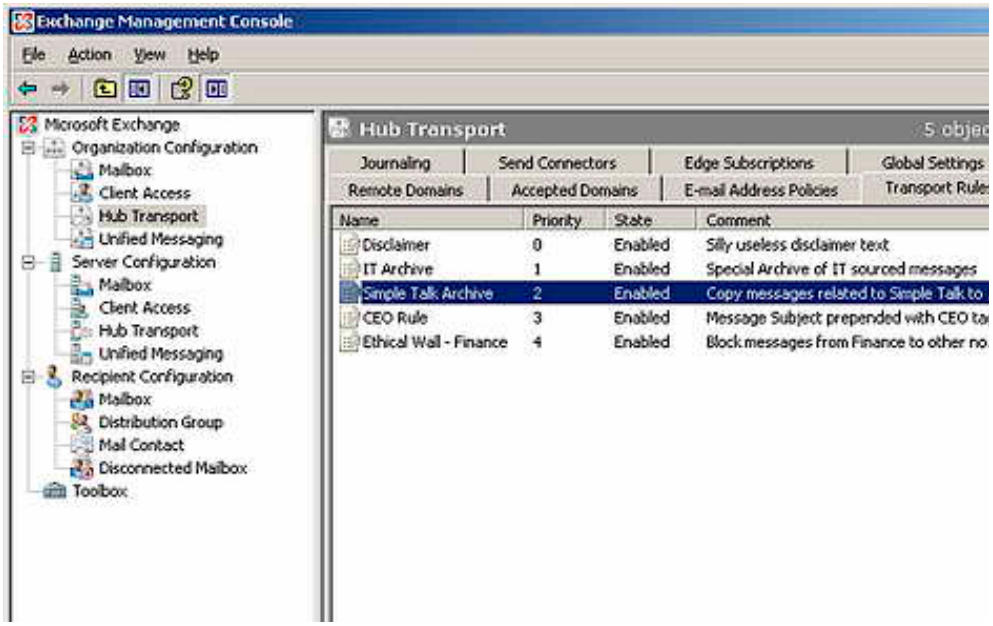

Figure 4.

Also in Figure 4, you can see the Transport Rules tab in the center pane is selected.

To launch the wizard, click on the New Transport Rule option in the Action pane of the EMC.

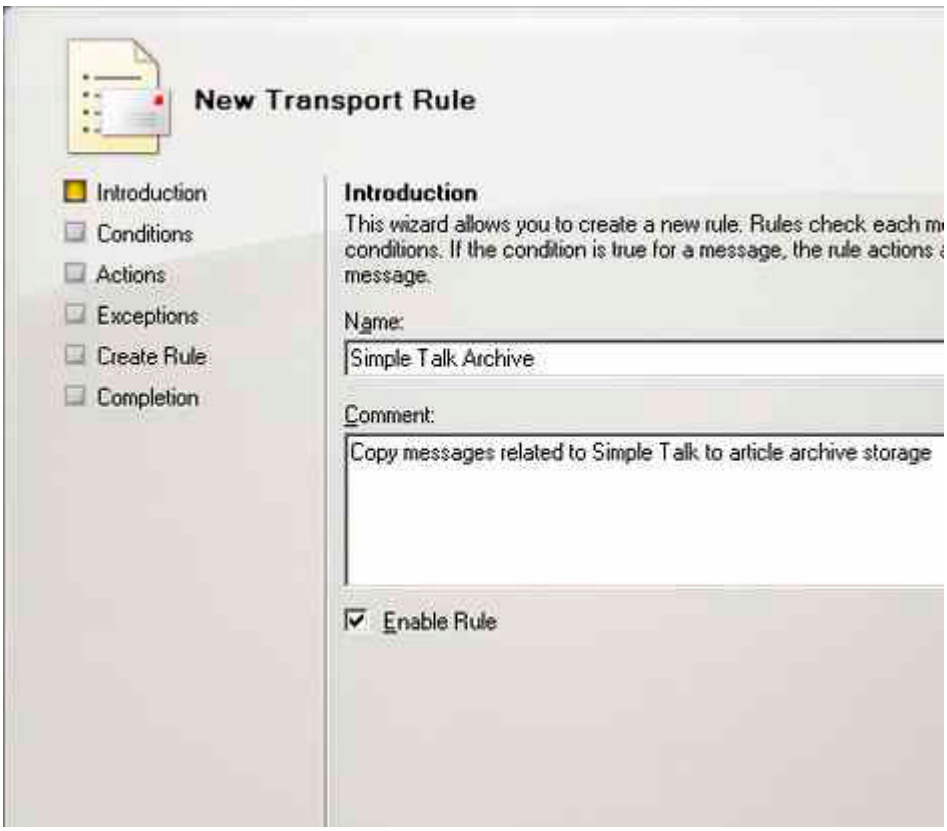

Figure 5.

Figure 5 shows the initial screen requiring a name for the rule. The description field is informational and optional and displayed in the EMC. The wizard walks through the conditions, actions, and exceptions for the new transport rule. For our example, the property that triggers the rule is the presence of the keywords "Red Gate" or "Simple Talk" in the message subject as shown in Figure 6.

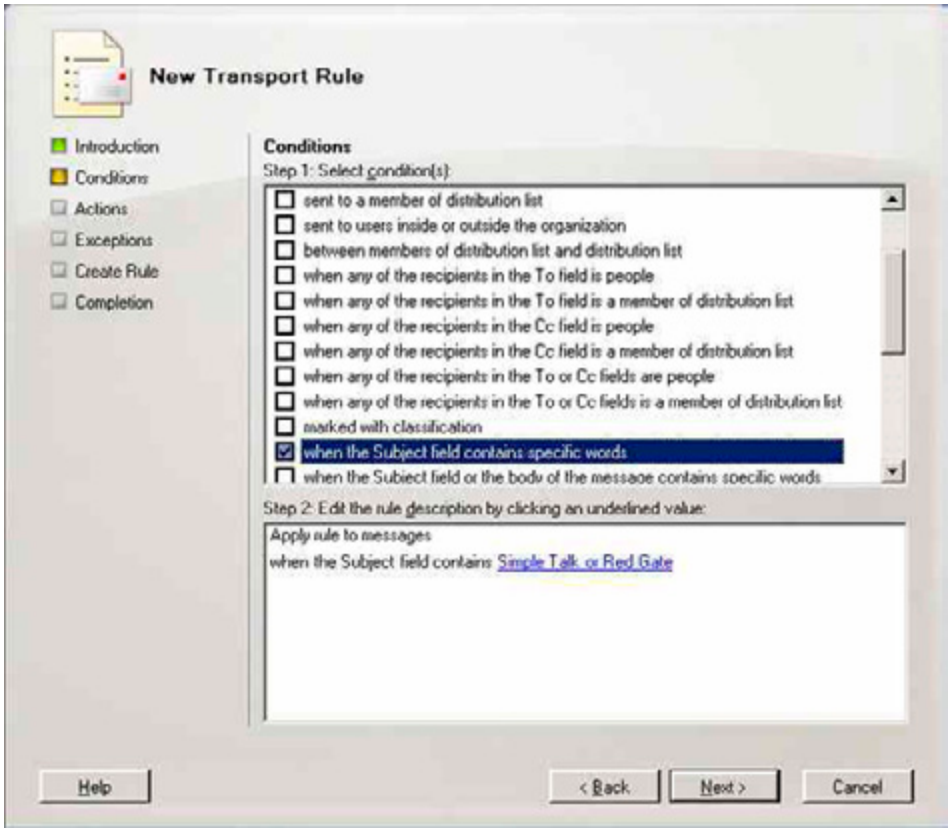

Figure 6.

This condition will result in the message being copied to an article archive mailbox (see Figure 7) unless the message has been tagged with the Message Classification "Company Confidential" (see Figure 8).

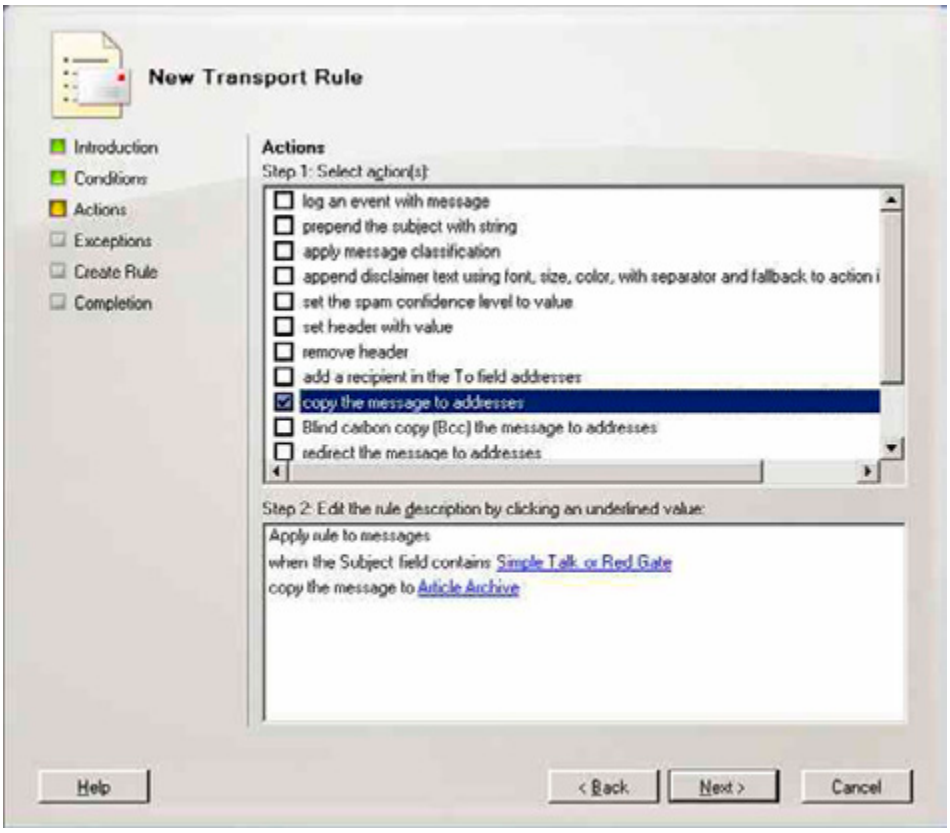

Figure 7.

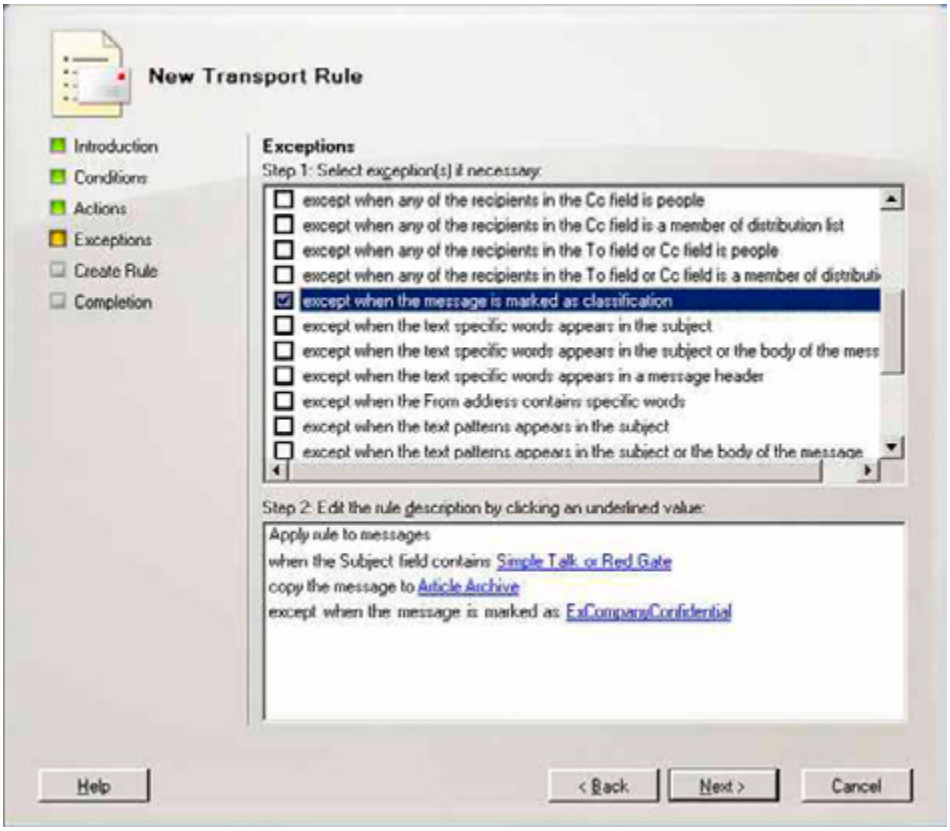

Figure 8.

The next window in the new rule wizard is a confirmation of what has been entered (see Figure 9).

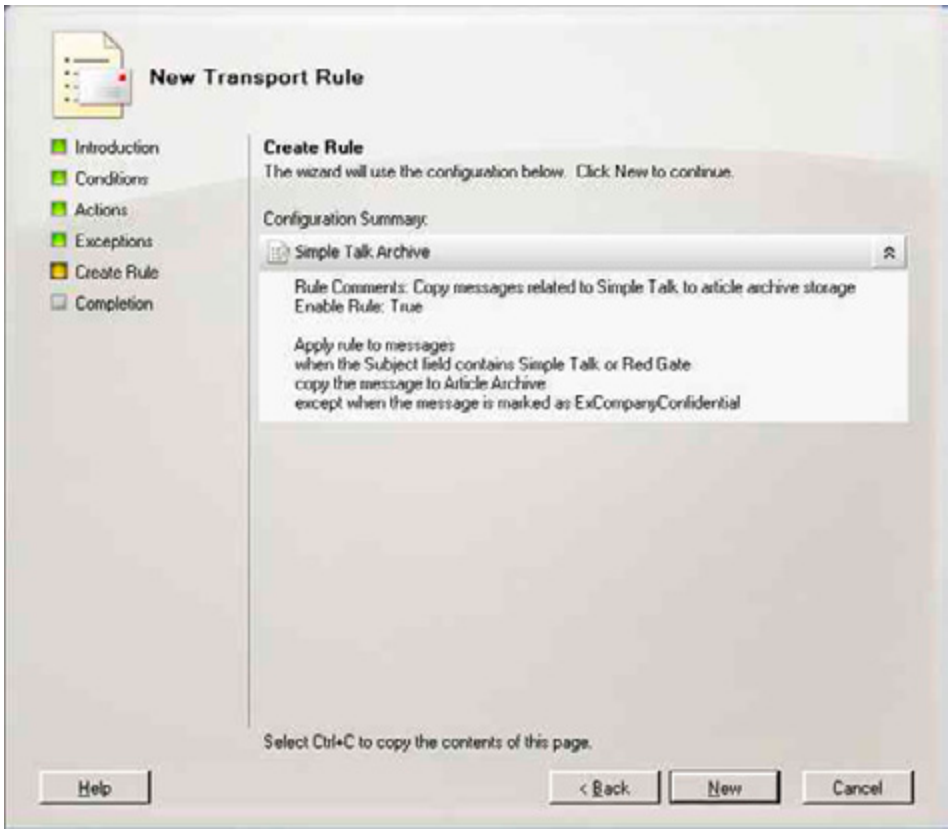

Figure 9.

Clicking New will complete the rule and present the EMS code that was used to create it (see Figure 10).

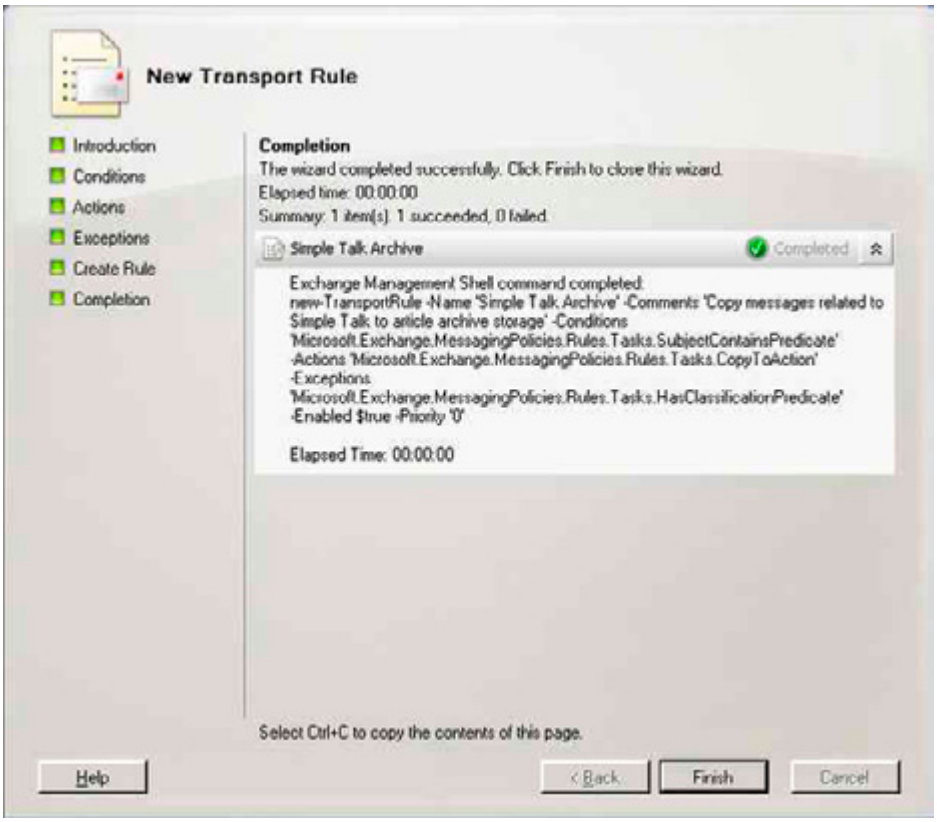

Figure 10.

A CTRL-C will copy this EMS command to the clipboard. In most places in the EMC, the cmdlet is complete; however, for transport rules, the variables are not displayed in the UI. They must be assigned manually if you are entering the transport rule using the New-TransportRule cmdlet. In our example, this command is as follows:

```
[PS]C:\>
$Condition = Get-TransportRulePredicate SubjectContains
$Condition.Words = ("Simple Talk,""Red Gate")
$Action = Get-TransportRuleAction CopyTo
$Action.Addresses = @((Get-Mailbox "Article Archive"))
$Exception = Get-TransportRulePredicate HasClassification
$Exception.Classification = (Get-MessageClassification ExCompanyConfidential).Identity

new-TransportRule -Name 'Simple Talk Archive' -Comments 'Copy messages related to Simple Talk to
article archive storage' -Conditions $Condition -Actions $Action -Exceptions $Exception -Enabled
$true -Priority '0'
```

The last parameter sets a transport rule priority. Transport rules are applied in order of priority starting with "0'. Rules are added in the order they are created. It may be necessary to move a rule up the list and is controlled using the -priority parameter in EMS for either *New-TransportRule* or *Set-TransportRule* cmdlets. This is also easily done in the EMC by using Change Priority option in the Actions pane when the desired rule is selected.

Figure 11.

Figure 11 shows the interface for entering a numerical priority from 0 (highest) to the number of rules less 1. Microsoft recommends a maximum of 1000 rules, mostly because that is where they stopped testing. This should be more than enough for most companies.

You modify existing transport rules in much the same way - either with the EMC or EMS. The EMS uses the *Set-TransportRule* cmdlet for updating rules. For a list of cmdlets for managing Transport rules in the EMS, see *Get-Help* as shown in Figure 12.

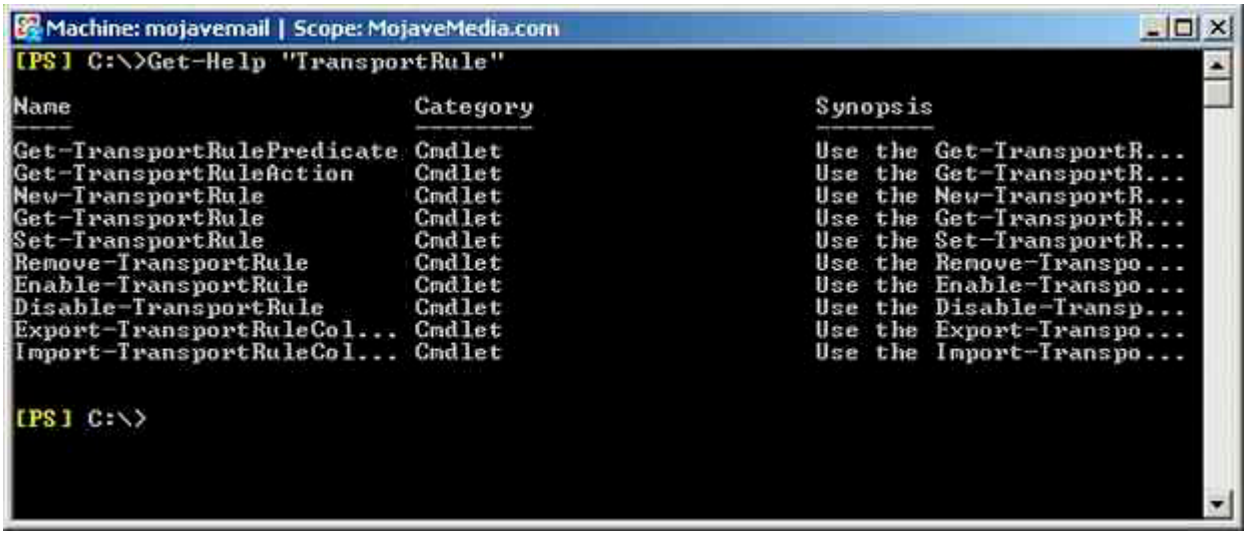

Figure 12.

Backing up Transport Rules

Internally, Transport Rules reside in AD, so they are backed up with AD. If you are going to make significant changes to transport rules, you might want to backup the set currently in place first. The EMS cmdlet *Export-TransportRuleCollection* bundles all of the transport rules into a file which can be imported later if needed. Importing overwrites the existing rules. On an Edge Server, exporting the transport rules is used as a backup mechanism or to import them into other "cloned" Edge Transport servers.

Summary

Exchange Server 2007 takes great steps forward in controlling messages and message transport from an administrator perspective. Some tasks that were challenging or required third party applications in previous versions of Exchange have been made more accessible through transport rules in Exchange 2007. With the Edge Server role concerned with security and the Hub Transport role focused on compliance and policy enforcement, Transport Rules provide a much improved set of tools for administrators to manage message flow in their Exchange organizations.

Exchange 2007 Mailbox Server Clustering

12 August 2008

by Jaap Wesselius

Note

The original, rather complex, Exchange clustering, 'Single Copy Cluster', protects you against a server failure but not a database failure. 'Local Continuous Replication' protects you from database failure, but not server failure. The more simple 'Clustered Continuous Replication' protects against both. If you use it with Transport Dumpster, you won't lose emails.

Exchange 2007 Mailbox Server Clustering

In the article High availability in Exchange 2007, several options in Exchange Server 2007 were discussed to create a high availability solution. In this article, all Exchange Server 2007 server roles are discussed. I discuss some of the technical challenges that you can expect when deploying clustering technologies, especially on the Exchange 2007 Mailbox Server role.

Exchange 2007 Single Copy Cluster

As explained in the previous article, an Exchange 2007 Single Copy Cluster (SCC) is basically a new name for the "traditional" Exchange clustering. An Exchange 2007 SCC environment is not very different from an Exchange 2003 cluster.

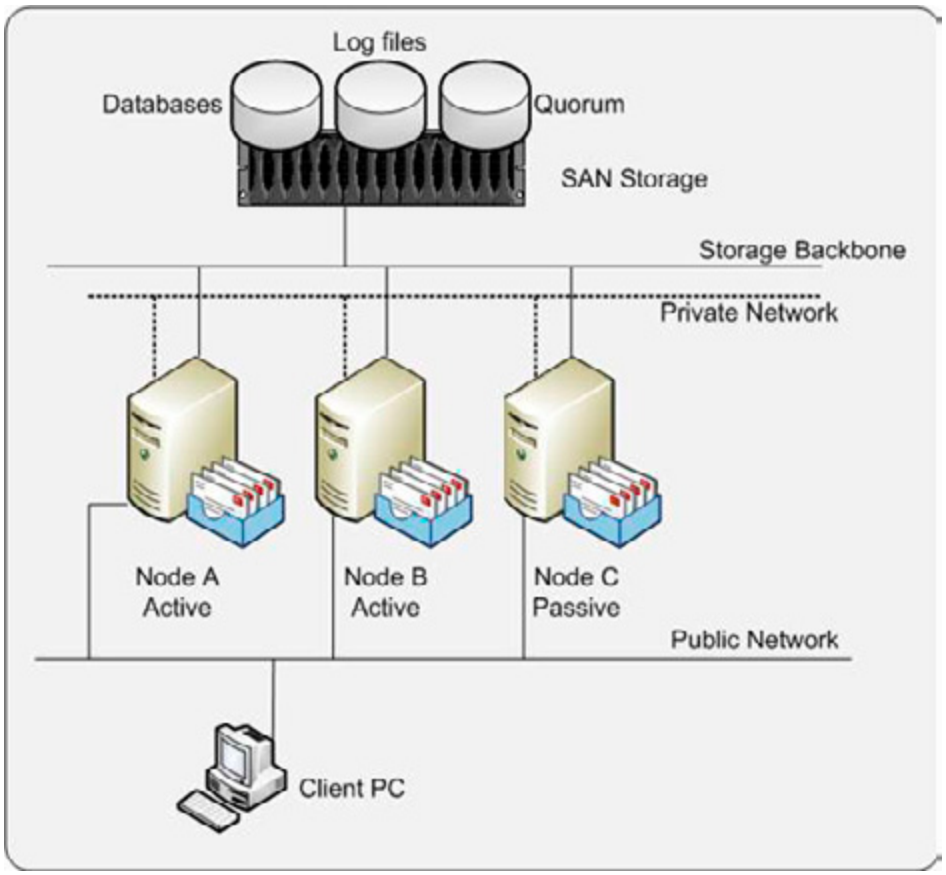

Figure 1. Single Copy Cluster.

The most important thing to remember in an SCC environment is regarding 'shared storage." Databases and log files are placed on a SAN or other central storage solution. This can be a Fiber Channel solution or an iSCSI solution. Data placed on the SAN is shared by all members of the cluster. Only one node is owner of the data and thus able to process the data at a given time. When this node fails, another node can take over the databases and log files and continue the processing of the data. For the end user, this is a fully transparent process; he or she will hardly notice that something has happened in the background.

Microsoft fully supports a clustered environment, but only if the complete configuration, including the Network Interfaces, host Bus Adapters and even the switches are listed in the Microsoft Logo'd Products List (http://microsoft.com/hcl). If the hardware configuration is listed, the hardware has been fully tested by both Microsoft and the vendor for use in a cluster.

Before installing Exchange, the cluster has to be built. Windows needs to be installed and the hardware needs to be properly configured. All nodes of the cluster need full access to the shared storage locations on the SAN. A special location on the SAN must be reserved for the quorum. This is basically the brains of the cluster, the location where the cluster configuration is actually stored. This is stored on the SAN, but is very limited in size. The smallest unit of allocation within your SAN is sufficient, and normally around 1 GB.

When the hardware is setup correctly you can setup your Windows Server 2003 cluster. This consists of the following steps:

- Create a cluster service account in Active Directory. This account will be used by the cluster service on the Windows Server 2003 cluster nodes. This can be a normal domain user account, but must be a member of the local administrators group on the Windows 2003 cluster node.
- Start the "Cluster Administrator" via the Administrative Tools menu and select "Create New Cluster."
- Using the New Server Cluster Wizard enter the new cluster name, the service account that you created in step 1, the public IP address of the cluster and the computer name of the first node of the new cluster.

- After verification of the data in the Proposed Cluster Configuration windows you can finish the new cluster configuration. The cluster will now be formed.
- Repeat steps 2 to 5, but instead of "Create New Cluster" select "Add Nodes" to cluster to start the "Add Nodes Wizard."

Please make sure that you have reserved sufficient IP addresses on the public network for your cluster configuration. You need an IP address for every node, an IP address for the Windows Cluster, and an IP address for every Clustered Mailbox Server (CMS) that you create. In figure 1 this means that six public IP addresses are needed.

For internal communication within the cluster a private network is used. This network should not be used for other communications, but only for the cluster. The cluster service uses this network for sending so called "keep alive" messages. Because of sending and receiving these messages a particular node in the cluster is aware of the presence and health of the other nodes.

If the first node in the cluster is configured correctly, you can start adding other nodes to the cluster. Just start the Cluster Administrator in the Administrative Tools menu and select 'Add node." The new node contacts the cluster, reads its information, and adds itself to the cluster. Naturally you need to repeat this step for all nodes you want to add to the cluster.

Installing Exchange Server 2007 on the cluster is similar to installing Exchange Server 2007 on a regular server. However you may only install one of two specific clustered mailbox roles. Instead of installing the normal Mailbox Role you have to install the 'Active Clustered Mailbox Role." This will install an instance of Exchange Server 2007 in the cluster.

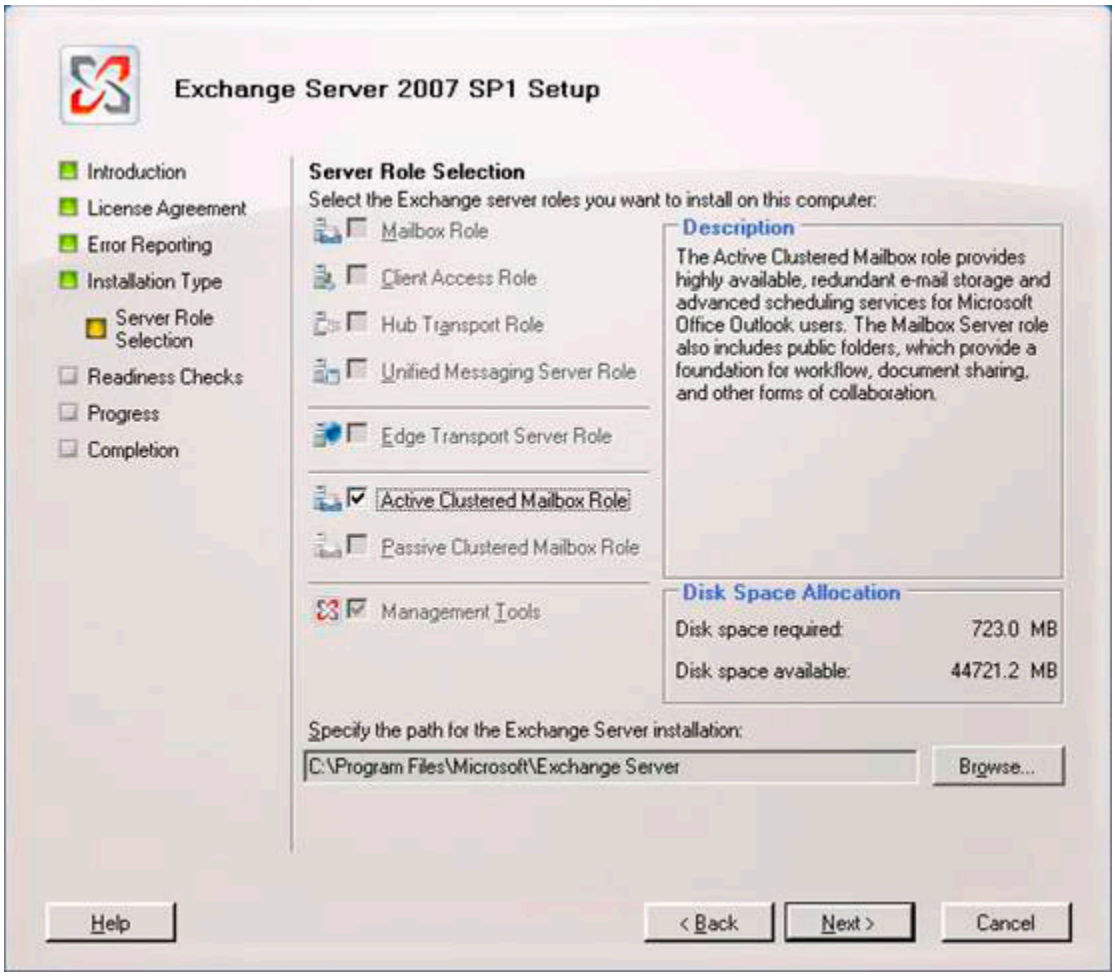

Figure 2. Installing the Active Clustered Mailbox Role.

Please be aware that only an Exchange 2007 mailbox server role can be installed on a Windows cluster. For the other roles, such as Client Access Server or Hub Transport Server, you have to install additional non-clustered servers as discussed in the prior article.

When finished installing the Active Node of the cluster, you can continue with the other nodes. In the example shown in Figure 1, you can continue installing another Active Node and finish with installing the Passive Node of the cluster.

Single Point of Failure

An SCC environment works great if a server is lost, but the databases and the log files are still a single point of failure. When the database is corrupt there is the possibility that you won't get it up and running again in minutes. When this happens there are two options:

- Restore from backup media
- Repair the database with Microsoft tools.

Whichever you choose, it will most probably be a lengthy operation. Suppose a normal tape backup device will restore at 50 GB/hour and you have a 150 GB database. In that case it will take at least three hours to restore your database from tape. Add to that the time needed to find the right tape, to start the restore itself and maybe to replay log files that are newer than the last backup. There have been situations where an Exchange Server in a clustered environment was out-of-order for more than 24 hours.

Using the Microsoft Exchange tools to repair a corrupt mailbox database can also be a lengthy operation. Tools like Microsoft's ESEUTIL can process between 5 and 10 GB/hour, depending on the underlying hardware. But 10 GB/hour processing related to the 150 GB database, ESEUTIL will take at least 15 hours to finish processing. And Microsoft PSS would like you to perform an offline defrag with ESEUTIL after you ran a repair with ESEUTIL. So this adds another 15 hours of processing your database. This totals 30 hours when your database and therefore your mailboxes are not available.

Implementing a SCC environment will help you protect against server failure, but since you only have a single copy of your data, it will not protect against database failure. You can lower the risk and possible impact by implementing multiple smaller databases in your environment instead of one large database. If a database failure occurs, only a part of your organization is impacted. But there's a part still impacted and this part will see an outage of the messaging service.

Replication Techniques

To prevent an outage due to a database failure, Microsoft built some techniques into Exchange Server 2007 that are based on replication techniques. With replication techniques a second copy of the database, or databases are maintained, therefore eliminating this single point of failure.

In Exchange Server 2007 RTM ("Release To Manufacturing," i.e. the original version of Exchange Server 2007) Microsoft introduced two types of replication:

1. Local Continuous Replication
2. Clustered Continuous Replication

In Exchange Server 2007 SP1 Microsoft introduced a third type of replication:

3. Standby Continuous Replication

Since Standby Continuous Replication is not dependent on any type of Windows Server 2003 clustering it is out-of-scope for this document.

Local Continuous Replication

The first solution from Microsoft for replication is Local Continuous Replication (LCR). As explained in the earlier article, LCR eliminates the single point of failure for the database, but it is still implemented on a single server. It does protect you against a database failure, but it does not protect you against a server failure.

To implement LCR on your Exchange 2007 mailbox server you need an additional disk (or LUN on a SAN) and preferably an additional controller to access this disk or LUN (this is in order to ensure that writing the LCR database copy does not negatively impact production operations). Format this disk with NTFS and give it a drive letter.

One prerequisite for implementing LCR is that a Storage Group can contain only one database. So if you need to implement multiple databases you automatically need to implement multiple Storage Groups.

Although not a hard prerequisite it is a best practice to implement volume mount points for designating where database files and transaction log files are stored. When using volume mount points it is much easier to replace the active copy of the database with the passive copy of the database as explained later in this article.

The next step is to enable LCR on the Storage Group on your Exchange 2007 mailbox server.

Open the Exchange Management Console and select the proper Mailbox Server under Server Configuration. Select the Storage Group you want to enable LCR for and select 'Enable Local Continuous Replication" in the tasks pane on the right hand side of the console.

When selected you have to enter the path for the system files, the log files, and the database file. Click finish to have the Exchange Server configure the LCR environment.

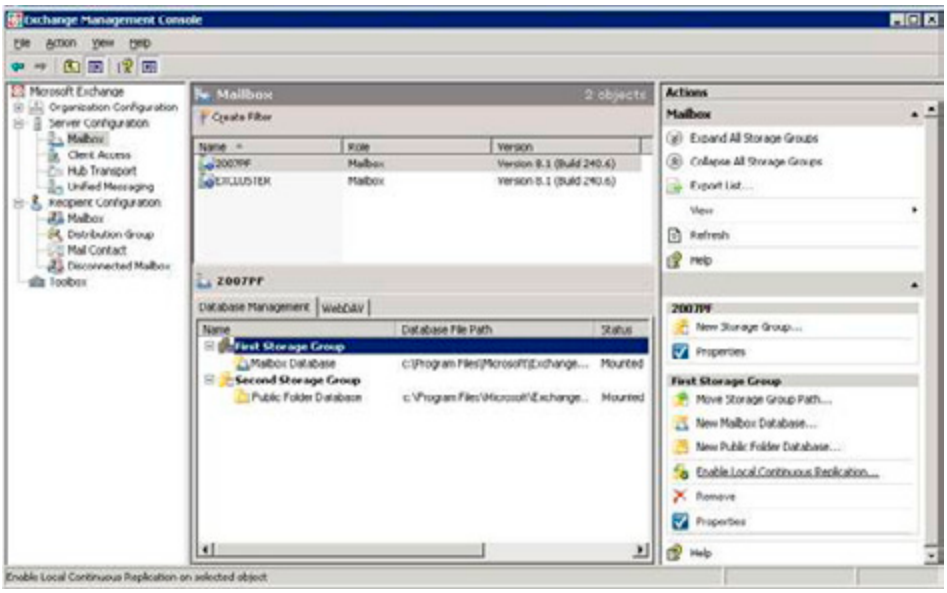

Figure 3. Enable LCR in the Exchange Management Console on the First Storage Group.

When finished configuring LCR, Exchange Server 2007 needs some time to replicate the database and the log files to the passive location and have the database duplicate up-and-running. This can be seen in Figure 4. Please note that the files are located on the c:\ drive instead of a different drive. This is configured on a demo system that has no additional disk. In a real world scenario LCR should use a physical separate disk. Depending on the size of the database the this can take a serious amount of time.

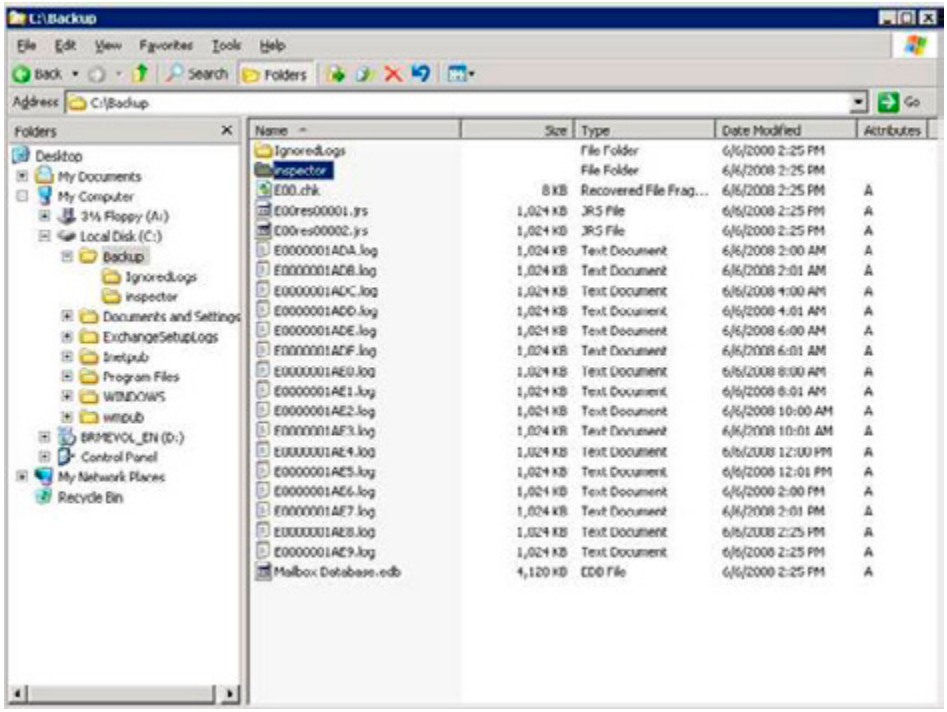

Figure 4. The additional LCR copy as seen in Explorer.

When the Exchange 2007 Mailbox Server role is processing normal message data, all information is flowing through the log files in the active copy. The active copy is the copy of the database that is used for normal operations. When Exchange has filled a log file and the log file is closed it is sent to the passive copy. This is the copy that is in use by LCR. The log file that was received is inspected and replayed into the passive copy of the database.

When the active copy of the database gets corrupt, or the database hardware fails, the end-users will experience an outage of the Exchange service. The messaging administrator now has to replace the active copy of the database, which is the unavailable copy, with the passive copy of the database.

In order to cause this replacement to occur, the following cmdlet should be run in the Exchange Management Shell:

```
Restore-StorageGroupCopy –Identity "SERVER\First Storage Group"
```

Replace the name of the server and the Storage Group with the actual names in your environment.

Using LCR protects you from a database failure by maintaining a copy of the database. Although manual interaction is needed in case of a failure, this works very well. Unfortunately it does not protect you from a server failure like SCC does.

For protecting against a server failure and a database failure, Microsoft combined the cluster technology with the replication technology. This resulted in the Clustered Continuous Replication (CCR).

Clustered Continuous Replication

CCR combines the server protection that the Windows Cluster Service offers in SCC and the database protection that LCR offers. This way your messaging environment is protected against a server failure and a database failure.

When a server failure occurs, the passive node of the cluster automatically takes over the Exchange service. This makes it fully transparent for end-users. Also, when a database failure occurs, the passive copy will take over the Exchange service since the passive copy has its own copy of the database.

A major difference with a traditional cluster is that the CCR environment does not have any shared storage. All data is placed on local storage. This can be just an additional disk or array of disk within your server or an external array attached to your server. This is called Direct Attached Storage or DAS. Naturally, the database can also be placed on a SAN (Fiber Channel or iSCSI) but there's no need to share this data between the nodes. Using DAS instead of SAN can make a major difference in the investment to make and may keep your messaging environment a lot simpler.

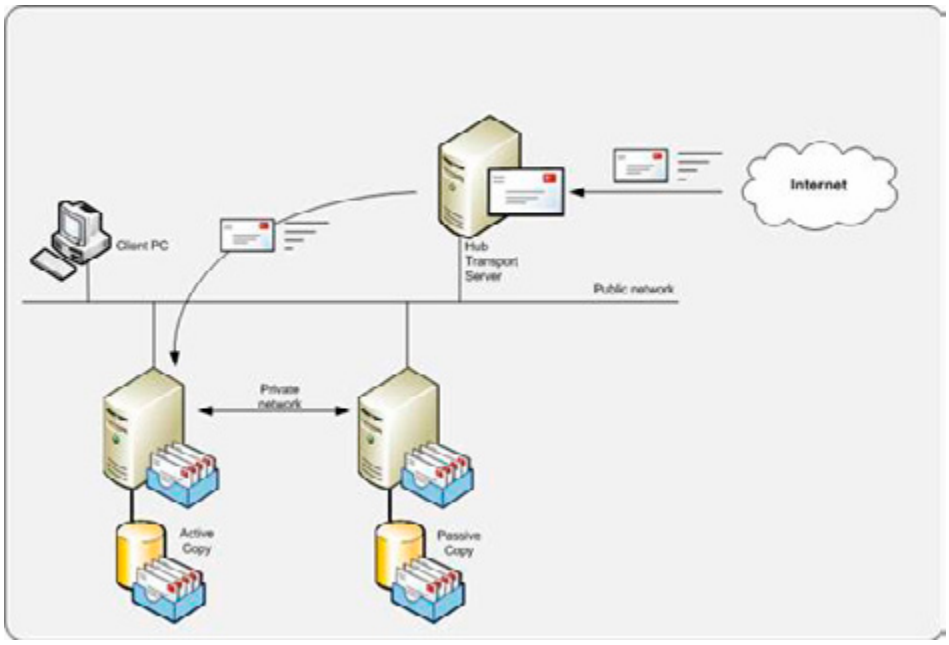

Figure 5. Clustered Continuous Replication.

When building a CCR environment, the Windows environment needs to be built first. A CCR environment does not use shared storage, so there's no quorum resource like in a SCC environment. Yes, there is a quorum, but it is shared between the nodes using normal networking. The cluster technology used for a CCR environment is a special version of clustering available in Server 2003 and Server 2008 called "Majority Node Set" or MNS. This is also a cluster version with an uneven number of nodes. In case of a failure, half of the number of nodes plus one should be available to keep the cluster up-and-running. For CCR, the MNS has changed to support only 2 nodes in a cluster.

A CCR environment can only have 2 nodes. More nodes is not supported and not built in the product.

When implementing a CCR environment, a special Windows share is created outside the cluster. This is normally created on a Hub Transport Server. This share is called the 'File Share Witness." In case of cluster problems, the passive node uses this share to determine whether the active node is still alive, or not. In the latter case, the passive node will start up.

Before implementing Exchange Server 2007 in a CCR environment, a cluster needs to be built. This can be using the cluster administrator or using a command line. The command line is useful if you have to build several clusters since you only have to change a couple of variables within a script. In this example we will build a Windows Server 2003 clustering using the command-line.

The first step is to create a cluster service account in your Active Directory. This account will be used by the cluster service on the individual nodes. This can be normal domain user account, but it must be a member of the local administrators group on the individual cluster accounts.

Then create the File Share Witness on the Hub Transport Server. The File Share Witness is used by the Windows 2003 Server cluster to determine which node will form the cluster when the individual nodes cannot communicate with each other anymore. The File Share Witness is created using this script:

```
Mkdir c:\MNS_Share1
Net share MNS_Share1=c:\MNS_Share1 /GRANT:Domain\ServiceAccount,Full
Cacls c:\MNS_Share1 /G BUILTIN\Administrators:F Domain\ServiceAccount:F
```

Note: replace ServiceAccount with the name of the account you created in the first step.

This script creates the directory on the local disk of the Exchange 2007 Hub Transport Server and shares it on the network. It also grants the built-in group Administrators and the cluster service account Full Access on the share and the directory itself.

When the File Share Witness is created we can start building the active node of the cluster.

```
::Step 3 - Initiate Cluster service
cluster WINCLUSTER /create /ipaddress:10.0.0.101 /User:domain\ServiceAccount /
password:<<password>> /verbose /unattended /minimum

::Step 4 – Make the cluster an MNS cluster
cluster /cluster:WINCLUSTER resource MNS /create /group:"Cluster Group" /type:"Majority Node Set"
/priv MNSFileShare="\\HUBSERVER\MNS_Share1" /online

::Step 5 – Delete the default Local Quorum
cluster /quorumResource:MNS
cluster /cluster:WINCLUSTER resource "Local Quorum" /Offline
cluster /cluster:WINCLUSTER resource "Local Quorum" /delete

::Step 6 – Add nodes to the cluster
cluster /cluster:WINCLUSTER /Addnodes:NODE1 /Minimum /Password:<<password>> /verbose
cluster /cluster:WINCLUSTER /Addnodes:NODE2 /Minimum /Password:<<password>> /verbose
```

This script creates the Windows cluster; you only have to replace the servername, the ServiceAccount and the IP Address according to your own environment. The next steps are to create a new MNS resource and define the File Share Witness on the Hub Transport Server as the MNS File Share within the cluster.

The original MNS quorum that's automatically created in step 3 is taken offline and deleted in step 5 since we created a new MNS quorum in step 4 based on the File Share Witness.

The last steps are adding the two nodes to the cluster. Only replace the names NODE1 and NODE2 with your own server names.

After creating the Windows cluster, Exchange Server 2007 can be installed and the Exchange 2007 CCR environment can be configured.

```
:: Step 7 – Install Exchange 2007 binaries and create Exchange cluster
<<source>>\setup.com /mode:Install /Roles:Mailbox /targetDir:"%programFiles%\Microsoft\Exchange
Server"

::Exchange Activecluster \\<source>\ setup.com /NewCMS /CmsIpAddress:10.0.0.110 /CmsName:EXCLUSTER
```

The first step installs the Exchange 2007 software in the appropriate directory. Please replace <<source>> with the location where the Exchange setup files reside, i.e. d:\ or \\anothershare\ if the software is on the network somewhere.

The second step creates the actual Exchange 2007 CCR environment.

The last step is to install the Exchange software on the passive node of the CCR environment using the following script:

```
:: Step 8 – Add Binaries on the Exchange Server 2007 passive node
<<source>>\setup.com /mode:Install /Roles:Mailbox /targetDir:"%programFiles%\Microsoft\Exchange
Server"
```

Your Exchange 2007 cluster is now ready to use and you can start configuring the Exchange 2007 environment using the Exchange Management Console or the Exchange Management Shell and the Windows Cluster Administrator.

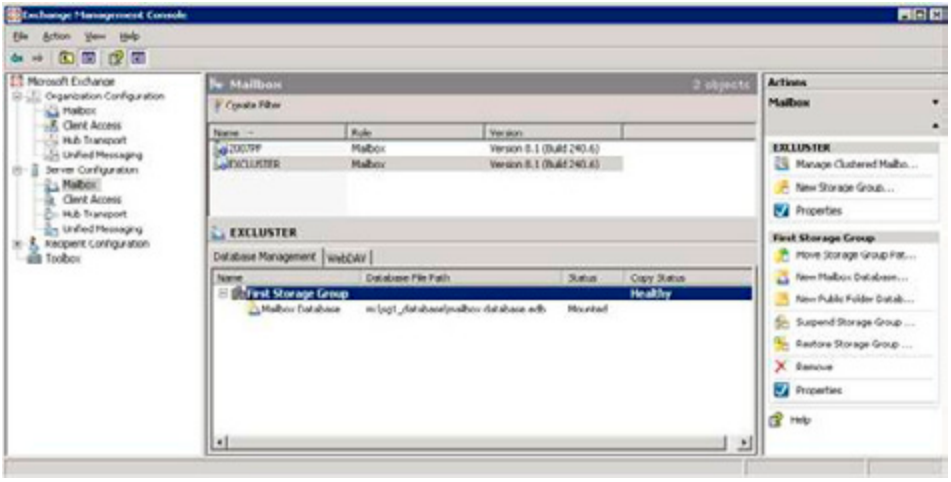

Figure 6. The newly created CCR environment as seen in the Exchange Management Console. Note the healthy state of the Copy Status.

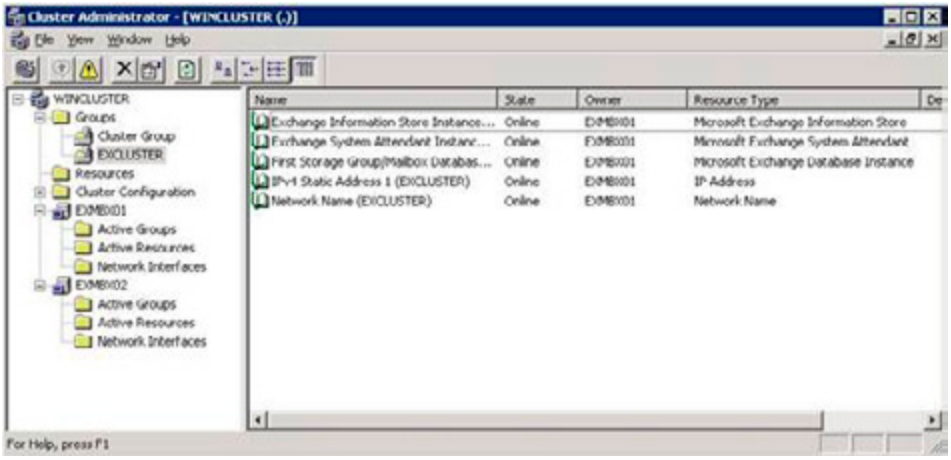

Figure 7. The newly created CCR environment as seen in the Cluster Administrator.

In contrast to LCR, where recovery is a fully manual process, recovery in a CCR environment is fully automatic. When the passive node fails, whether it is a server failure or a database failure, the passive node automatically takes over.

The replication process is identical to LCR. When the active node has finished processing an Exchange log file, the log file is replicated to the passive node. After receiving the log file, it is inspected and replayed into the passive copy of the database. There's a delay however, between the active and the passive copy. Depending on the utilization of the active Exchange 2007 mailbox server this delay can be more than 20 log files.

This means that when a fail-over occurs and the passive node takes over, there is a potential loss of e-mail. To prevent this, the Exchange 2007 Hub Transport server keeps a cache of all messages sent to the active node of the cluster. This cache is called the Transport Dumpster. When a fail-over occurs, the replication service contacts the Transport Dumpster to request a redelivery of the last messages. This way a near zero mail loss fail-over can be achieved.

The health of a Storage Group copy can be retrieved by the following Exchange Management Shell Command:

```
Get-StorageGroupCopyStatus –identity "NODE1\First Storage Group" | fl
The |fl option gives you detailed output from the commandlet:
[PS] C:\Documents and Settings\Administrator>Get-StorageGroupcopystatus |fl

Identity              : EXCLUSTER\First Storage Group
StorageGroupName         : First Storage Group
SummaryCopyStatus        : Healthy
NotSupported           : False
NotConfigured          : False
Disabled              : False
ServiceDown            : False
Failed               : False
Initializing           : False
Resynchronizing          : False
Seeding              : False
Suspend              : False
CCRTargetNode           : EXMBX02
FailedMessage           :
SuspendComment           :
CopyQueueLength          : 0
ReplayQueueLength         : 0
LatestAvailableLogTime       : 6/6/2008 7:55:38 PM
LastCopyNotificationedLogTime  : 6/6/2008 7:55:38 PM
LastCopiedLogTime         : 6/6/2008 7:55:38 PM
LastInspectedLogTime        : 6/6/2008 7:55:38 PM
LastReplayedLogTime        : 6/6/2008 7:55:38 PM
LastLogGenerated          : 60102
LastLogCopyNotified        : 60102
LastLogCopied           : 60102
LastLogInspected          : 60102
LastLogReplayed          : 60102
LatestFullBackupTime        : 6/5/2008 8:01:07 PM
LatestIncrementalBackupTime   : 6/6/2008 6:00:54 PM
LatestDifferentialBackupTime   :
LatestCopyBackupTime        :
SnapshotBackup           : True
SnapshotLatestFullBackup     : True
SnapshotLatestIncrementalBackup : True
SnapshotLatestDifferentialBackup :
SnapshotLatestCopyBackup     :
OutstandingDumpsterRequests   : {}
DumpsterServersNotAvailable   :
DumpsterStatistics         :
IsValid              : True
ObjectState            : Unchanged

[PS] C:\Documents and Settings\Administrator>
```

It can always happen that for some reason the replication breaks and needs to be repaired. This should be performed on the passive node since this is the location where the replica resides.

The replication can be fixed by entering the following commands in the Exchange Management Shell:

```
Suspend-StorageGroupCopy -Identity "excluster\second storage group"
```

Next step is to delete the passive copy of the database file and the corresponding log files and let the replication process update the passive copy by entering the following command in the Exchange Management Shell:

```
Update-StorageGroupCopy -Identity "excluster\second storage group"
```

A new copy of the database will be created on the passive copy and the log files will be replication again to the passive node. After convergence, both the Exchange Management Console (as can be seen in Figure 6) as well as the output from the Get-StorageGroupCopyStatus commandlet should read "healthy" again.

Conclusion

Exchange clustering techniques can be used to increase the availability of your Exchange Mailbox Server. The traditional clustering technique, now called Single Copy Cluster, does protect you against a server failure but does not protect you against a database failure. Also the requirement of shared storage can make the configuration complex, especially when your cluster hosts a large number of servers.

Replication techniques in Exchange Server 2007 can safeguard you from a database failure. With Local Continuous Replication you are protected against a database failure, but since LCR is built on top of a single server it does not protect against a server failure. Furthermore, recovery from a database failure is a fully manual process.

CCR is a very nice solution, protecting against both a server failure as well as a database failure. Also the combination with the Transport Dumpster will lead to an almost zero mail loss scenario.

Since the CCR solution is not built on shared storage, the CCR solution is far less complex than a SCC environment and it even works well with Direct Attached Storage.

The best option to implement is a CCR solution when you want to achieve a high availability Exchange Sever 2007 environment.

Top Tips for Exchange Admins

22 August 2008

by Michael Francis

Our top tips competition this month has been dominated by one man, an admin of rare enthusiasm.

Ben Lye of The MathWorks has sent us three tips this week, and these are as follows.

In the spirit of the Olympics....the Gold goes to....

Disabling PST files – straightforward I know, but this is something every admin should do. The effort this takes, vs. the effort it could save, means we have a clear victor. Ben writes....

"Microsoft Outlook supports a couple of options for restricting PST files. The first is a registry change called PSTDisableGrow prevents new data being added to existing PST files. The second is a registry change called DisablePST which prevents users creating or opening PST files altogether. The settings can be applied individually or together to phase out the use of PST files, and they can both be applied through group policy.

PSTDisableGrow doesn't have a great deal of documentation, but DisablePST is pretty well documented here: http://support.microsoft.com/kb/896515."

Silver goes to a hardworking runner-up, striving for automation and convenience.....

Smallest database script – Put in the effort, quite useful, but doesn't quite have the sheer vital, everyone-must-do-this-now urgency of the first place.

"It's based on one we use as part of our mailbox provisioning process, and it returns the smallest database on a specified server. By default it sums the sizes of all the mailboxes in each database, but it can also look at the size of the EDB file instead. (Summing the mailbox size will avoid counting whitespace in the EDB file.) It would be pretty easy to modify to look at specific storage groups, or use other search criteria (we look for databases which are named according to physical office locations, and pick the smallest database for that office).

The script takes two parameters – the server name and the optional flag to look at EDB file sizes instead of mailbox sizes. If the server name isn't specified you will be prompted for it.

To get the smallest database on a server name SERVER01 using mailbox sizes:

```
Get-SmallestDatabase.ps1 -server SERVER01
```

To get the smallest database using the EDB file size:

```
Get-SmallestDatabase.ps1 -server SERVER01 -edb
```

The script returns the database object as the output, but again that could easily be changed to suit a particular need."

```
# Script to return the smallest database on a specified server
# Written by Ben Lye - 20th August 2008

# By default the smallest database is determined by summing the size of all the mailboxes in each
database.
# Optionally the -edb flag can be specified to make the script look at the EDB file size instead.
# If no server name is specified on the command line one will be prompted for

# Usage:
# Get-SmallestDatabase.ps1 [-Server <Server name>] [-edb]

# Get the command line parameters
Param ([string]$server,[switch]$EDB)

# Load the Exchange 2007 snap-in if they are not already loaded
Add-PSSnapIn -Name Microsoft.Exchange.Management.PowerShell.Admin -ErrorAction SilentlyContinue

# Check that the Exchange 2007 snap-in loaded
$snapin = Get-PSSnapin -Name Microsoft.Exchange.Management.PowerShell.Admin -ErrorAction
SilentlyContinue
If (-not $snapin) {
    Write-Host "Error: Exchange 2007 snap-in not found" -ForegroundColor "Red"
    Write-Host
    break
}

# Prompt for a server name if one wasn't passed in
if (-not $server) {
    $server = Read-Host "Server name"
}

# Find any databases on the specified server
$databases = Get-MailboxDatabase -Server "$server" -ErrorAction SilentlyContinue

# Error if there are no databases found
If (-not $databases) {
    Write-Host "Error: No databases found for server $server" -ForegroundColor "Red"
    Write-Host
    $break
}

If ($databases) {
    # Prepare some variables for storing the name and size of the smallest database
    $smallestdbsize = $null
    $smallestdb = $null

    # Loop through each of the databases
    Foreach ($database in $databases) {

       If ($EDB) {
           # Get the size of the .edb file
           $dbsize = (get-childitem ("\\" + $database.Server + "\" + $database.EDBFilePath.
PathName -replace(":,""$")) | select-object name,length).length
       } Else {
           # Get the database size in bytes by summing the size of all mailboxes in the database
           $dbsize = (Get-MailboxStatistics -Database $database.Identity | Measure-Object
```

```
-Property TotalItemSize -Sum).Sum
        }

        # Compare the sizes to find the smallest DB
        if (($dbsize -lt $smallestdbsize) -or ($smallestdbsize -eq $null)) {
            $smallestdbsize = $dbsize
            $smallestdb = $database
        }
    }

    # Return the smallest database
    $smallestdb
}
```

And finally, the Bronze....

Change the mailbox information cache refresh – Neat, but I guess most of you that are in a hurry to apply mailbox information already caught onto this. Third place.

"By default the Exchange Information Store service caches information about mailboxes for two hours, meaning that any changes to mailbox quotas can take up to two hours to take effect. You can change the cache refresh interval to a recommended value of 20 minutes which means that you only have to wait a maximum of 20 minutes for mailbox quota changes to take effect.

This TechNet article describes the problem and the registry changes required to reduce the cache refresh interval: http://technet.microsoft.com/en-us/library/bb684892.aspx

Ben has won the grand prize of a $50 Amazon voucher for his efforts (that's about 1/6th the value of a real Olympic Gold, fact fans).

If you want a chance of winning a $50 Amazon voucher, or you simply feel the need to share knowledge with your fellow humans, then send your Exchange top tips for next month's competition to michael.francis@simple-talk.com

Exchange Database Technologies

22 August 2008

by Jaap Wesselius

> ***One of the most misunderstood technologies in Exchange Server, regardless of its version, is the database technology. Most people, even Exchange administrators know it is something to do with* ESE *and tools like* ESEUTIL*, but once it's running they leave it that way for the rest of their lives. It's too difficult and you'd better not touch it in case it breaks....***

In this article I'll explain some of the fundamentals of the Exchange Server database technology. The primary focus is on Exchange Server 2007, but when appropriate I'll refer to Exchange Server 2003. This article is the first in a series about Exchange database technologies, backup, restore and disaster recovery.

What's on the disk?

When you've installed Exchange Server 2007 you will find some database files on the C:\ drive, typically in "C:\Program Files\Microsoft\Exchange Server\Mailbox\First Storage Group\" as can be seen in figure 1.

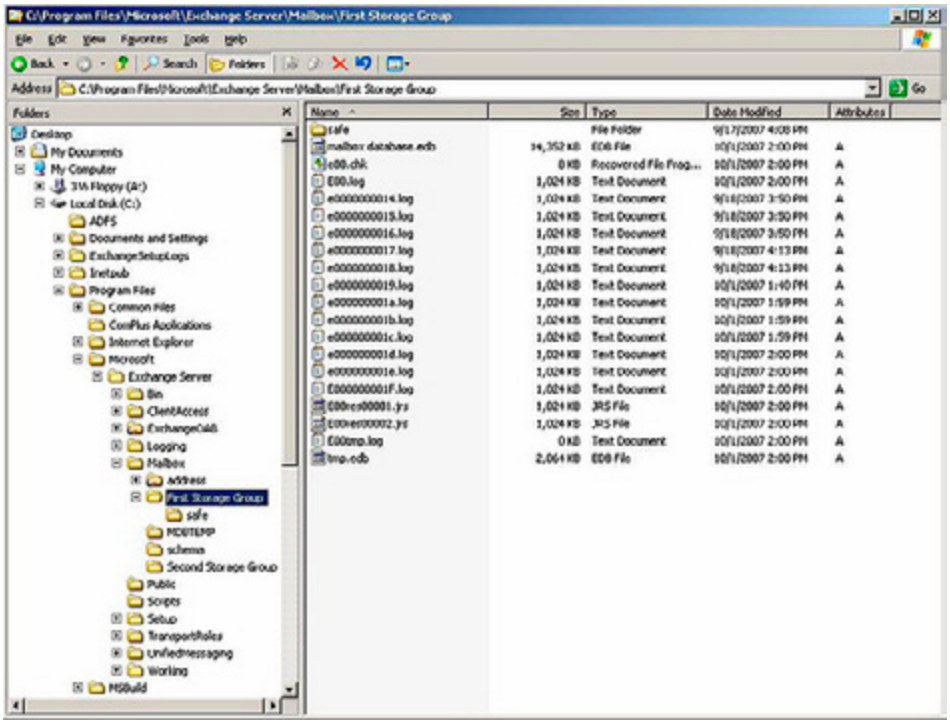

Figure 1. Database files in a standard Exchange setup.

We can see several files in this directory:

- **mailbox database.edb** – this is the actual mailbox database where all the mailboxes are hosted. One or more mailbox databases can be hosted in one directory
- **E00.log** and subsequent log files – these are log files that are used by Exchange server for transactional processing of all information
- **e00.chk** – a checkpoint file that keeps track of the relation between the information in the log files and in the database
- **E00tmp.log** – a temporary log file
- **E00res00001.jrs** and **E00res00002.jrs** – temporary reserved log files used by Exchange server in case a disk full situation occurs.

These files belong together and create a so-called "Storage Group." During the installation of Exchange Server one database is created in a Storage Group. In Exchange 2007 the Public Folder database is created in a second Storage Group. When replication technology (like LCR, CCR or SCR) in Exchange 2007 is used than only one database per Storage Group can exist. If no replication technology is used, up to five databases can be created in each Storage Group.

So, a Storage Group is a set of one or more database files that share a common set of log files.

ESE – Extensible Storage Engine

The underlying technology used in Exchange Server is the Extensible Storage Engine, or ESE. ESE is a low-level database technology, sometimes referred to as JET database. ESE has been used for Exchange since the first version of Exchange, version 4.0 in 1997. But Active Directory, WINS, and DHCP also use a form of ESE.

The ESE database follows the ACID principle. ACID stands for:

- **A**tomic – A transaction is all or nothing, there is no "unknown state" for a transaction
- **C**onsistent – the transaction preserves the consistency of the data being processed
- **I**solated – a transaction is the only transaction on this data, even when multiple transaction occur at the same time
- **D**urable – the committed transactions are preserved in the database.

Transactions can be seen in normal life as well. Suppose you go to the bank to transfer money from your savings account to your normal account. The money is withdrawn from your savings account and then added to your normal account and both actions are recorded and maybe printed on paper. This can be seen as one transaction. You don't want it to end in an unknown state, where the money is withdrawn from your savings account but not added to your normal account.

The same principle goes for Exchange Server. Suppose you move a message from your inbox to a folder named "Simple Talk." From a transaction point of view it starts by adding the message to the "Simple Talk" folder, then it updates the message count from this folder, it deletes the message from the Inbox and updates the message count for the Inbox. All these actions can be seen as one transaction.

The mailbox database

The mailbox database is the primary repository of the Exchange Server information. This is where all the Exchange Server data is stored. On disk it is normally referred to as "mailbox database.edb"; in older versions of Exchange Server it is called Priv1.edb, but it can have any name you want.

In Exchange Server 2000 there was also a file called Priv1.stm, referred to as the streaming file. It was meant to store Internet messages like SMTP. These messages were saved in the streaming file and a pointer was set in the .edb file. An .edb file and a .stm file belong together and cannot exist without each other. The streaming file was removed from Exchange 2007, which was possible because of the improvements to ESE, though Exchange 5.5 also had no .stm file.

In theory the mailbox database can be 16 TB in size, but it is normally limited to a size you can handle within the constraints of your Service Level Agreement or SLA. The recommend maximum database size of a normal Exchange Server 2007 (or earlier) is 50 GB, for an Exchange Server 2007 using Local Continuous Replication it is 100 GB, and for an Exchange Server 2007 using Continuous Cluster Replication it is 200 GB. These are sizes that can readily be used in a normal backup cycle and can be restored in an appropriate timeframe when needed.

The data within a database is organized in a Binary Tree, or B+ Tree. A Binary Tree can be seen as an upside down tree where the leaves are in the lower part. The actual mail data is stored in the leaves. The other pages only contain pointers. This is a very efficient way of storing data since it requires only two or three lookups to find a particular piece of data and all pointers can be kept in memory.

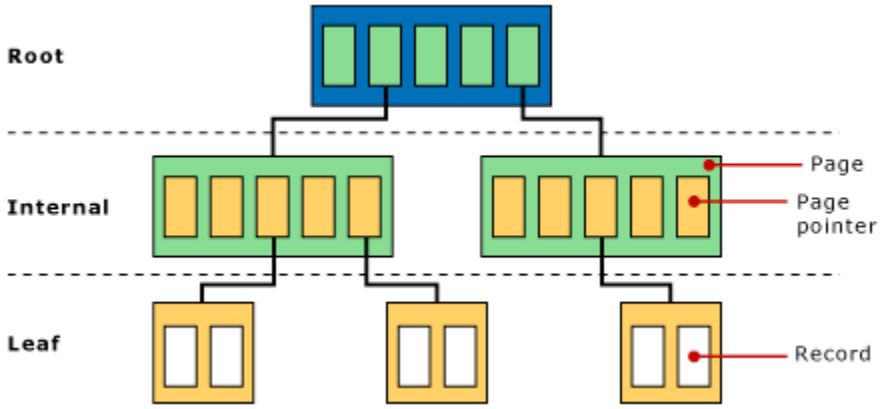

Figure 2. A Binary Tree used in Exchange Server. The actual data is stored in the leafs.

One or more Trees in a database make a table. There are several kinds of tables in an Exchange server:

- Folders table
- Message table
- Multiple Message Folder tables
- Attachment table.

The tables hold the information we can see in Outlook (Web Access). The tables consist of columns and records, the columns are identified as MAPI properties, the records contain the actual information.

Multiple Trees exist in one database and sometimes Trees need to be split when the Tree fills up. Exchange Server's internal processes will reorganize all information from one Tree into two Trees. This is called a "split." It is not possible to predict how many Trees will be in a particular database, but it will be hundreds or thousands of Trees. This can be seen in the header of the database which will be explained later in this article.

Database pages

A page is the smallest unit of data in an Exchange Server environment. For Exchange Server 2007 a page is 8KB in size, all earlier versions use a 4KB page size. Increasing the page size from 4KB to 8KB results in a dramatic increase in performance of the database technology. There are also several hardware storage solutions that perform much better with a page size of 8KB.

A page consists of a header, pointers to other pages, checksum information and the actual data from Exchange Server regarding messages, attachments or folders for example. A database file can consist of millions of pages. The total number of database pages can easily be calculated by dividing the total size of the database by the page size 8KB. If a database for example is 100 GB in size, it consists of 100 GB/8KB = approximately 13.1 million pages. Each page is sequentially numbered. Whenever a new page is created it gets a new incremented number. When pages are read from the database and altered, they get a new number before being written to the log file and flushed to the database file. Needless to say, this sequential number must be a very large number. It's a 64-bit number which means 18 quintillion changes can be made to a database.

How does it fit together?

There are four parts that are import with respect to Exchange Server data:

The **internal server memory** - this is the location where all the processing of data takes place. Exchange server creates new database pages when needed and processes them. When a database page is needed from the database file it is read from the disk into memory.

The **log files** - as soon as Exchange Server has finished processing the database pages for a transaction they are immediately written to the log file in use at that moment. This is the log file called E00.log, or E01.log, E02.log etc. depending on the Storage Group. Please remember that every Storage Group has its own set of log files starting with its own prefix like E00, E01, E02 etc. The database pages are kept in memory though because the pages might be needed in the near future. If so and it is still in memory it saves the Exchange Server a disk read. A disk read is a valuable disk action that is much slower that a page read from memory. When a log file is filled it is closed and renamed to a different name. This name consists of the prefix, followed by a sequential number which is incremented every time the log file is saved. This is way you can see files like E0000001.log, E0000002.log, E0000003.log etc. Be aware that the sequence use HEX numbering - the log file that follows E0000009.log is going to be E000000A.log, not E0000010.log. The sequence number in Exchange terms is called the lGeneration number. This process is called the log file roll-over.

The **database file**, the part where the database pages are stored eventually. As stated before, a transaction is written to the log file first but it is kept in memory. When Exchange needs more memory the database pages are flushed to the database file. They are flushed from memory to the database file, they are not read from the log files and then written to the database. This is a common misunderstanding! Flushing data to the database file also occurs when pages are kept too long in memory and the gap between the log files and the database file become too large.

The difference between the transactions in the log file and the transactions in the database is monitored by the **checkpoint file**, E00.chk (for the first storage group). The checkpoint file records the location of the last transaction written to the database file. As soon as a new transaction is written to the database file the checkpoint file location is advanced to the next location.

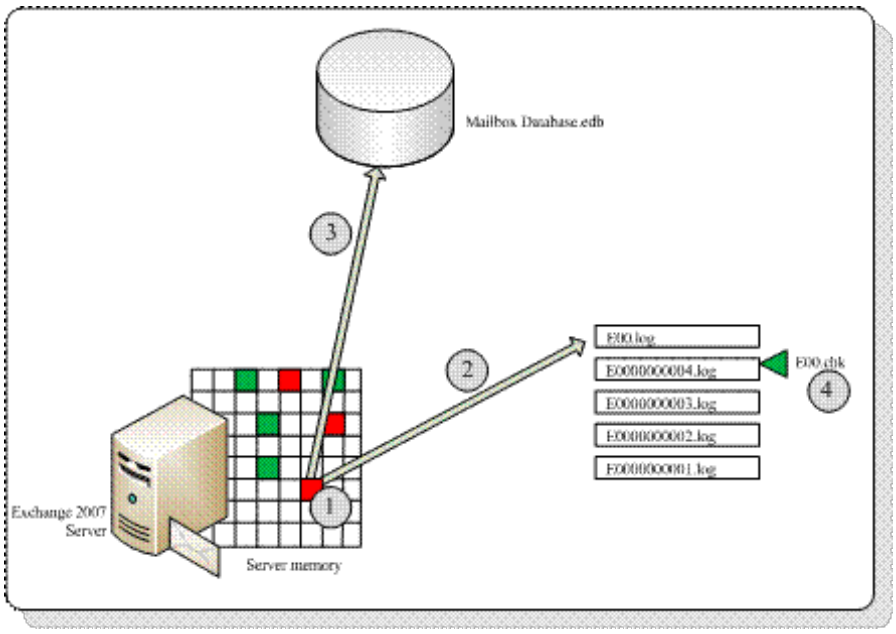

Figure 3. Schematic overview of the working of the Exchange Server database.

The above mentioned process is the same in all versions of Exchange Server.

Note

One serious constraint in Exchange 2003 and earlier is the 32-bit platform which is limited to only 4 GB of memory. All processes will run in this 4 GB memory, separated into 2 GB of system and 2 GB of user space. This is relatively tight for caching purposes, especially when a large number of mailboxes are consolidated on one Exchange 2003 server. It can be tweaked using the /3GB and the /USERVA switch in the boot.ini file of the Exchange Server, but it still isn't an optimal situation.

Memory constraints will always result in a higher disk IO, and when the storage is not designed properly this will result in a bad performance. Most of the Exchange 2003 performance issues are related to bad storage design. Exchange 2007 is a 64-bit platform and memory usage on a 64-bit server is only limited by your budget. Exchange 2007 will use as much memory as possible for caching, which will result in fewer IO operations on the disk. So, the more memory in your Exchange 2007 server, the less disk IO you will see on your storage.

All transactions in an Exchange server flow through the log files. This is the only way to achieve a consistent environment. This includes everything - new messages as they arrive, the creation of new folders, moving messages between folders, deleting messages, appointments, notes etc. All information in the mailbox goes through the log files. The creation of the mailbox (not the mailbox properties in Active Directory though!) is recorded in the log files, and even the creation of a mailbox database is recorded in the log files. This automatically means that if you lose your database and you still have all your log files available (maybe in combination with your last backup set) you can reconstruct everything up to the point you lost your database!

Note

It is very important to separate your log files and your database files. Separate them on multiple disks or multiple LUNs. If you separate them on multiple LUNs on your SAN, make sure that under the hood multiple, separated disks are used. This will have a positive impact on your performance and it will make sure you have a recovery path when either one of the disks is lost!

One important point to remember is that the log files are always in advance of the database, so there's always data not yet written into the database. This data can cover a number of log files and this number of log files or amount of data is called the "checkpoint depth." This automatically means the database in a running state is always in a non-consistent state. To get it in a consistent state all log files in this range are needed. This is because data is already written to the log files, but not yet to the database file.

In Exchange Server this is called a "dirty shutdown" state. When a database is dismounted it is brought into a consistent state. All data not yet written to disk is flushed to the database and all files are closed. All information is now written into the database. This is called a "clean shutdown" state of the database.

The log files needed to get the database into a consistent state or " clean shutdown" is recorded in the header of the database. The header of a database is written into the first page of the database file and contains information regarding the database. The header information can be retrieved using the ESEUTIL tool. Just enter the following command in the directory where the database file resides:

```
ESEUTIL /MH "Mailbox Database.edb"
```

Which will result in an output like this:

```
F:\SG2>eseutil /mh mbx2.edb

Extensible Storage Engine Utilities for Microsoft(R) Exchange Server
Version 08.01
Copyright (C) Microsoft Corporation. All Rights Reserved.

Initiating FILE DUMP mode...
     Database: mbx2.edb

    File Type: Database
  Format ulMagic: 0x89abcdef
```

```
  Engine ulMagic: 0x89abcdef
 Format ulVersion: 0x620,12
 Engine ulVersion: 0x620,12
Created ulVersion: 0x620,12
   DB Signature: Create time:07/03/2008 21:44:30 Rand:136947877 Computer:
     cbDbPage: 8192
      dbtime: 8851359 (0x870f9f)
      State: Dirty Shutdown
   Log Required: 14647-14658 (0x3937-0x3942)
  Log Committed: 0-14659 (0x0-0x3943)
  Streaming File: No
     Shadowed: Yes
    Last Objid: 6345
   Scrub Dbtime: 0 (0x0)
    Scrub Date: 00/00/1900 00:00:00
   Repair Count: 0
   Repair Date: 00/00/1900 00:00:00
 Old Repair Count: 0
 Last Consistent: (0x2F16,2D,181) 07/29/2008 19:34:29
   Last Attach: (0x2F17,9,86) 07/29/2008 19:34:29
   Last Detach: (0x0,0,0) 00/00/1900 00:00:00
       Dbid: 1
  Log Signature: Create time:07/03/2008 21:44:29 Rand:136951962 Computer:
    OS Version: (5.2.3790 SP 2)

Previous Full Backup:
    Log Gen: 8697-8722 (0x21f9-0x2212)
      Mark: (0x220D,1EC,162)
      Mark: 07/19/2008 22:57:43

Previous Incremental Backup:
    Log Gen: 0-0 (0x0-0x0)
      Mark: (0x0,0,0)
      Mark: 00/00/1900 00:00:00

Previous Copy Backup:
    Log Gen: 0-0 (0x0-0x0)
      Mark: (0x0,0,0)
      Mark: 00/00/1900 00:00:00

Previous Differential Backup:
    Log Gen: 0-0 (0x0-0x0)
      Mark: (0x0,0,0)
      Mark: 00/00/1900 00:00:00

Current Full Backup:
    Log Gen: 0-0 (0x0-0x0)
      Mark: (0x0,0,0)
      Mark: 00/00/1900 00:00:00

Current Shadow copy backup:
    Log Gen: 0-0 (0x0-0x0)
      Mark: (0x0,0,0)
      Mark: 00/00/1900 00:00:00

   cpgUpgrade55Format: 0
```

```
  cpgUpgradeFreePages: 0
cpgUpgradeSpaceMapPages: 0

    ECC Fix Success Count: none
  Old ECC Fix Success Count: none
     ECC Fix Error Count: none
   Old ECC Fix Error Count: none
  Bad Checksum Error Count: none
Old bad Checksum Error Count: none

Operation completed successfully in 0.578 seconds.


F:\SG2>
```

Plenty of interesting information regarding the database can be found in this output:

- **DB Signature** – a unique value of date, time and an integer that identifies this particular database. This value is also recorded in the log files and the checkpoint files and this ties them together.
- **cbDbPage** – the size of the pages used in this database. In Exchange 2007 this is 8KB, in earlier versions of Exchange Server this is 4KB.
- **Dbtime** – (part of) the number of changes made to this database.
- **State** – this files shows the state of the database, i.e. is it in a consistent state or not. The database in this example is in a "dirty shutdown" state (I crashed it to get the information) and it needs a certain amount of log files to get in a "clean shutdown" state.
- **Log Required** – If it is not in a consistent state, these log files are needed to bring it into a consistent state. To make this database a consistent state again, the log files E000003937.log through E000003942.log are needed. Exchange Server will perform the recovery process when mounting a database, so under normal circumstances no Exchange Administrator action is needed.
- **Log Committed** – This entry is for a new feature in Exchange 2007 called "Lost Log Resiliency" or LLR. Under normal operation the E00.log is just an open file. When the Exchange Server crashes there is the possibility that the E00.log will be corrupted. When this happens it is no longer possible to recover the database to a consistent state because E00.log is in the "Logs Required" range. The LLR feature takes the E00.log out of the "Logs Required" range, making it possible to do a recovery even if the E00.log is lost. One important thing to note is that all information already contained in the E00.log will be lost!
- **Last ObjID** – the number of B+ Trees in this particular database. In this example there are 6345 B+ trees in the database.
- **Log Signature** – a unique value of date, time and an integer that uniquely identifies a series of log files. As with the database signature this ties the database file, the log files and the checkpoint file together.
- **Backup information** – Entries used by Exchange Server to keep track of the last full or incremental (or VSS) backup that was made on this particular database.

The same kind of information is logged in the header of the log files (ESEUTIL /ML E00.LOG) and in the header of the checkpoint file (ESEUTIL /MK E00.CHK). As these file are grouped together you can match these files using the header information.

Conclusion

A Storage Group in an Exchange Server is one or more database files, a set of log files and a checkpoint file. These files belong together and have a tight relationship. But if you understand the basic principles of transactional logging it isn't too difficult. All information in a database is first recorded in the log files for recovery purposes, so deleting "unneeded" log files to create additional space is a bad idea since it will break the recovery process. Don't just delete these log files, but have them purged by a backup application. This way you have a recovery process using this backup and the log files that are generated after the last backup.

Bad things can happen to databases and log files in recovery scenarios. These bad things are not normally caused by Exchange Server itself, but by administrators that do not know what they are doing and the implications of their actions.

Install a test server, create some storage groups and mailboxes and start playing around with the databases. See for yourself what happens during a crash and discover the databases and the header information. It will be the first step in a solid disaster recovery plan!

In my next article I will explain what happens during recovery, replay of log files, and offline backups

Message Classifications in Exchange 2007

18 September 2008

by Neil Hobson

In Exchange 2007, you can now classify your messages in any way you wish, so that, for example, you can flag messages as being sensitive information that should not be sent outside the company. You can also create transport rules for all messages of a particular category. It is an easy way of implementing email policies within a company.

Introduction

Exchange 2007 has a new feature known as message classification that allows users to apply a classification to a message in order that the actual usage of that message is understood by both the sending and receiving parties. In previous versions of Exchange the concept of message classification has been restricted to marking messages with high, normal or low importance. Now it's possible to not only choose between the default message classifications that ship with Exchange 2007 but also to create custom classifications, as I'll show you later.

It's important to understand that message classifications are also an Outlook 2007 feature and therefore this is the version of Outlook you need to deploy to take advantage of this feature. However, there are several configuration changes required in order to make message classifications available to Outlook 2007 clients. If you have Outlook Web Access 2007 clients, these can use message classifications without any further modifications.

As I've just mentioned, there are message classifications that are provided by default. In all, there are six default message classifications:

1. A/C Privileged
2. Attachment Removed
3. Company Confidential
4. Company Internal
5. Originator Requested Alternate Recipient Mail
6. Partner Mail.

In the next section within this article I'll be showing you how to export the message classifications to an XML file that the Outlook 2007 clients within your environment can locate. This is to allow Outlook 2007 to display the classifications within email messages. Therefore, if you plan on creating new custom message classifications, you should do so before you export the classifications to an XML file. I'll be covering the creation of custom message classifications later in this article.

Creating the XML File

The first part of the process is the creation of the classification XML file that Outlook 2007 will reference. Fortunately Microsoft has made this part of the process easy by providing a PowerShell script that can do this for you. The script is installed along with Exchange 2007 and can be found in the *\Program Files\Microsoft\Exchange Server\Scripts* folder on the drive where you installed Exchange 2007. The script name is *Export-OutlookClassification.ps1* as you can see from Figure 1.

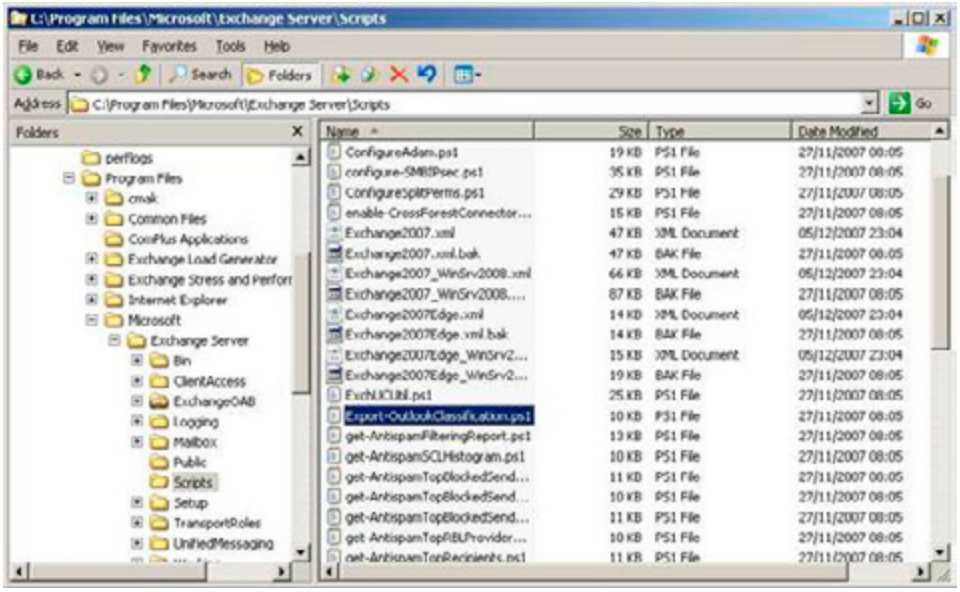

Figure 1: Exchange Scripts Folder.

To use this script, bring up the Exchange Management Shell and run the following cmdlet:

```
./Export-OutlookClassification > c:\classifications.xml
```

As you can see from the cmdlet example, the output of the PowerShell script is redirected to a file called *c:\classifications.xml*. Of course, you are free to use a different file name if you choose. If running the cmdlet has been successful, you should be returned straight back to the command prompt; in other words, there is no *"success"* message per se. To prove that the cmdlet has been successful, open the classification.xml file with Internet Explorer and check for valid contents. An example of what this file looks like when the six default message classifications have been exported is shown in Figure 2.

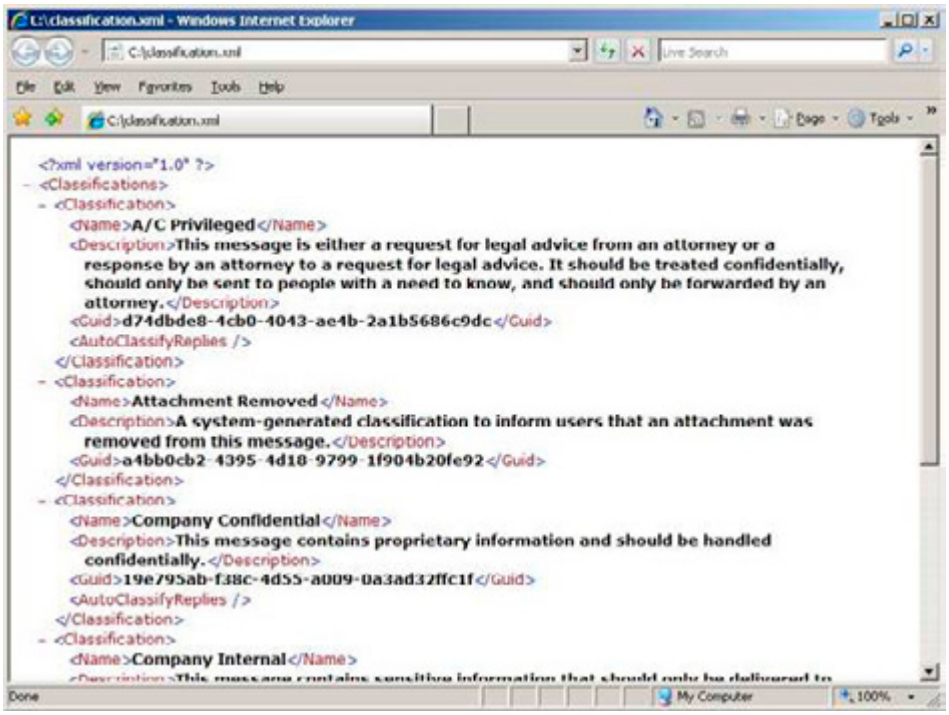

Figure 2: Contents of classification.xml.

The classifications that you have exported to the XML file are those classifications that can be chosen by the users who are sending the message; they have nothing to do with the type of message classification that a user can receive. I will expand on this later in the article. Now that you have exported the classifications.xml file, there are two additional parts of the overall message classification configuration to complete. First, you need to store the classifications.xml file in a location that each Outlook 2007 client can access and second, you need to make a registry change to each Outlook 2007 client to enable message classifications. I'll cover these two configuration elements in the next two sections of this article.

Locating the XML File

With regards to the location of the classification.xml file, you might think at first that the best location is on a network share, since you will only need to copy the file once to a specific location. However, it's actually better if you copy the file locally to each Outlook 2007 client that requires the use of message classifications. You have to consider the case of Outlook 2007 clients that run in cached mode. Outlook 2007 clients that are running in cached mode are sometimes disconnected from the corporate network, such as those users connecting via Outlook Anywhere when working from home. I'm not suggesting that the model of copying the classification XML file to every Outlook 2007 client is the best model that Microsoft could have come up with, but at the same time this is what we, as IT professionals, currently have to work with. Therefore, you'll need to produce a good working method, such as login scripts, to distribute the XML file to all Outlook 2007 clients along with the registry change that is detailed in the next section.

Required Registry Modification

Once you have copied the XML file to each client machine that requires the message classification functionality, you also need to create several registry values on these same client machines. The required registry information is as follows:

Key:
HKEY_CURRENT_USER\Software\Microsoft\Office\12.0\Common\Policy
String Value:
AdminClassificationPath
Value:
c:\classifications.xml
This is the location of the XML file and therefore must match the file name and location of your classification XML file.
DWORD Value:
EnableClassifications
Value:
1
This setting simply controls whether message classifications are enabled or not. Set this to 1 to enable message classifications or 0 to disable them.
DWORD Value:
TrustClassifications
Value:
1

The TrustClassifications setting should be set to 1 when the user's mailbox is on an Exchange 2007 server. This setting can also be used to control the prepending of text to the message classification when sending messages to mailboxes on legacy versions of Exchange, since these versions of Exchange do not support message classifications. I will not be covering this area any further within this article.

The *Policy* key is not present by default, and so must be created. Once the new information has been entered, the registry should look like the one shown in Figure 3.

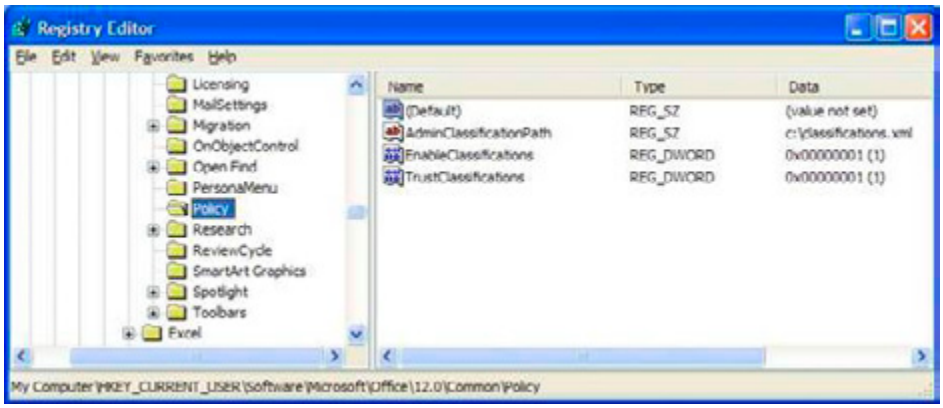

Figure 3: Registry Modifications.

Creating a Classified Message

Having set up classifications, creating a classified message couldn't be easier. Once you've copied the classification XML file to the Outlook 2007 client and created the required registry settings, launch Outlook 2007 and compose a new message. If you had Outlook 2007 open when making the registry changes, restart Outlook 2007 to start using message classifications.

In the new message window, you'll find the *Permission* button on the ribbon as you can see from Figure 4.

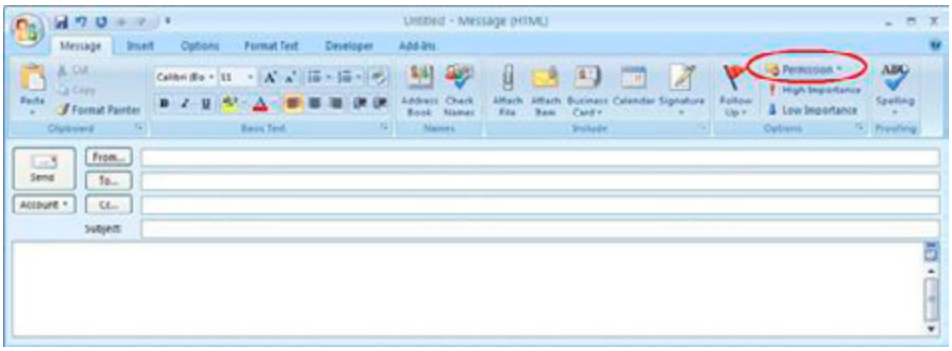

Figure 4: Outlook 2007 Permission Button.

Click the small down-arrow to the right of the Permission button and you will be presented with the six message classification options as defined in the XML file created earlier. You can see this in Figure 5.

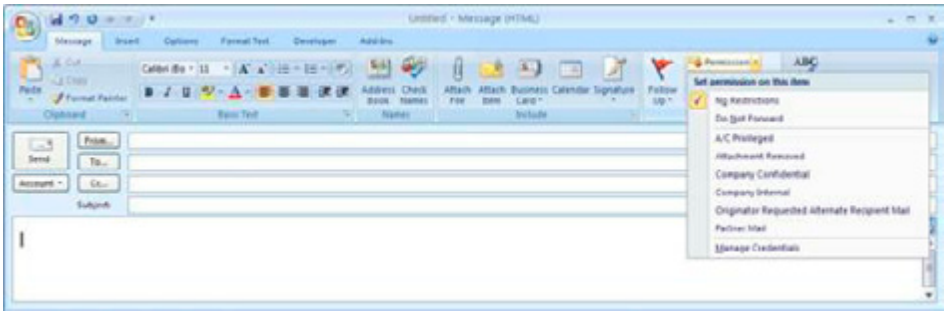

Figure 5: Default Classifications in Outlook 2007.

Let's say that I choose to classify this new message as *Company Confidential*. Once I've classified my message, it appears as shown in Figure 6 below.

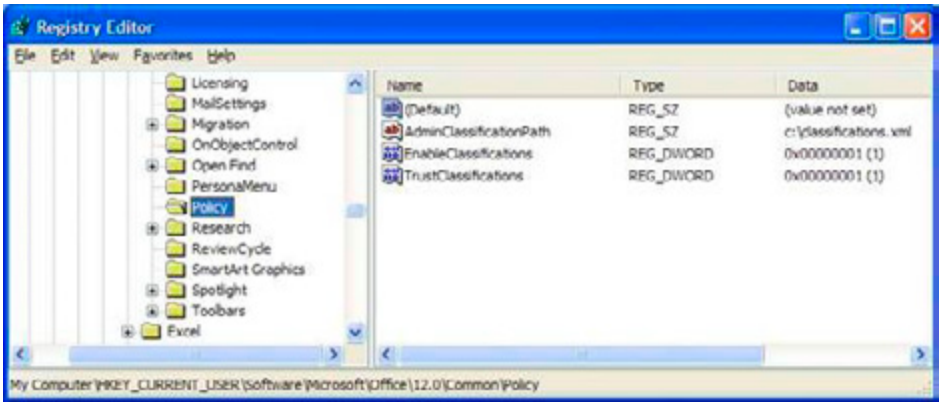

Figure 6: New Message Classified as Company Confidential.

What if the recipient, a user called Ann in this case, isn't enabled for message classifications and therefore doesn't have the required registry modifications in place? In this case, Ann just sees an ordinary message as shown in Figure 7.

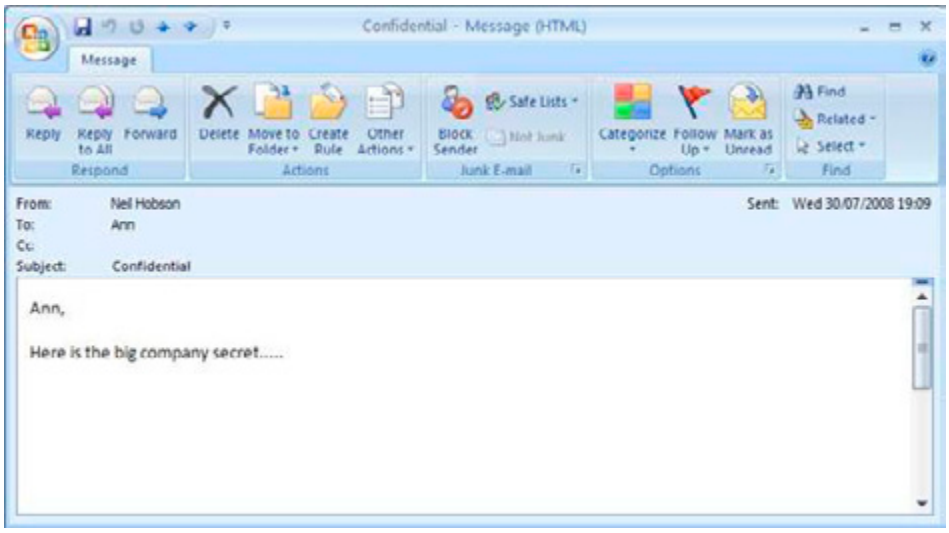

Figure 7: Received Message Without Classification.

The message classification metadata is still associated with the message even if Ann's client is not able to show it. We can determine that this is true by adding the required registry changes and restarting Ann's Outlook 2007 client. Once this has been done, we can see the message classification is now shown as you can see in Figure 8.

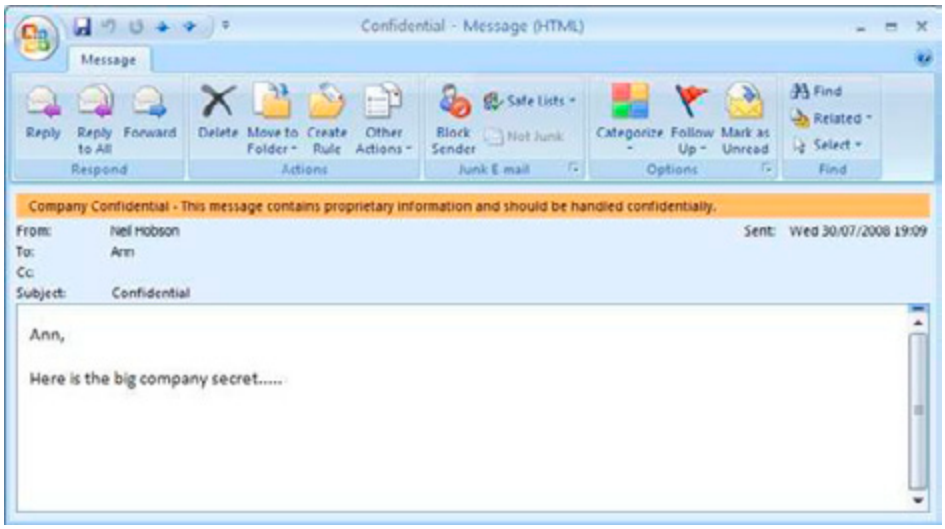

Figure 8: Received Message With Classification.

Creating Custom Classifications

The six default classifications may suffice for your needs, but there is always the chance that you will need something a little different. With that in mind, let's look at creating custom message classifications.

To create a new message classification you can use the *New-MessageClassification* cmdlet. In order to run this cmdlet, the account you are using must be delegated the Exchange Organization Administrator role, since you are making changes that affect the entire Exchange organization. Before we run the *New-MessageClassification* cmdlet, let's run the *Get-MessageClassification* cmdlet to confirm the presence of the default six message classifications. This is shown in Figure 9.

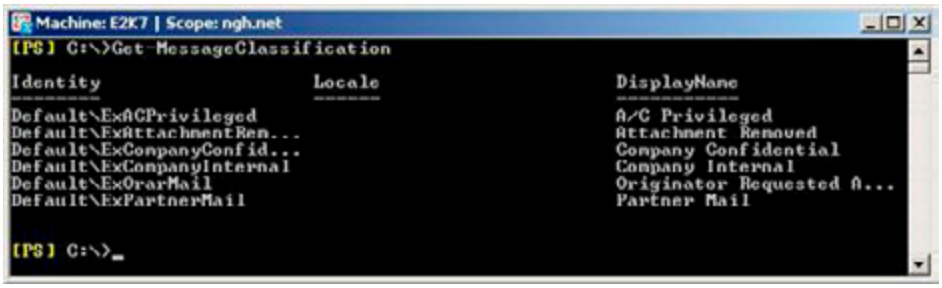

Figure 9: Default Message Classifications.

If you run the *New-MessageClassification* cmdlet without any additional parameters, you'll be prompted for three parameters to complete the creation process. They are the *Name*, *DisplayName* and *SenderDescription* parameters.

- **Name**. This is the administrative name of the classification. For example, if you want to retrieve details about the message classification with a Name attribute of Custom, you can use the Get-MessageClassification –Identity Custom cmdlet.
- **DisplayName**. The DisplayName attribute is the name of the classification as seen in Outlook 2007, as you have seen earlier in this article in Figure 5.
- **SenderDescription**. This is the description that the sender of the message sees in Outlook 2007. This is the orange bar that you can see in Figure 6.

Figure 10 below shows the process of creating a new message classification using just the three basic parameters.

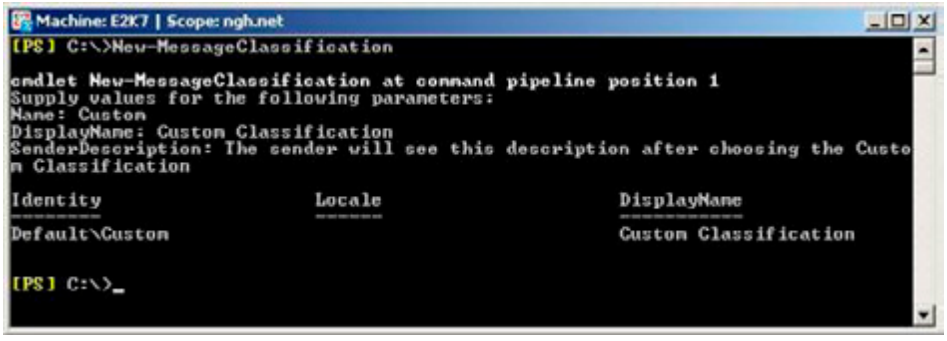

Figure 10: Creating a New Message Classification.

Here's something to note. Immediately after you have created this new classification, run the following cmdlet:

Get-MessageClassification custom | fl

This obtains full details about the newly created *Custom* message classification as you can see from Figure 11. What you may notice is that

the *RecipientDescription* attribute is populated with the same text that we supplied for the *SenderDescription* attribute, even though we never had to specify the *RecipientDescription* information during the creation of this new message classification. This is expected behavior if you do not specify the *RecipientDescription* text during the creation of the message classification.

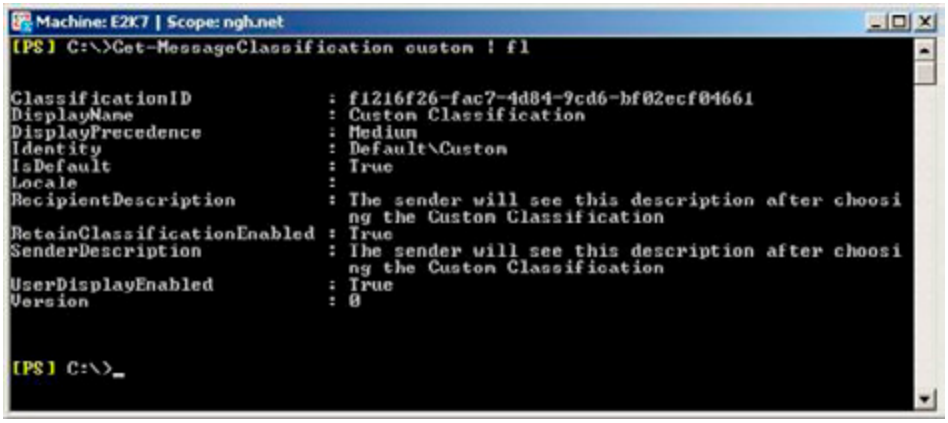

Figure 11: Custom Message Classification Parameters and Values.

As you can probably guess, the *RecipientDescription* attribute contains the text that the recipient of the message will see when opening the classified message. Once you've created the message classification, you can easily alter the parameters as with any other Exchange Management Shell cmdlet. For example, the following cmdlet alters the *RecipientDescription* attribute on the *Custom* message classification that we've recently created.

Set-MessageClassification Custom –RecipientDescription "The recipient will see this description after opening a message sent with the Custom Classification"

Once you have configured your message classifications, you need to re-export the entire list of message classifications into a new XML file and re-distribute to the Outlook 2007 clients. Therefore, you should ideally plan your custom message classifications before you export the list of classifications into an XML file for the first time. As you can see from Figure 11, once Ann opens a new message that has been classified with the *Custom* message classification, the new recipient description text is now displayed.

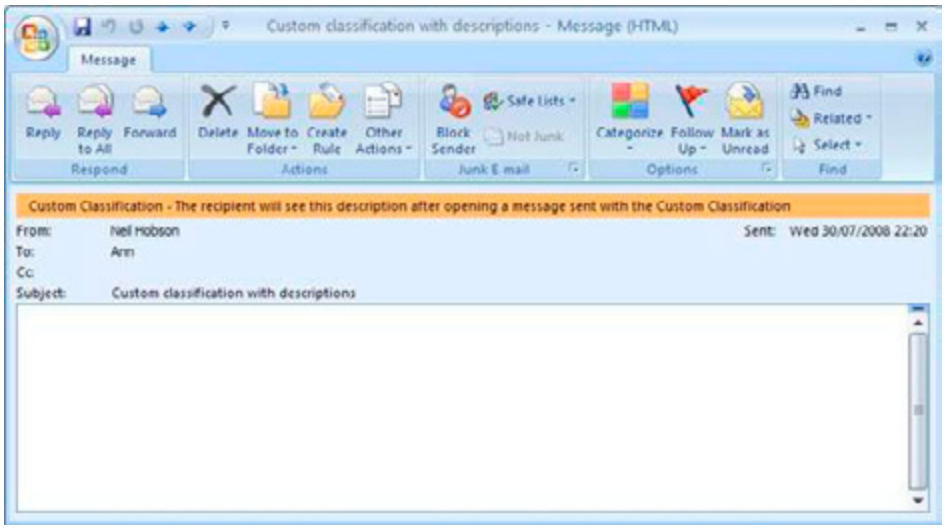

Figure 12: Recipient Description Information Displayed.

Manipulation With Transport Rules

With the INTRODUCTION OF TRANSPORT RULES within Exchange 2007, you can now begin to perform really useful administrative tasks that have previously been unavailable in legacy versions of Exchange, such as adding a disclaimer to all outbound email messages or perhaps copying messages from certain individuals to an additional mailbox.

You can also use transport rules to further extend the ability of message classification. For example, suppose that we need to add specific text to the subject of a message that has been marked with our custom message classification. Let's see how we can use transport rules to do this. I'm going to use the Exchange Management Console in this example. Here's what to do.

- Run the Exchange Management Console and navigate to the Organization Configuration container.
- Under the Organization Configuration container you will see the Hub Transport container. Click this and then choose the Transport Rules tab as you can see in Figure 12. Note that the Action pane has been removed for clarity.

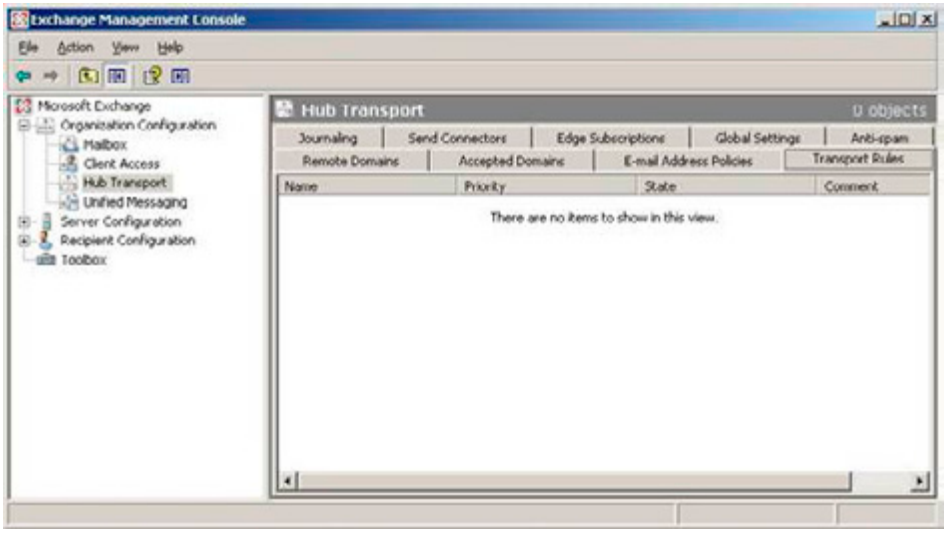

Figure 13: Transport Rules Tab.

- Right-click the Hub Transport container and choose New Transport Rule... from the context menu. This invokes the new transport rule wizard.
- On the opening screen of the transport rule wizard, give your rule a suitable name and make sure that the Enable Rule checkbox remains selected. Click Next to advance to the next screen.
- The next screen of the wizard is the Conditions screen. Here, choose the marked with classification condition in the Step 1 area of the screen. You should now see that, in the Step 2 area of the screen, the marked with classification condition has now been added. An example is shown in Figure 13 below.

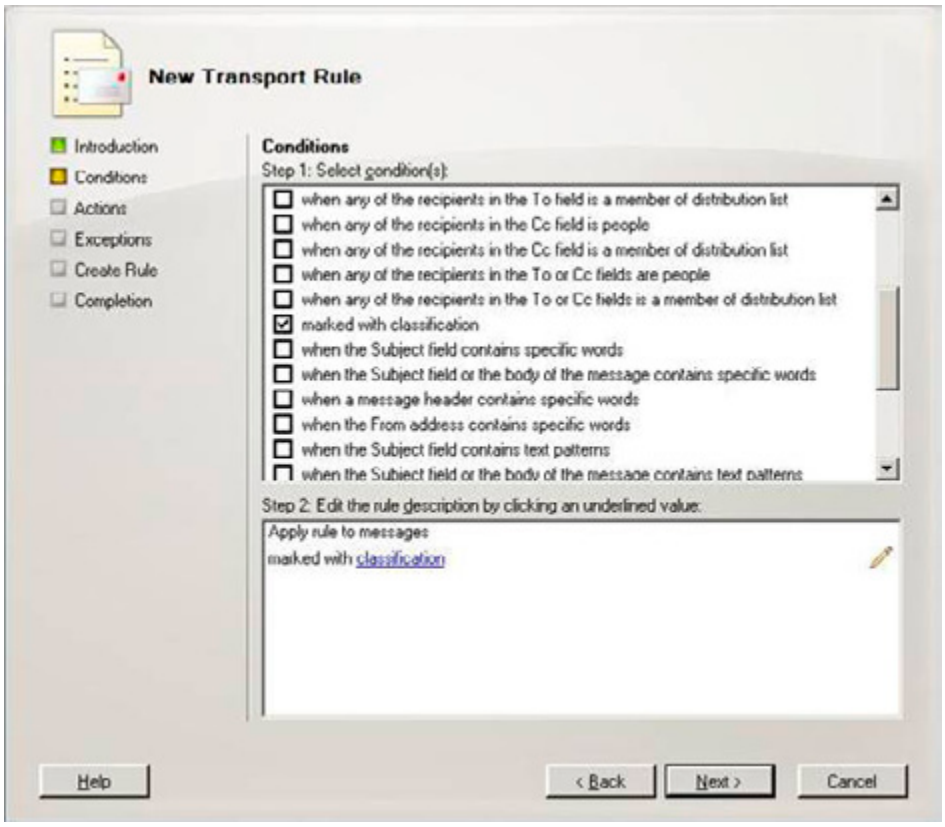

Figure 14: Transport Rule Conditions.

- In the Step 2 area of this screen, click the underlined word "classification." This brings up the Select message classification window as you can see from Figure 14. Select the relevant message classification, which in this example is the Custom Classification, and then click OK.

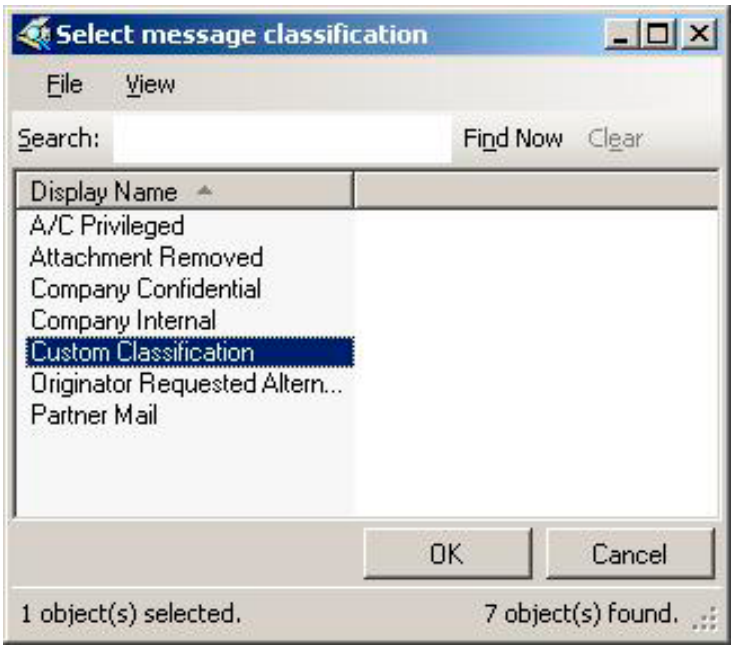

Figure 15: Select Message Classification Window.

- Back at the Conditions window of the transport rule wizard, you should now see that the Step 2 area of the screen shows your condition as marked with Custom. Click Next to proceed through the wizard.
- You are now presented with the Actions screen of the wizard. In this example we are going to add additional text to the subject line of the messages, so choose the prepend the subject with string option in the Step 1 area of the screen.
- In the Step 2 area of the Actions screen, click the underlined word "string" and in the resulting Specify subject prefix window enter your desired text to be prepended to the subject. In this example, I'm going to add the text "CUSTOM CLASSIFICATION:'. If you've done everything correctly, your Actions screen should look like the example shown in Figure 16.

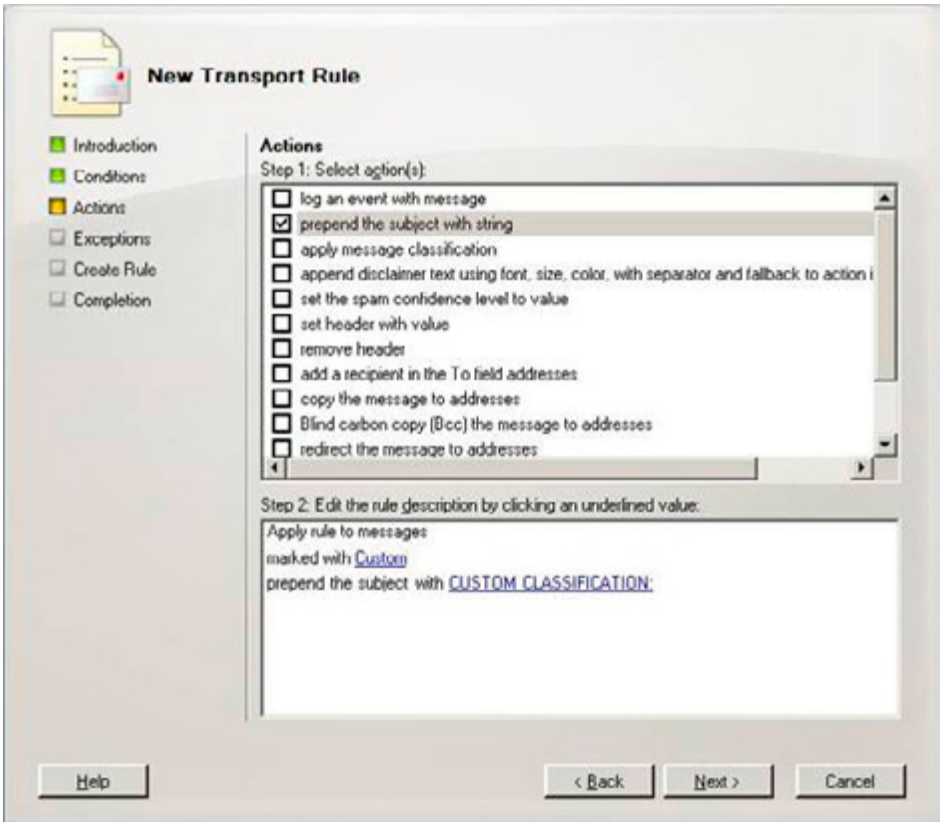

Figure 16: Transport Rule Actions.

- Clicking Next on the Actions screen takes you to the Exceptions screen where you can choose to apply exceptions to the rule. For the example within this article I'm not going to add any exceptions so I will simply click Next and move on to the next screen.
- Finally we are now at the Create Rule screen that allows you to review your selections. If you are happy with your selections, click New to create the new transport rule.
- If everything has been successful, you are presented with the Completion screen informing you of a successful creation as you can see in Figure 17.

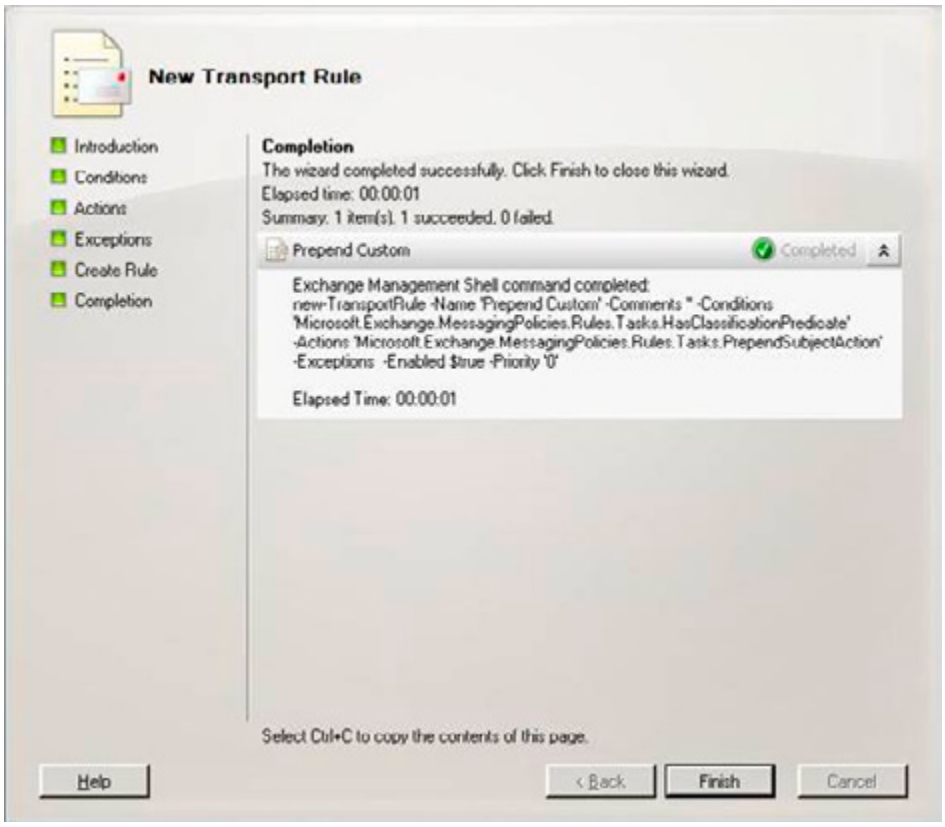

Figure 17: Transport Rule Completion Screen.

The expected outcome of this transport rule is that whenever a message is sent and is marked with the Custom message classification, the subject line of that message should be prepended with the text "CUSTOM CLASSIFICATION:" As you can see from Figure 18, the transport rule works perfectly.

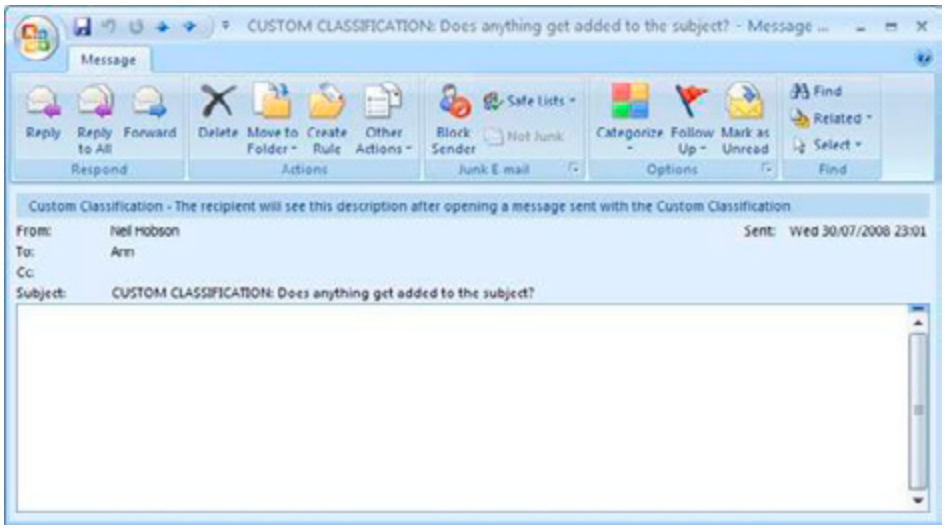

Figure 18: Transport Rule in Action.

Summary

It's now possible for the end users of an Exchange 2007 system to classify their messages such that recipients of those messages understand that there may be restrictions relating to the message content. For example, the message may contain sensitive information that should not be sent externally from the company. The flexibility of message classifications is further extended when you consider that transport rules can be created to perform specific actions on messages that have been classified by the users as you have seen in this article. In my experience not every company that deploys Exchange 2007 implements message classifications but nonetheless it is a useful and welcome addition to the Exchange 2007 feature set, particularly for those organizations that need to comply with regulations or other email policies.

Deploying Exchange 2007 on Windows Server 2008

19 September 2008

by Nicolas Blank

Nicolas Blank recounts his experiences deploying Exchange 2007 SP1 on Windows Server 2008, when not everything worked out of the box as it should have. In this article Nicolas writes about the fixes to the issues he faced when installing on Server 2008.

My customer's scenario wasn't quite typical – he had an unstable mail server running Exchange 2003, as well as Active Directory issues, one of which included the requirement to rename the directory tree. The customer wanted a brand new environment and in order to realize the scalability and security benefits of Microsoft's 64 bit OS decided on Windows Server 2008. This meant I was called in to perform a "Green Fields" migration, where a new target environment is built and all users, machines and mail are migrated to it. To complicate matters, the customer was on a tight hardware budget, meaning he could only afford a single large machine for a 200 user site.

Designing a solution

The design was relatively straightforward – since I only had a single machine available, I had to place the HUB, CAS and mailbox role onto that machine. Having all of the roles on one machine is well catered for in the available design guidance from the Exchange team at Microsoft. The machine had 16 GB of memory and a quad core processor. I also had enough disks to create a decent set of mirrors for the OS, page file, logs and a RAID 5 array for the Exchange Database. SPAM handling was done at the ISP, which meant one less burden for the HUB role to handle, since the budget did not allow for additional hardware for an edge server. AV would be handled by the ISP, though this did not preclude internal attack, and I chose Forefront to handle AV on the Exchange server to scan both existing mail in the stores and transmitted mail via the HUB role.

Building a new mail server

Server 2008 is much "lighter" on a default install than Server 2003, with fewer components deployed by default. However this default install requires me to add the Roles and Features required to build a multi role Exchange Server. Instead of adding each feature through the GUI by hand, I built a batch file containing the required commands. The only reboot required would be after the installation of Active Directory Domain Services remote management tools. From the command line I ran "ServerManagerCmd -i RSAT-ADDS" to install this service, and rebooted. After which I ran the following commands in listed order in a batch file.

```
ServerManagerCmd -i PowerShell
ServerManagerCmd -i Web-Server
ServerManagerCmd -i Web-ISAPI-Ext
ServerManagerCmd -i Web-Metabase
ServerManagerCmd -i Web-Lgcy-Mgmt-Console
ServerManagerCmd -i Web-Basic-Auth
ServerManagerCmd -i Web-Digest-Auth
ServerManagerCmd -i Web-Windows-Auth
ServerManagerCmd -i Web-Dyn-Compression
ServerManagerCmd -i RPC-over-HTTP-proxy
```

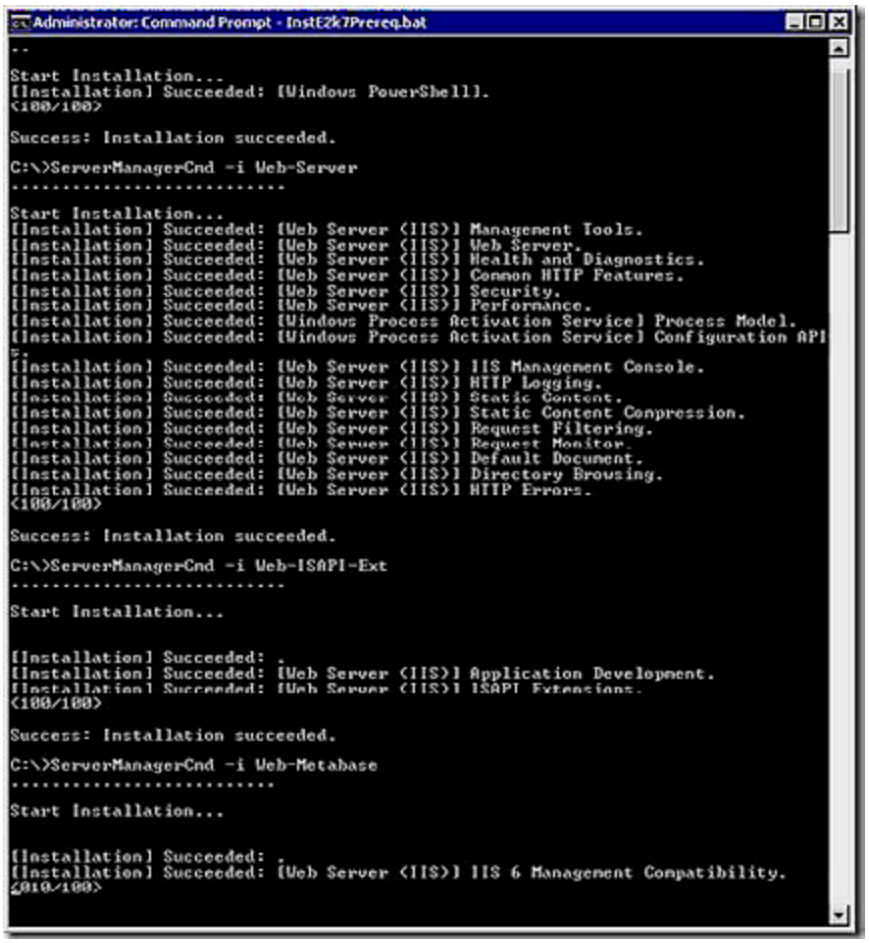

Using the command line was dramatically faster than using the GUI would have been and allowed me to script all of the required prerequisites, thereby eliminating any potential mistakes installing the prerequisites. After this I invoked the Exchange installer and since all of the prerequisites were met, Exchange had no issues installing.

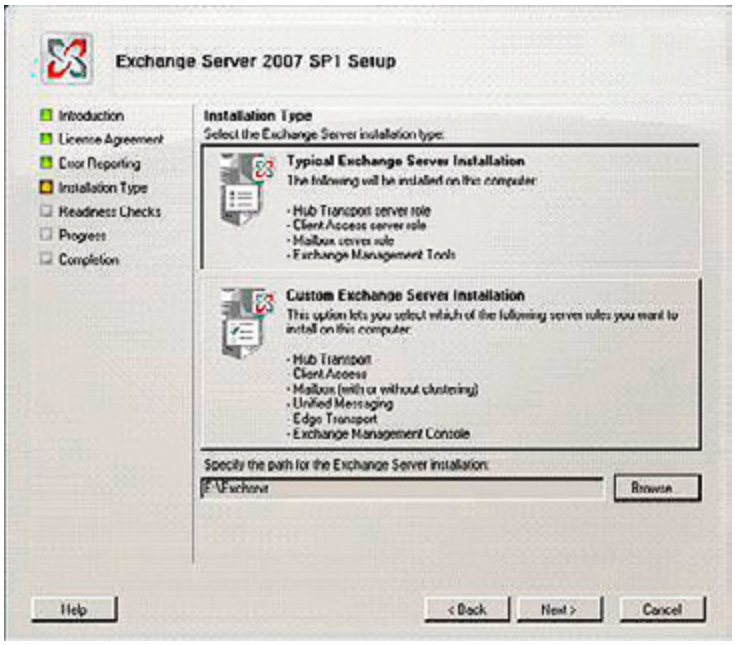
Exchange Server 2007 SP1 Setup
Introduction
License Agreement
Error Reporting
Installation Type
Readiness Checks
Progress
Completion
Installation Type
Select the Exchange Server installation type:
Typical Exchange Server Installation
The following will be installed on this computer:
- Hub Transport server role
- Client Access server role
- Mailbox server role
- Exchange Management Tools
Custom Exchange Server Installation
This option lets you select which of the following server roles you want to install on this computer:
- Hub Transport
- Client Access
- Mailbox (with or without clustering)
- Unified Messaging
- Edge Transport
- Exchange Management Console
Specify the path for the Exchange Server installation:
Browse...
Help
< Back
Next >
Cancel

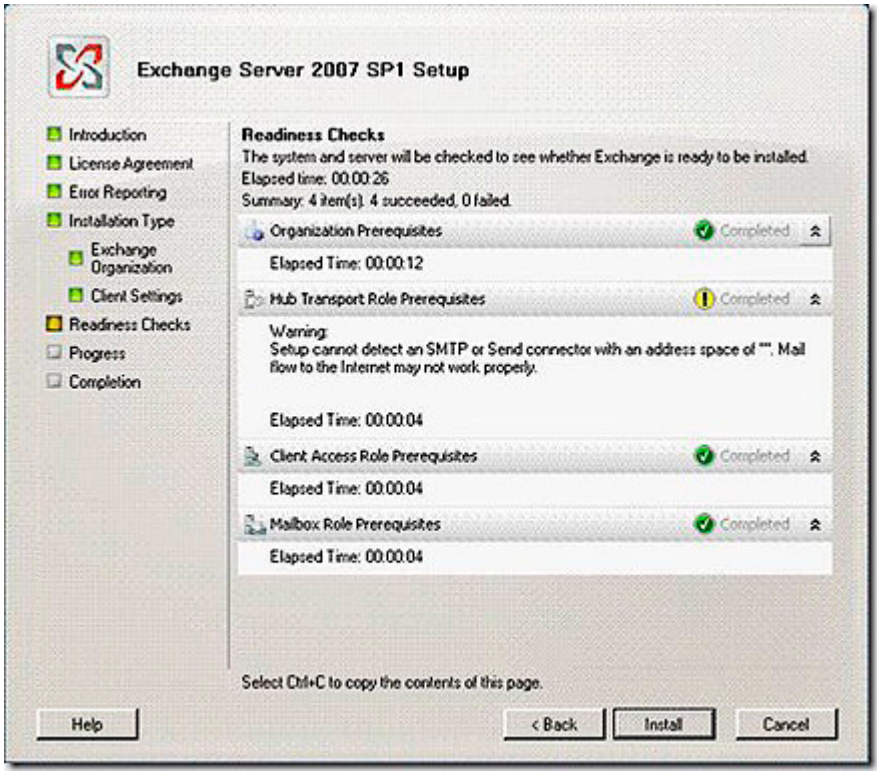
Exchange Server 2007 SP1 Setup
Introduction
License Agreement
Error Reporting
Installation Type
Exchange Organization
Client Settings
Readiness Checks
Progress
Completion
Readiness Checks
The system and server will be checked to see whether Exchange is ready to be installed.
Elapsed time: 00:00:26
Summary: 4 item(s). 4 succeeded, 0 failed.
Organization Prerequisites
Completed
Elapsed Time: 00:00:12
Hub Transport Role Prerequisites
Completed
Warning:
Setup cannot detect an SMTP or Send connector with an address space of "*". Mail flow to the Internet may not work properly.
Elapsed Time: 00:00:04
Client Access Role Prerequisites
Completed
Elapsed Time: 00:00:04
Mailbox Role Prerequisites
Completed
Elapsed Time: 00:00:04
Select Ctrl+C to copy the contents of this page.
Help
< Back
Install
Cancel

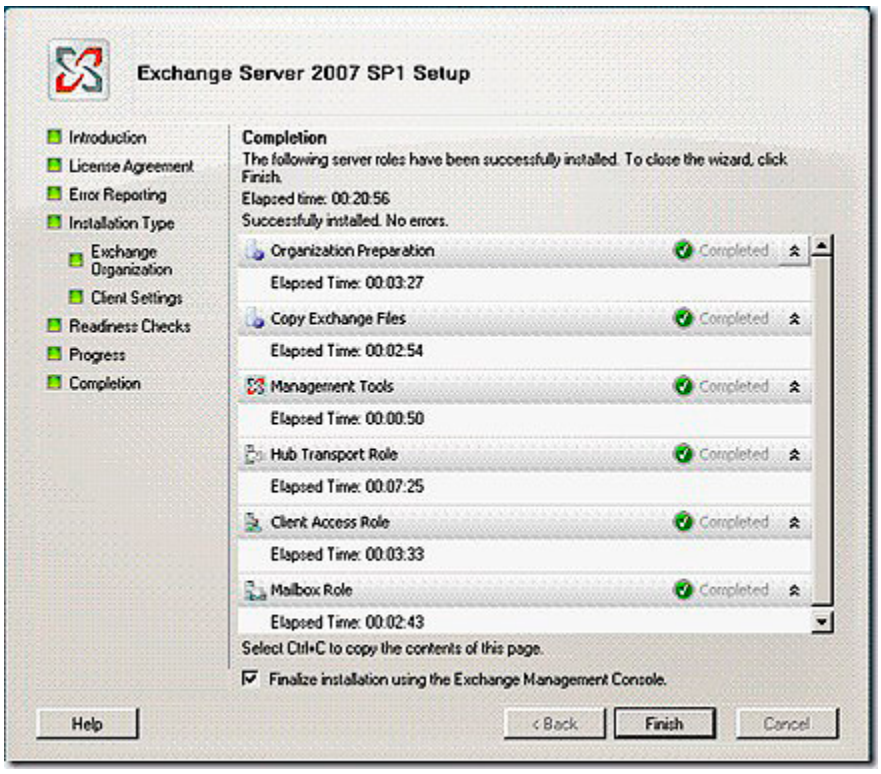

The problem, IPv6.

The last thing I added before installing Exchange was the prerequisite for the CAS role to host Outlook Anywhere and mobile clients, namely RPC over HTTP, using this command

```
ServerManagerCmd -i RPC-over-HTTP-proxy
```

I noticed that RPC over HTTP didn't always work. The solution lay in the limited support for the CAS role and IPv6. Running

```
Netstat -a -n
```

From the command line gave the following results

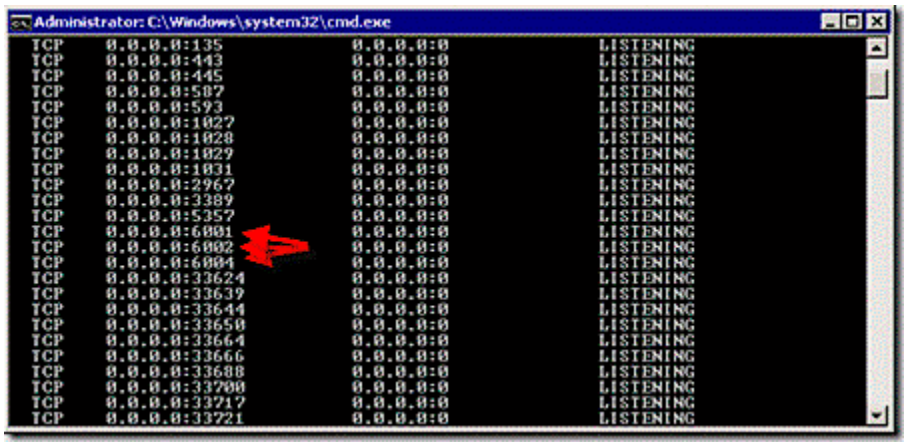

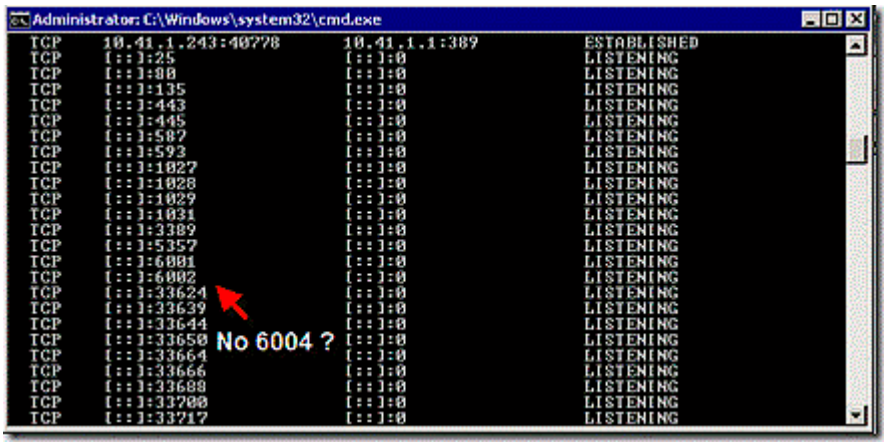

If you're familiar with IPv4, you'll know that in the first picture the IP stack is listening on open ports 6001, 6002 and 6004, but these ports were missing on the IPv6 stack on the same address "[::]:" . This meant that one of the core requirements for RPC over HTTP, communication with the local server, had been compromised. At first glance, the fix seems simple, surely you just disable IPv6? Correct, but that wasn't as easy as you might think.

First I had to unbind IPv6 from the Network Adapter, but just like Vista, Server 2008 requires a registry hack in order to disable the protocol altogether.

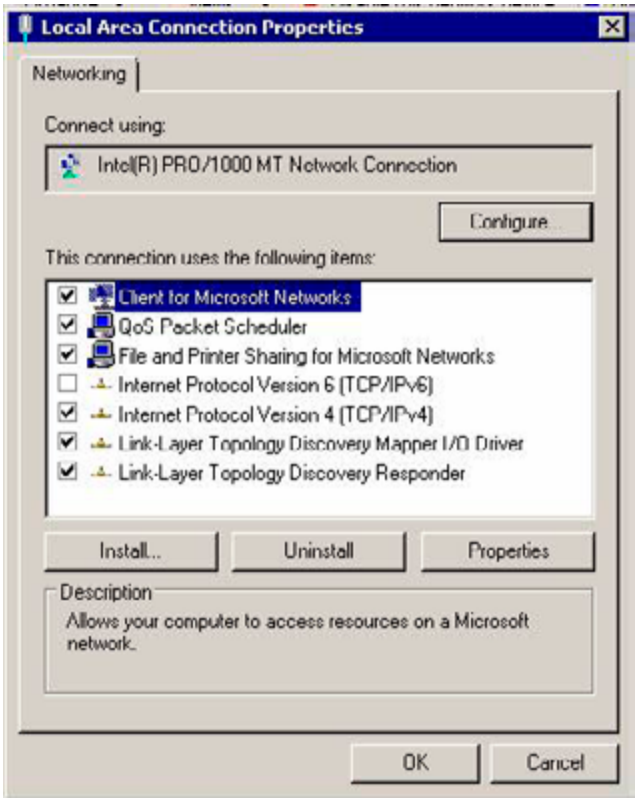

Using Regedit I navigated to: `HKLM\SYSTEM\CurrentControlSet\Services\Tcpip6\Parameters`, added a DWORD32 called DisabledComponents and gave it the following value: `0xff`, effectively disabling all IPv6 components. See THIS ARTICLE from the MS Exchange team for background.

When installing Exchange 2007 on Server 2008, using Outlook Anywhere requires using this value, but refer to the table below for other possible values.

Function	Value
Disable all IPv6 tunnel interfaces, including ISATAP, 6to4, and Teredo tunnels	0x1
Disable all 6to4-based interfaces	0x2
Disable all ISATAP-based interfaces	0x4
Disable all Teredo-based interfaces	0x8
Disable Teredo and 6to4	0xA
Disable IPv6 on non tunnel interfaces including all LAN and PPP interfaces	0x10
Disable IPv6 on all LAN, PPP, and tunnel interfaces	0x11
Prefer IPv4 to IPv6 when attempting connections	0x20
Disable IPv6 over all interfaces and prefer IPv4 to IPv6 when attempting connections	0xFF

Have a look at the IPv6 Transition Technologies Whitepaper for more details.

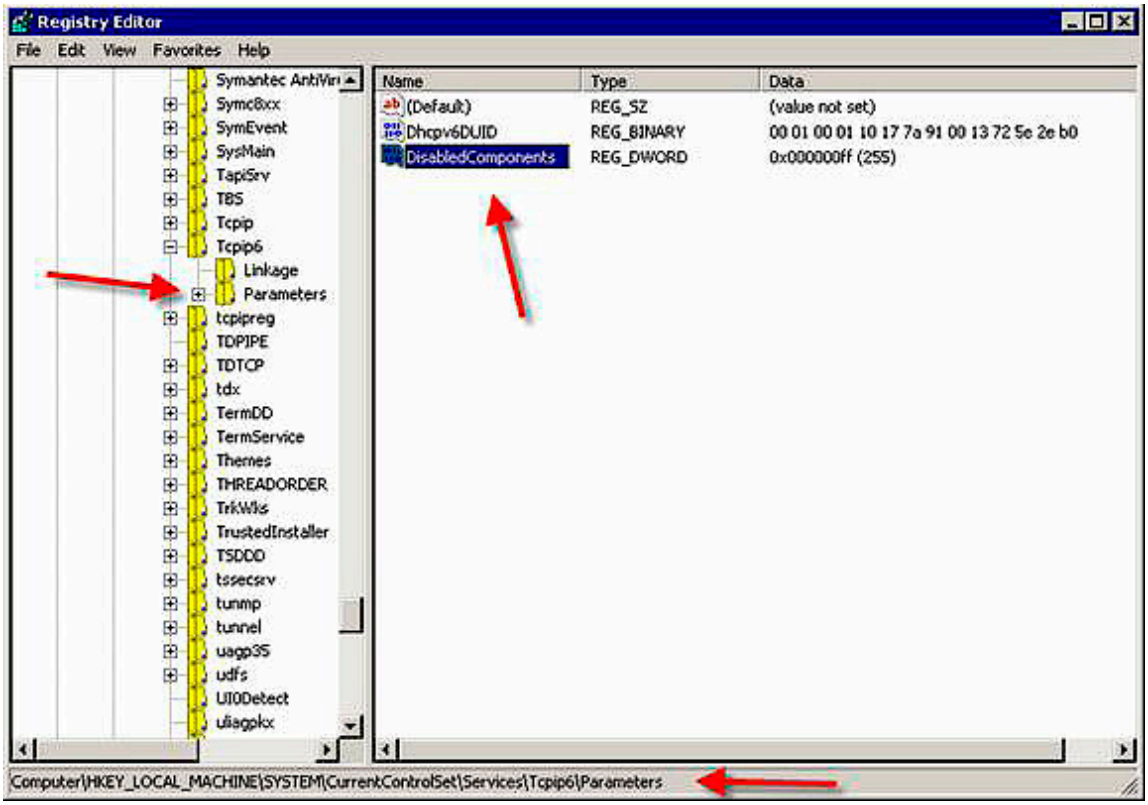

One final step was required, namely editing the hosts file to remove the IPv6 "localhost" equivalent. This meant commenting out the ::1 line by placing a # in front of it as well as manually adding the Netbios and FQDN names

10.41.1.243 mpmx01.ds.customerAD.com

10.41.1.243 mpmx01

127.0.0.1 localhost

::1 localhost

Note that the last line comments out the IPV6 address

> ***Note***
>
> *This issue is current as of Exchange 2007 SP1 Rollup 3, though it should be resolved with Rollup 4 (IF IT COMES OUT FOR REAL!). None of my Exchange customers run IPv6 ,and even with this issue resolved, I would still disable IPv6 or any other protocol not actively used in the environment. After a reboot running netstat –a –n again revealed that IPv6 was indeed gone for good.*

Finishing the Org

With that out of the way, configuring Exchange was straightforward. I added the same SMTP namespace as the original org, configured a SAN certificate for the CAS role, allowing OWA and Outlook Anywhere to communicate securely and allowed anonymous mail submission to the Receive connector, thereby enabling internet mail. Final testing showed that I could communicate with Exchange both internally and externally. The migration proceeded smoothly after that, using Quest Migration Manager (QMM) to move both the AD user accounts, and the Exchange Mailboxes. The advantage in using this toolset over native tools, was that there was virtually no user impact, and it required no desktops visits. Depending on the timeframe required, the complexity of the migration and the amount of mail that needs to be moved, I generally prefer using third party utilities to native utilities. I have had particular success with QMM, since it supports single or many object rollback. This allowed me to build Disaster Recovery plans that fitted the overall business requirement into the migration plan. Native tools can often be "fire and forget" and you have to hope that the end result is the one you hoped for.

it was worth noting that the original Exchange server suffered massive hardware failure the day after the migration completed and was signed off. The server drive subsystem failed catastrophically, requiring a complete replacement of all drives in the array. One of the original migration drivers was to move off the old hardware platform. Had the business decided to wait to migrate any longer we might have experienced the hardware failure while migrating.

Conclusion

If you get the chance to upgrade, Windows Server 2008 offers a number of enhancements in the OS which benefit Exchange 2007 deployment and management greatly. Security and resilience are enhanced and Windows ships with a better IP stack allowing more RPC connections, amongst other features. This "Green Fields" migration path is particularly straightforward, but even the more complex methods are well worth following if you have the budget. A few things remain incompatible, for example, Server 2008 contains no native backup utility for Exchange 2007, and Exchange 2007 does not support the new Read Only Domain Controller feature in Server 2008. The first of these at least is likely to change in the near future. IPv6 is irritating , but it is quickly disabled since it offers no value over IPv4 at this point. Server 2003 is still available at the time of writing, but I wouldn't hesitate to deploy Server 2008, and gain advantages such as Hyper-V support or "free" geo-clustering with CCR and SCR clusters replicating over the WAN. It is worth remembering that Exchange is a large application, making every deployment worth planning for, irrespective of which operating system it is deployed against.

Exchange Server Log File Replay

22 September 2008

by Jaap Wesselius

Exchange Server stores information in a database and uses log files for transactional processing. To restore, defragment or repair a database, the ESEUTIL tool is essential . You can always recover data when the database is lost, if you have backed up the database.

In my previous article, Exchange Database Technologies, I discussed the underlying database technology in Exchange Server, the Extensible Storage Engine or ESE. One of the most important points in that article was that all changes to the Exchange Server database go through the log files. This is done for recovery purposes. Let's look at the log files, and the replay of log files in case of a recovery scenario...

Creation of the database

After you create a Storage Group in the Exchange Management Console, we have an empty directory on disk. The only thing that actually happens is setting a property for the Storage Group in Active Directory, so no log files have been created yet.

When a database is created in the Exchange Management Console there is still an empty directory on the disk. Again, the only thing that happened is setting a property in Active Directory.

- When the database is actually mounted, these things will happen on disk.
- Active Directory is checked for the location of the log files.
- When no log files are found a new set of log files is created with an lGeneration number of 1.
- Active Directory is checked for the location of the database.
- Since this is an initial mounting a new database is created.

At this point log files, a database file and a checkpoint file have been created and the Exchange Server's database is ready to use.

Tip

In the old Backoffice 4.x resource kit there was a small utility called "mailstorm." This MAPI utility was very useful to send a massive amount of messages in a short timeframe. Using mailstorm it is possible to see the creation of log files. You can still find the original mailstorm utility here:. *Unfortunately mailstorm doesn't work with Exchange Server 2007 anymore, but there's a PowerShell script available having the same functionality. Using this script a variety of of predefined messages can be sent to your Exchange Server 2007 to check the log file creation. The PowerShell version of mailstorm can be downloaded* here:.

Create a user account in Active Directory and create a mailbox using the Exchange Management Console. Log on to the mailbox and start sending messages until you have a couple of messages and a couple of log files.

What happens when the database file is lost? If we have all the log files still available it should be possible to retrieve all information. Remember what I wrote in my previous article: **"everything is logged in the log files, even the creation of database files!"**

If you dismount the database, delete the database file "mailbox database.edb" or whatever name you have given the database and try to mount the database again, the following error is raised:

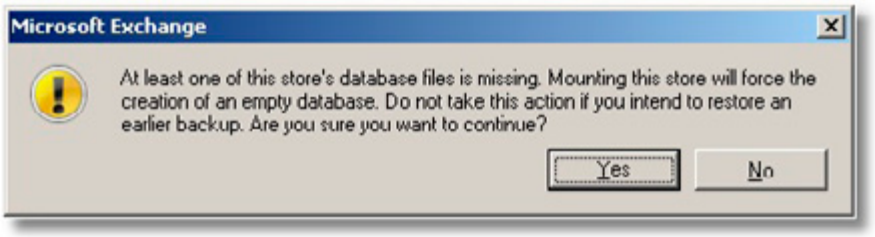

Figure 1. Error message when a mailbox database appears to be missing.

In my humble opinion, the yellow exclamation mark should be replaced with a very large red "X" since this is a very important message. When you click "Yes" a new mailbox will be created in the same location as the "old" database. This is a completely new database. Although it has the same name "*mailbox database.edb*" as in the previous step, it has new signatures, as explained in my previous article, which Exchange Server uses to link databases and log files file together. Recovery of old log files will **not** result in information being replayed into the database because it is another database in this scenario. And remember, since all information is logged into the log file the creation of this new database is also logged. Choose *No*, and then delete the checkpoint file E00.chk and try to mount the database again. No error message is raised and the database is mounted. Even better, when you log on to your test mailbox you will see that **no information is lost** either!

This is what happens when you do click "Yes": during the mount process Exchange Server cannot find the database file and it cannot find the checkpoint file. Therefore it starts to recover all information by replaying the available log files. It starts with the oldest log file E0000000001.log which also contains the creation of the initial database file. All information in the other log files is replayed into this database until it reaches the end of the last log file E00.log. The database is mounted and ready to use.

When Exchange Server cannot find the database file but it does find the checkpoint file it will not replay the log files. It starts at the end of the last log file E00.log and it will create a new database.

How can you tell which files belong together?

Dismount the Exchange Server's database, open a command prompt in the database's directory and check the header information using the ESEUTIL tool. You will find the following information:

```
K:\sg1>eseutil /mh "mailbox database.edb"

Extensible Storage Engine Utilities for Microsoft(R) Exchange Server
Version 08.01
Copyright (C) Microsoft Corporation. All Rights Reserved.

Initiating FILE DUMP mode...
     Database: mailbox database.edb

   DB Signature: Create time:08/25/2008 22:54:52 Rand:4416518 Computer:

  Log Signature: Create time:08/25/2008 22:54:50 Rand:4412688 Computer:

Operation completed successfully in 0.140 seconds.


K:\sg1>
```

When you check the header information of the log files with the ESEUTIL tool you will find corresponding information:

```
L:\sg1>eseutil /ml e000000001a.log

Extensible Storage Engine Utilities for Microsoft(R) Exchange Server
Version 08.01
Copyright (C) Microsoft Corporation. All Rights Reserved.

   Signature: Create time:08/25/2008 22:54:50 Rand:4412688 Computer:

   1 k:\sg1\Mailbox Database.edb
         Signature: Create time:08/25/2008 22:54:52 Rand:4416518 Computer:

Operation completed successfully in 0.60 seconds.

L:\sg1>
```

Note: both screen outputs have been edited for readability

As you can see both the log file signature and the database signature match, so these files belong together. When you (accidentally) create a new mailbox you will find other information in the database header:

```
L:\sg1>eseutil /ml e000000001c.log

Extensible Storage Engine Utilities for Microsoft(R) Exchange Server
Version 08.01
Copyright (C) Microsoft Corporation. All Rights Reserved.

   Signature: Create time:08/25/2008 22:54:50 Rand:4412688 Computer:

   1 k:\sg1\Mailbox Database.edb
         Signature: Create time:09/02/2008 07:27:28 Rand:963593 Computer
:
Operation completed successfully in 0.70 seconds.

L:\sg1>
```

As you can see the log signature hasn't changed (still the same set of log files) but the database has a new signature, meaning that although the database has the same name and it is in the same location it is a new database!

Key take-a-way: the checkpoint file determines if log files are replayed and where log file replay will start and therefore what happens during the mounting process.

Offline backups

You can create backups by copying the database files to a safe location. The steps to do so are:

- Dismount the database (meaning it is not available!).
- Copy the database file to a safe location.
- Mount the database.
- Perform a consistency check on the database copy.
- If everything is OK, delete the log files.

The first three steps do not need any further explanation. But what log files can you safely delete?

When dismounting the database all information in the log files that is not yet committed to the database is flushed to the database file. When all data is flushed the files are closed.

You can check in the database header information when the database was dismounted by looking at the "last detach" information. Also check that the database is in a clean shutdown and does not need any log files for mounting.

```
K:\sg1>eseutil /mh "mailbox database.edb"

Extensible Storage Engine Utilities for Microsoft(R) Exchange Server
Version 08.01
Copyright (C) Microsoft Corporation. All Rights Reserved.

      State: Clean Shutdown
   Log Required: 0-0 (0x0-0x0)
  Log Committed: 0-0 (0x0-0x0)

   Last Detach: (0x11,27,35) 09/02/2008 16:34:34

Operation completed successfully in 0.150 seconds.

K:\sg1>
```

So in this specific scenario, all log file older than E0000000011.log can safely be deleted. All information in these log files is flushed to the database. Why not log file E0000000011.log itself? There can be information logged in the log file in this same log file beyond this point after mounting the database again.

The offline copy of the database has also been checked for consistency. A database can contain corrupt pages, and as long as these pages are not read by the Exchange server you will never know these pages are corrupt. Suppose somebody is on a one year sabbatical leave and his mailbox data is never accessed, corrupt pages can exists for this period of time without being noticed.

So you have to check the offline copy for any inconsistencies using the ESEUTIL tool with the /K option:

```
K:\backup>eseutil /k "mailbox database.edb"
Extensible Storage Engine Utilities for Microsoft(R) Exchange Server
Version 08.01
Copyright (C) Microsoft Corporation. All Rights Reserved.

Initiating CHECKSUM mode...
    Database: mailbox database.edb
 Temp. Database: TEMPCHKSUM5592.EDB

File: mailbox database.edb

           Checksum Status (% complete)e)

     0  10  20  30  40  50  60  70  80  90 1000000
     |----|----|----|----|----|----|----|----|----|----|
     ..................................................


1282 pages seen
0 bad checksums
0 correctable checksums
388 uninitialized pages
0 wrong page numbers
```

```
0x7dd5 highest dbtime (pgno 0x193)

161 reads performed
10 MB read
1 seconds taken
10 MB/second
17350 milliseconds used
107 milliseconds per read
280 milliseconds for the slowest read
0 milliseconds for the fastest read

Operation completed successfully in 0.471 seconds.
K:\backup>
```

When using the /K option all pages in the database are read and their checksum information is read and checked. When everything is ok, you can safely continue.

Note

Starting with Exchange Server 2003 Service Pack 1 Exchange has error correcting code for checksum errors. If a certain page contains a checksum error, which is usually caused by only one bit being "flipped," Exchange Server can automatically correct this bit. When the page is flushed to disk the error is automatically corrected. Please check the Microsoft kb article 867626 (Microsoft kb article 867626) for more background information regarding this change.

Offline Restore

When something happens we have to restore the offline backup on the Exchange server. An offline restore requires the following steps.

- Copy the offline backup mailbox file to the correct location.
- Replay the log files
- Mount the database.

Now remember the first part of this article. If we leave the checkpoint file in its original location the Exchange Server will check this file and will determine that no replay is needed and the copy of the database will be mounted.

If we delete the checkpoint file the Exchange Server will start replaying log files from the oldest log file available. This will result in all information after the offline backup being replayed into the copy of the database. The result will be a complete up-to-date database.

If you check the event log of the Exchange server you can see which log files were replayed. For every log file an entry is logged with ID 301 from the ESE Source.

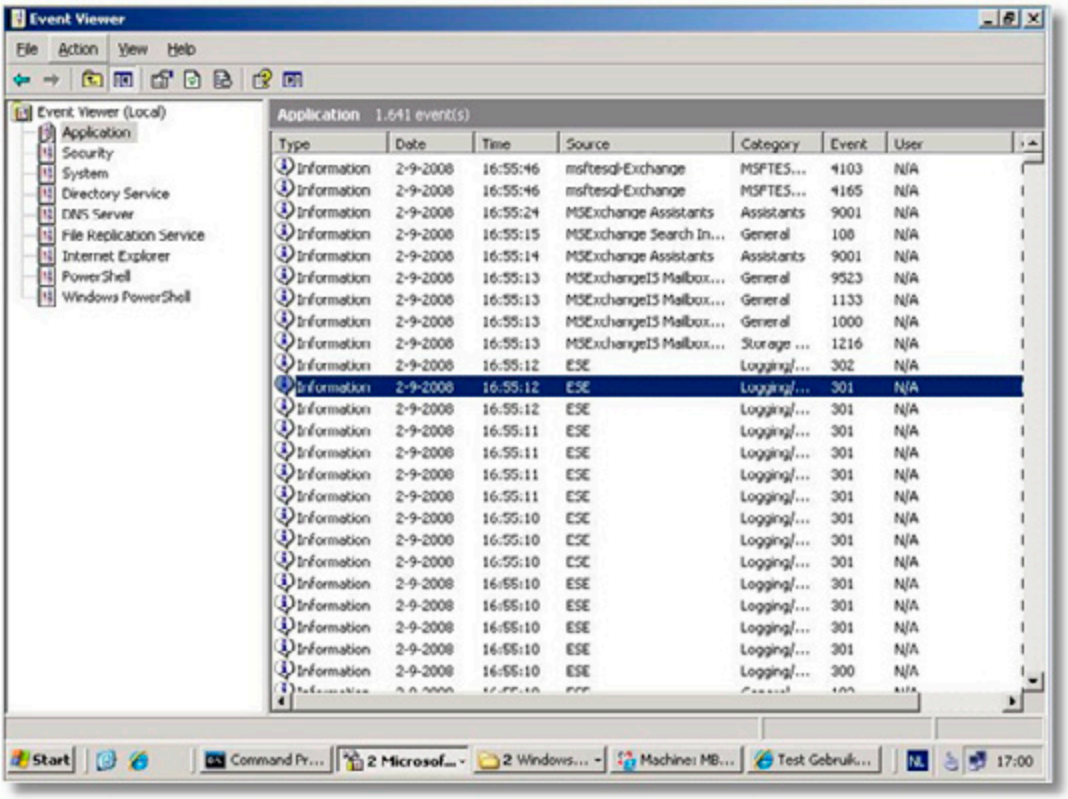

Figure 2. Log file replay events in the event viewer.

Besides just mounting the database it is also possible to replay the log files into the database manually using the ESEUTIL tool. This can be achieved using the /R option for Recovery.

```
L:\sg1>eseutil /R E00 /i /dK:\SG1

Extensible Storage Engine Utilities for Microsoft(R) Exchange Server
Version 08.01
Copyright (C) Microsoft Corporation. All Rights Reserved.

Initiating RECOVERY mode...
  Logfile base name: E00
      Log files: <current directory>t;t;
    System files: <current directory>t;t;
  Database Directory: K:\SG1

Performing soft recovery...
           Restore Status (% complete)

     0  10  20  30  40  50  60  70  80  90 100
     |----|----|----|----|----|----|----|----|----|----|
     ..................................................

Operation completed successfully in 5.819 seconds.

L:\sg1>
```

Note. the drive letter K: immediately follows the /d option, there's no space in between.

Online Maintenance

There's an internal process in the Exchange server that's responsible for the maintenance of the mailbox database. This is called the online maintenance, or OLM. OLM runs every night, by default between 1 am and 5 am. Online maintenance consists of 10 tasks that are performed during this period. Five of them are specifically for the Public Folder database and five of them are for the Public Folder database and the mailbox database. I want to highlight two of them in this article:

OLM checks the mailbox database for deleted messages that are beyond their retention time. If there are, they are fully deleted from the database and their pages are freed up (that is, marked as "white space"). OLM is also responsible for permanently deleting mailboxes that are beyond their retention time. If messages and mailboxes are beyond their retention time they are really gone, the only way to get these back is by restoring an old backup of the mailbox database. Archiving solutions can help here, but they are out of the scope of this article. When the mailboxes and messages are permanently deleted the pages they used are freed up and available for new messages. But the free pages are scattered across the database. The OLM also performs an online defragmentation of the database file by consolidating freed up pages from the previous step and if possible, placing them in one contiguous block of free data. This makes the addition of new data to the database more efficient since the data doesn't have to be split and separated across multiple segments of the database.

Online Maintenance is a very disk intensive application and it should be monitored very closely. If the Exchange server cannot finish the OLM in the designated timeframe errors are logged in the event log. Carefully examine these errors and check what causes them.

Make sure that no other processes are working with the mailbox database at the same time. Typically maintenance processes are scheduled at night, such as backups or archiving. I will cover online backups in the next article in this series. The movement of large numbers of mailboxes, which is also often scheduled during the night, will also have a negative impact on the OLM.

For a complete overview of the online maintenance please visit the Microsoft Exchange team blog.

Key take-away: Online Maintenance is responsible for online defragmentation; it defragments inside the database only. This does not result in a decrease of the physical size of the mailbox database.

Offline defragmentation

Please note that offline defragmentation should never be considered a part of normal maintenance of an Exchange database. Only if your free space exceeds 30% of your total database size, or if you are told to do a defragmentation by Microsoft Customer Support Services, should you consider doing an offline defragmentation.

While the online maintenance does defragmentation within in the database, it does not compact the database in physical size. If you have a 100 GB database and it has only 10 GB of data in it, the online maintenance will eventually free up 90 GB within the database (it will be marked as "white space" or "available space" within the logical tables of the database). The file "mailbox database.edb" as it resides on disk still remains 100 GB. To compact the database we have to perform an offline defragmentation. This can only be done using the ESEUTIL tool with the /d option. As the name implies, this has to be done when the database is offline.

```
K:\sg1>eseutil /d "mailbox database.edb"

Extensible Storage Engine Utilities for Microsoft(R) Exchange Server
Version 08.01
Copyright (C) Microsoft Corporation. All Rights Reserved.

Initiating DEFRAGMENTATION mode...
      Database: mailbox database.edb

         Defragmentation Status (% complete)
```

```
     0  10  20  30  40  50  60  70  80  90 100000000
     |----|----|----|----|----|----|----|----|----|----|
     ..................................................


Moving 'TEMPDFRG2256.EDB' to 'Mailbox Database.edb'... DONE!

Note:
 It is recommended that you immediately perform a full backup
 of this database. If you restore a backup made before the
 defragmentation, the database will be rolled back to the state
 it was in at the time of that backup.

Operation completed successfully in 8.913 seconds.

K:\sg1>
```

If you look closely at this output you can already see what happens during an offline defragmentation. A new mailbox database file is created is with a random name, in this case TEMPDFRG2256.EDB. The data from the original mailbox database is copied to this temporary database. Of course existing white space in the original database is not copied to the new database. When finished copying the temporary file is renamed to the original file and is now called "mailbox database.edb." So this is a new file, containing only the database that resided in the original file. In the previous example with the 100 GB database with only 10 GB of data, when an offline defragmentation is performed we will be left with a new 10 GB database. There is some overhead in this, so the sizing can vary +/- 10%.

You must realise that this is a new database with new signatures etc. so there is no possibility of recovering theto recover data in existing log files into this new database. This is the reason you should make a new backup immediately after performing an offline defragmentation.

Please note that if you do not redirect the temporary file created by Offline Defragmentation, the disk volume containing the Exchange mailbox database should have at least 110% of the size of the mailbox database for available disk space.

In the Exchange Server 5.5 timeframe it was a best practice to perform an offline defragmentation every months. This would create a new database with new tables, indices etc. Microsoft has spent a tremendous amount of development in the database engine. Nowadays the database engine is that good and stable that an offline defrag has only to be performed after a repair action, or when large amounts of data can be freed up. This can be the case after the deletion of a large number of mailboxes or after moving a lot of mailboxes to another database.

Worst case scenario: Repairing your database

The worst thing that can happen when running Exchange is that your server crashes and your database refuses to mount. After checking if the Information Store is running and maybe rebooting the server it still won't mount and the following error is raised:

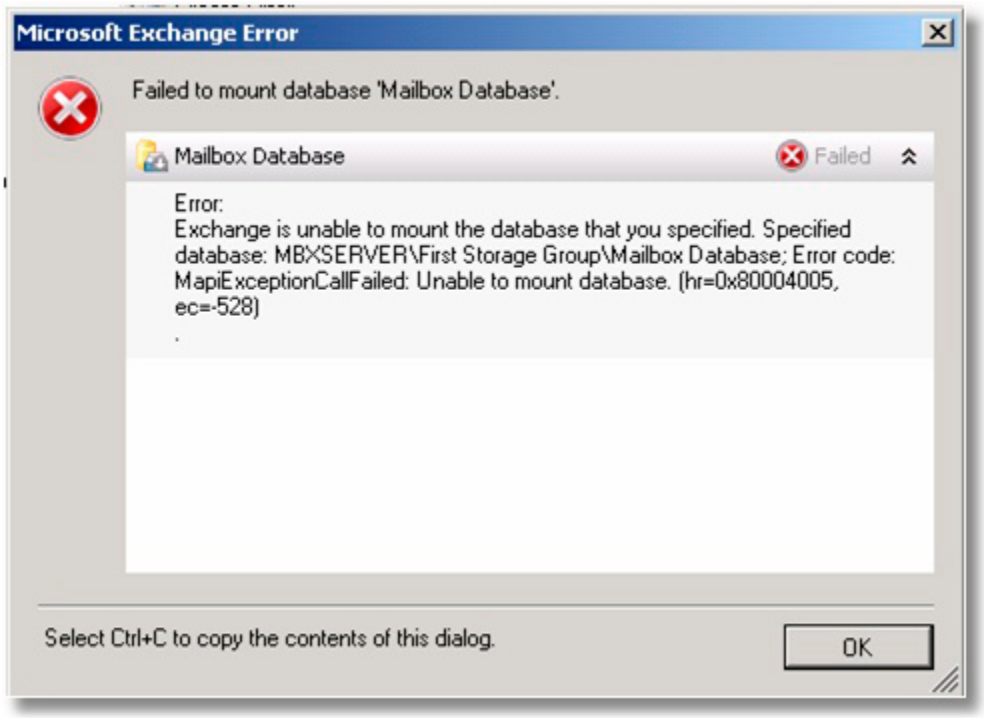

Figure 3. Error message when the database won't mount after a crash.

In this example one of the log files needed for the mounting process is not available anymore (I deleted one of these log files for demonstration). If this occurs it can be checked by comparing the values in the "logs needed" section of the database header against the actual log files that are on disk.

So now we have a database that is not consistent anymore and that has to be fixed before it can be used again. To accomplish the ESEUTIL tool should be used again but now with the /P option (for repair).

Note. Only perform this step when you have made a copy of all databases, log files and checkpoint file and placed them in a safe location. Performing a repair action can result in data loss!

When you enter this command an error is raised stating that if you continue information might be lost. What actually happens is that ESEUTIL checks all information in the database and it checks all tables and pointers. If pointers do not point to the correct location, or information in this location cannot be read, the pointer is deleted and the index entry is cleaned up. All information that the pointer points to will be lost!

```
K:\sg1>eseutil /p "mailbox database.edb"

Extensible Storage Engine Utilities for Microsoft(R) Exchange Server
Version 08.01
Copyright (C) Microsoft Corporation. All Rights Reserved.

Initiating REPAIR mode...
    Database: mailbox database.edb
```

```
 Temp. Database: TEMPREPAIR1888.EDB

Checking database integrity.

The database is not up-to-date. This operation may find that
this database is corrupt because data from the log files has
yet to be placed in the database.

To ensure the database is up-to-date please use the 'Recovery' operation.


           Scanning Status (% complete)

     0  10  20  30  40  50  60  70  80  90 100
     |----|----|----|----|----|----|----|----|----|----|
     ..................................................


Integrity check successful.

Note:
 It is recommended that you immediately perform a full backup
 of this database. If you restore a backup made before the
 repair, the database will be rolled back to the state
 it was in at the time of that backup.

Operation completed successfully in 150.747 seconds.


K:\sg1>
```

At the end of this operation you will have a database that is in a consistent state, but indexes in the database might be non-contiguous and thus less efficient..As a Microsoft best-practice an offline defragmentation should be performed after a database repair. This will create a new database, with the existing data but with new tables, indices etc. Both ESEUTIL /P and ESEUTIL /D will fix issues on a low (database) level but not on an application level. After performing a repair and on offline defragmentation on your database you have to run the ISINTEG tool. This will check the database on an application (=logical) level and if issues are found fix them. When completely finished do not forget to backup your database immediately.

> ***Note***
>
> *Perform these steps only if you have backed up the database(s), log files and the checkpoint file since performing a repair action can lead to data loss!*

To be 100% sure that all issues are fixed on all levels you can create a new mailbox database and move all mailboxes from the repaired database to the new database. When all mailboxes are moved you can delete the repaired database. But please be aware that a repair action with ESEUTIL can delete data, so despite all activities there might be circumstances that end users mail can be lost. Unfortunately it is not predictable what and how much will be lost.

You can find more information on the Microsoft website: http://technet.microsoft.com/en-us/library/aa997152(EXCHG.80).aspx and http://msexchangeteam.com/archive/2004/06/18/159413.aspx.

Note

Of course it is also possible to restore a backup of the database and recover the existing log files up to the point of the missing log files. If possible this is most likely the best and fastest option you will have.

When using ESEUTIL, there is one very important issue. ESEUTIL is very powerful, but not very fast. Depending on the hardware in use for the Exchange server, as a rule of thumb ESEUTIL will process the Exchange database with a rate of between 5 GB/hour and 10 GB/hour. When you have a 250 GB Exchange database, a repair action on your database would take 25 hours or longer to complete A following offline defragmentation can also take up to 25 hours, so you will see an outage of at least 50 hours. I don't take the ISINTEG tool and moving the mailboxes to a complete new database into account, but I'm pretty confident that this will not match your Service Level Agreement (SLA). On an Exchange server using a single copy database, i.e. a single server or a single copy cluster, it is best to use a maximum database size of 50 GB. This will keep your database in a manageable state from a service level point of view.

So why does Microsoft suggest using a maximum 200 GB database size in an Exchange Server 2007 Continuous Cluster Replication (CCR) environment? As explained in my first article, you have an extra copy of your database available in your Exchange 2007 cluster. So when the database of the active node crashes and is beyond recovery (i.e. cannot be mounted anymore) the passive copy will take over and start servicing your (Outlook) clients. When this happens you have time to start investigating the broken node of your cluster and repair it. You are not tied to the 5 GB/hour processing time of ESEUTIL, which ought to make you more comfortable with your Service Level Agreement.

Conclusion

Exchange Server stores information in a database and uses log files for transactional processing. This way it is always possible to recover data when the database is lost, if you have backed up the database. The tool for database maintenance is ESEUTIL. Using ESEUTIL you can check your database, recover data from the log files and replay it into your database or repair your database when it's corrupt.

When you need to compact your database, for example after you deleted or moved multiple mailboxes, you have to perform an offline defragmentation, but remember to create a new backup after performing an offline defragmentation.

One important thing to remember is the checkpoint file. Although a tiny file only 8KB in size it has a major impact on the results when either recovering data using ESEUTIL or mounting the database in the Exchange Management Console.

When you run into issues in your test environment that isn't a problem. You can and you should start working with database and ESETUIL to see what exactly happens in the different scenarios. Make sure you fully understand (and document!) all the necessary steps for a successful recovery when you lose your database. However, do this in your TEST environment, not in your live/production environment.

When problems happen in your production environment and you face a major outage, ask Microsoft Customer Support for help. Microsoft CSS will guide you through a successful recovery, but when you have some experience in your test environment you know at least what's happening and why Microsoft Customer Support asks you to perform the various steps.

In my last article I will talk about online backups and VSS backups and explain the differences between these and the steps documented in this article.

Configuring Exchange Server 2007 to Support Information Rights Management

02 October 2008

by Brien Posey

In Exchange Server 2007, Information Rights management is easy to set up once you have set up the prerequisites. It is also much cheaper, and easier to use. This is just as well, because of the increasing statutory regulations to prevent the mishandling of confidential information in emails.

Exchange Server has supported the use of rights management since the first days of Exchange Server 2003. Even so, the Exchange Server 2003 implementation tended to be a bit annoying for users. When users attempted to open messages subject to rights management, they were prompted to enter their authentication credentials every time, prior to being able to view the message. As you can imagine, this seriously reduces efficiency if a user needs to open a significant number of rights managed messages.

Fortunately, things have improved in Exchange Server 2007. Users who are running at least Microsoft Office Outlook 2007 or Windows Mobile 6.0 now have the ability to open rights managed messages without being prompted to enter their credentials. In this article, I will explain why this is possible, and how to configure rights management in an Exchange Server 2007 environment.

The Rights Management Server

As mentioned before, rights management has been available for quite some time. Prior to the release of Windows Server 2008, Microsoft's rights management solution was a product named Rights Management Services (RMS). Aside from the Exchange Server issues that I mentioned earlier, Rights Management Services was a good product. The only problem was that it came with a price tag that put it out of reach for many small and medium-sized companies.

Since the time that rights management was first introduced, there have been several new federal regulations introduced that affect the way that certain types of data must be handled. Since many of these regulations provide stiff penalties for the mishandling or accidental disclosure of data, Microsoft decided to incorporate Right Management Services directly into Windows Server 2008, rather than requiring customers to buy an entirely separate product. Before you get too excited though, keep in mind that there is a Client Access License requirement associated with using Rights Management Services.

New Rights Management Features

Obviously the greatest change in Rights Management Services is that it is no longer a separate product from Windows Server. Aside from this, there have been a substantial number of other changes that you need to be aware of.

One very welcome change is the new self enrollment feature. The previous version of Rights Management Services required administrators to connect to a special Microsoft enrollment server via the Internet. The enrollment server would issue a Server Licensor Certificate (better known as a SLC). The SLC gave the server the right to issue licenses.

The new self enrollment feature does away with this requirement. Windows Server 2008 uses the self-enrollment certificate to sign the SLC without the aid of an external certificate Authority.

Another change to Rights Management Services is that it now supports the delegation of roles. You can use these roles to provide IT staff with varying degrees of administrative control over the Rights Management Server. The table below lists the roles that are available for use:

Role Name	Function
AD RMS Service Group	The AD RMS Services Group is simply a group that contains the AD RMS service account. When you install the Rights Management Service, the service account that you specify during the installation process is automatically added to this group.
AD RMS Enterprise Administrator	As the name implies, AD RMS Enterprise Administrators have all of the necessary rights to manage rights management related policies and settings.
AD RMS Template Administrator	AD RMS Template Administrators have permission to create new rights management templates, or to modify existing templates. They do not have the rights to modify non template related server settings.
AD RMS Auditor	The AD RMS Auditor role is a role that Microsoft created in order to help organizations to comply with various federal regulations. It is a read only administrative role that allows an auditor to view rights management logs and create reports.

The other new Rights Management Services feature is its integration with the Active Directory Federation Service (ADFS). ADFS is beyond the scope of this article, but I wanted to at least mention it.

Rights Management Functions

A Windows Server 2008 Rights Management Server provides three primary functions. Exchange Server 2007 is a rights managed application, which means that it utilizes these three primary functions, but also builds on them.

The first and most obvious function of RMS is that it allows authorized users to create rights managed files and templates. In order to perform this function, the application that is used to produce the rights protected document that is to be rights protected must also support rights management. Fortunately, all of the Office 2007 applications (including Outlook 2007) support rights management.

This might be a good time to point out that rights management is different from Digital Rights Management (DRM). Rights management, as I am referring to it, is a mechanism for encrypting documents (or in this case, e-mail messages), and assigning permissions to control who can do what with the protected document. DRM, on the other hand, is basically a copy protection mechanism that is used mostly by content providers. For example, many websites that offer online video use DRM to prevent subscribers from downloading all of the site's content, and then canceling their subscriptions.

The second primary function of the Rights Management Server is its ability to identify trusted users and groups. The Rights Management Server accomplishes this task by issuing a rights account certificate to the users and groups who have been granted permission to create rights protected content.

The third primary function of the Rights Management Server involves allowing the use of protected content. If a user wants to open a protected document, they must possess the proper rights account certificate. Beyond that, the Rights Management Server must cross reference the user or group with Active Directory in order to determine whether or not they have the rights to decrypt the protected document, and what the document's policies allow them to do with the document (modify, copy, print, etc.).

Rights Management Services and Exchange Server 2007

As I mentioned earlier, Exchange Server 2003 worked with RMS, but it was a pain to use. When Outlook users opened a protected message, they were prompted to enter their credentials, even though they had already been authenticated by Active Directory during their initial login. Typically, Windows Mobile users were completely unable to open rights managed E-mail messages, because they could not connect to Rights Management Services.

Rights Management Service support works a bit differently in Exchange Server 2007 than it did in Exchange Server 2003. When you install the Hub Transport role on an Exchange 2007 server, the installation process automatically installs a component called the AD RMS Prelicensing Agent. The basic idea behind this component is that it certifies the identity of Outlook 2007 and Windows Mobile 6.x users at the Hub Transport level. That way, Exchange is able to perform the various rights management related tasks before the user attempts to open a protected message. This means that Outlook no longer needs to prompt users for their credentials every time that they attempt to open a rights managed message. This new architecture also means that it is even possible for users to open RMS protected messages when they are working offline, something business travelers are sure to appreciate.

Some Caveats to Be Aware of

Although Exchange Server 2007 automatically installs the Prelicensing Agent on any Hub Transport servers, there are a few gotchas that you need to be aware of. For starters, the Prelicensing Agent is only included in Exchange Server 2007 SP1. Therefore, if you are installing the RTM version of Exchange Server 2007 onto a Hub Transport server, the Prelicensing Agent will not be installed.

Another caveat that you need to be aware of is that the Prelicensing Agent does not work with older versions of Rights Management Services. If you've got a legacy RMS Server in place, you can use it, but you will have to install SP2 for Rights Management Services 1.0. Of course, if you are using the Windows Server 2008 version of RMS (often referred to as AD RMS or AD RMS 2008), then you are good to go.

An additional requirement that you must be aware of is that if your organization has multiple Hub Transport servers, then each of those Hub Transport servers must be running Exchange Server 2007 SP1 or higher. The reason for this is that the Prelicensing Agent works at the Hub Transport server level. If multiple Hub Transport servers exist, then any one of them could potentially be called upon to deliver a protected message, and must therefore contain the Prelicensing Agent, which of course is only available in SP1.

Preparing to Use the Prelicensing Agent

So far I have talked about some prerequisites for installing the Prelicensing Agent, but there are a couple more requirements that must be met before your organization can use the Prelicensing Agent. The first of these requirements is that there is a client component that must be installed on your Hub Transport servers. The client component that you must use depends on the version of Windows Server that the Hub Transport server is running on.

If your Hub Transport server is running Windows Server 2003, then you will have to install the Windows RMS Client 1.0 SP2. This client will only run on 64-bit versions of Windows Server 2003 (and Windows XP). That shouldn't be a problem though, since you should never run the 32-bit version of Exchange Server 2007 in a production environment. You can download the client component at: Microsoft . Windows Rights Management Services Client with Service Pack 2 - X64 Edition. .

If you are running your Hub Transport server on Windows Server 2008, then the server must be running the Active Directory Rights Management Client. Fortunately, this client component is installed by default on 64-bit versions of Windows Server 2008.

The other thing that you will have to do prior to using the Prelicensing Agent is to enable it. Although Exchange Server 2007 SP1 installs the Prelicensing Agent by default, it does not enable the agent by default. Thankfully, the technique for enabling the Prelicensing Agent is fairly simple. In order to enable the Prelicensing Agent though, you will have to log into the Hub Transport server using a domain account

that is both a member of the Exchange Organization Administrators group and a member of the local Administrators group for the Hub Transport server.

Once you have logged in, open the Exchange Management Shell. I recommend getting started by verifying the Prelicensing Agent's status. You can do so by entering the following command as shown in Figure A:

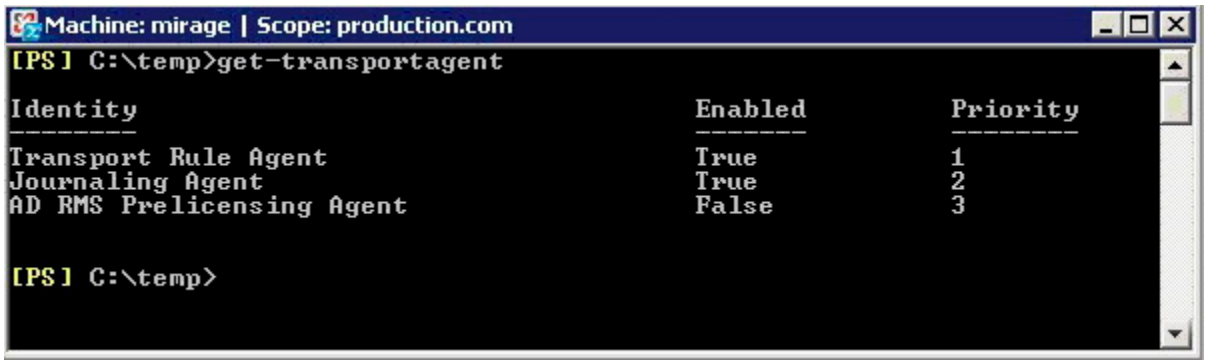

Figure A: Executing Get-TransportAgent.

```
Get-TransportAgent
```

Keep in mind that this command may also report the status of some unrelated transport agents, such as the Transport Rule Agent or the Journaling Agent.

Now that you have verified the Prelicensing Agent's status, you can enable it. Once again, you will perform this task through the Exchange Management Shell. The process involves enabling the Prelicensing Agent, and then stopping and then restarting the Exchange Transport Agent. After the Exchange Transport Agent starts back up, I recommend checking the Prelicensing Agent's status one more time, just to make sure that it is enabled. This entire process can be completed by entering the following four commands:

```
Enable-TransportAgent "AD RMS Prelicensing Agent"
Net Stop MSExchangeTransport
Net Start MSExchangeTransport
Get-TransportAgent
```

Using Microsoft Outlook with Rights Management Services

Both Outlook 2003 and Outlook 2007 can use rights management, but for the purposes of this article, I am going to limit my discussion to talking about Outlook 2007. Having said that, there are two different ways in which your end users can use Outlook to send rights managed messages. One method involves using your Rights Management Server, and the other involves using a Rights Management Server that is provided by Microsoft.

This can actually be a bit confusing for users, because Outlook is designed to use Microsoft's Rights Management Server by default. To see what I am talking about, open Outlook 2007, and compose a new message. Now, click on the Microsoft Office button (sometimes called the orb or the Jewel) that is located in the upper left corner of the New Message window. Next, choose the Permissions command from the resulting window. When you do, you will be given the option of assigning the Do Not Forward attribute to the message that you are composing, or of managing the message's credentials. If you select either of these options though, then you will be prompted to sign up for Microsoft's free Rights Management Service, as shown in Figure B.

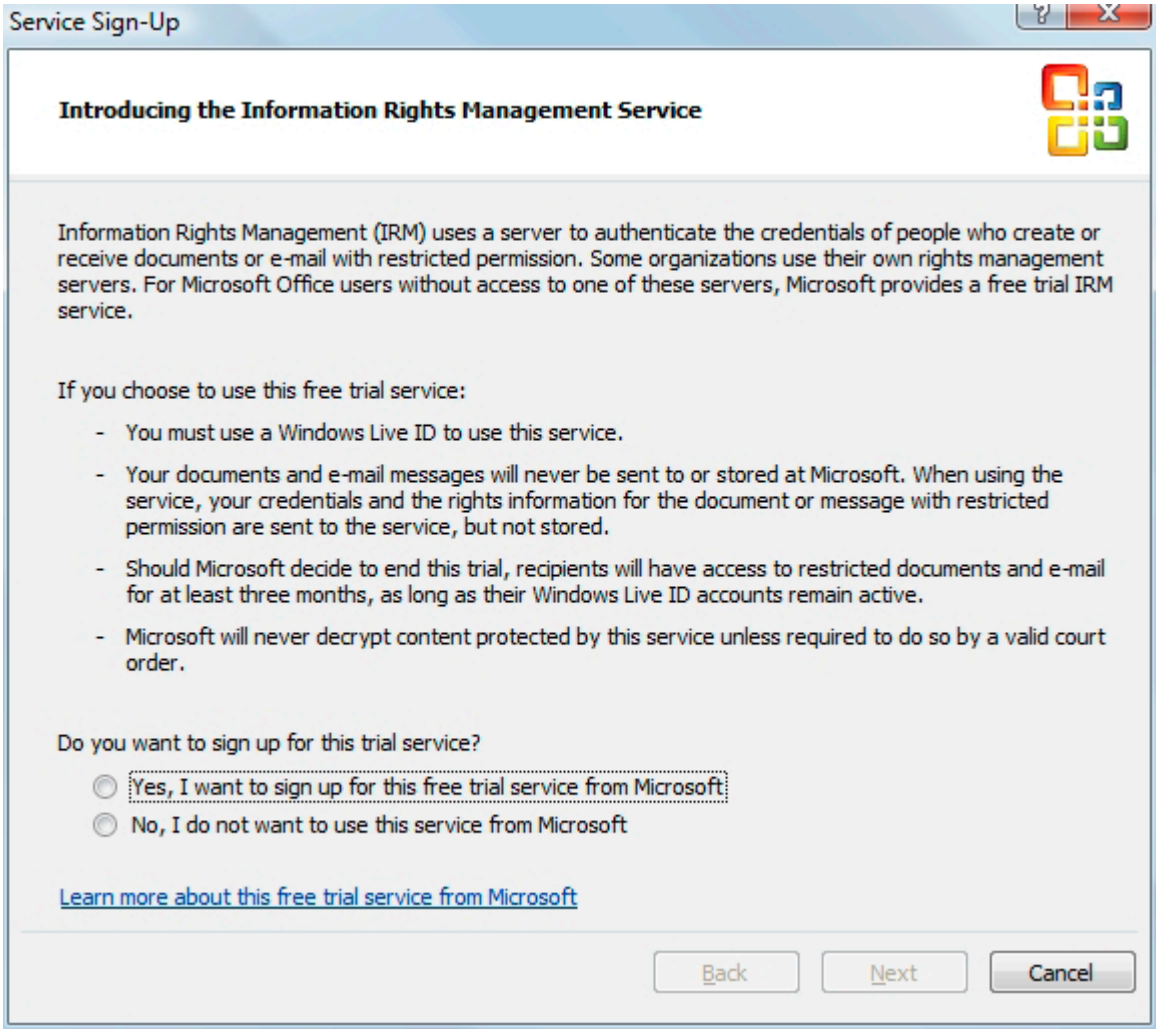

Figure B: Sign-up screen for Rights Management Service.

Of course this raises an interesting question. Why should you bother with the hassle of setting up your own Rights Management Server if Microsoft offers one that you can use for free? Well, there are a couple of reasons why it is better to have your own Rights Management Server.

For starters, if you read the text on the dialog box that I talked about earlier, it indicates that this is a "trial" service that Microsoft may choose to discontinue at any time. Granted, this trial service has been available for at least a couple of years now (I can't recall exactly when Microsoft first started offering it), but you never know when Microsoft may choose to end the trial program.

Microsoft has indicated that if they decide to discontinue the trial program, then users will continue to have access to rights managed documents for three months after the program ends. From a business prospective, it is simply a bad idea to encrypt e-mail messages using a service that may be discontinued at the drop of a hat. This is especially true since business needs or federal regulations often mandate the long term storage of e-mail messages. If a message is important enough to protect with RMS, then there is a good chance that it may also be important enough to hang on to for an extended period of time.

If you are still thinking of using Microsoft's free RMS service, remember that they have only committed to making documents accessible for three months after the service ends. They have not committed to notifying you that the service is no longer being offered. In all likelihood, Microsoft will probably notify their customers, if or when the free RMS service is ended, but you just never know whether or not your spam filter could snag the notification. Microsoft probably isn't going to wait around for you to confirm that you have received notification of the service's termination. More than likely, you would simply lose access to protected documents and messages after three months.

Even if the trial RMS service were guaranteed to go on forever, there is still a good reason why you should use your own RMS server rather than the free one. Remember that part of RMS's job is to authenticate each user's identity. In the Exchange Server 2007 implementation of RMS, this is done at the Hub Transport server level. In the free version of RMS, Microsoft uses Microsoft Passports as the means of identifying users.

The reason why this is such a big deal is because Microsoft Passports are completely outside of your control. You have no way of centrally managing Microsoft Passports or of requiring your users to use a standard passport name. Besides that, there are so many people in the world using Microsoft Passports, that it would be easy to inadvertently assign permissions to the incorrect person by making a typo.

There is one thing about the free version of RMS that is better than the commercial version. The free version is global in nature. If you wanted to send a rights protected message to a recipient who works for another company, there is nothing stopping you from doing so, as long as you both have Microsoft Passports and are registered with Microsoft's RMS trial program.

The Windows Server 2008 version of AD RMS also allows you to send rights protected messages to recipients outside of your Exchange organization, but doing so is much more complicated. In order to accomplish this, both organizations must have AD RMS servers in place, and you must use Active Directory Federation Services (AD FS) to establish a relationship between the two organizations. Only then can you send rights protected messages between the two organizations.

Assuming that you choose to use the commercial version of Rights Management Services, you may have to install a client component. As you may recall, earlier I mentioned that AD RMS would not work unless you installed a client component onto the Hub Transport server. This same client component is also required for any computer that will be using Outlook to send or receive protected messages.

The good news is that if your workstations are running Windows Vista then you already have everything that you need, because the client component for RMS is installed by default. If you have workstations that are running Windows XP, then you will need the SP2 version of the Windows Rights Management Services Client. I provided a link to the 64-bit version earlier in this article, but if you need the 32-bit version, you can download it at: Microsoft Windows Rights Management Services Client with Service Pack 2 - x86.

Conclusion

In this article, I have explained that configuring Exchange Server 2007 to support rights management is a fairly simple process, although there are a number of prerequisites that must be met. If you need additional assistance with the initial deployment and configuration of AD RMS in Windows Server 2008 then I recommend checking out my book *The Real MCTS / MCITP Upgrading Your MCSA on Windows Server 2003 to Windows Server 2008* (Syngress ISBN 978-1-59749-236-2). The book is designed to serve as a study guide for Microsoft's 70-648 certification exam, but it covers AD RMS deployment step by step.

Reporting on Mobile Device Activity Using Exchange 2007 ActiveSync Logs

10 October 2008

by Ben Lye

Top Tips for SysAdmins No 1.

In this new column giving practical advice on all things Sys Admin related, Ben Lye takes on the often difficult task of keeping track of mobile device activity.

I was recently asked to generate reports on how many users are using mobile devices to access their Exchange mailbox, what kind of devices are being used, and how that use has changed over time. Fortunately, Exchange 2007 includes a PowerShell cmdlet which will parse the IIS log files on a client access server and produce CSV output files detailing the Exchange ActiveSync usage. So, with a small amount of effort it's possible to extract the relevant data from the Exchange logs and produce some interesting reports.

The command for exporting ActiveSync logs is intuitively called **Export–ActiveSyncLog**. It takes an IIS log file as input, and generates six CSV files as output.

Output Filename	Description
Users.csv	ActiveSync activity by user, including items sent and received, as well as device type and ID
Servers.csv	ActiveSync activity by client access server
Hourly.csv	Hour-by-hour report of ActiveSync activity
StatusCodes.csv	Summary of the HTTP response codes issued in response to ActiveSync requests
PolicyCompliance.csv	Report on device compliance with ActiveSync policy
UserAgents.csv	Summary of the different user agents used to access Exchange

For my purposes Users.csv is the most interesting part of the output as it can be used to identify who the users are, which device types are the most popular, and how much use the service is getting. *It's worth noting that the data in the reports is taken from the server's perspective, so "Total Emails Sent" refers to the number of emails that the server sent to the client device.*

In an Exchange environment with multiple client access servers (such as an environment with servers in multiple Active Directory sites, or one using an internet-facing network-load-balancing array) you will need to export the logs from all client access servers which mobile devices connect to. If you have a single client access server exposed to the internet which all mobile devices connect to, you'll only need to export the logs from that one.

To use Export-ActiveSyncLog you need:

- The Exchange Server Administrator role
- Read-only access to the directory that contains the IIS log files

This example will export the ActiveSync data from the IIS log file of September 1st 2008. It will use UTC times, and will put the output in C:\Temp\EASReports.

```
Export-ActiveSyncLog -FileName "C:\Windows\System32\LogFiles\W2SVC1\ex080901.log" -UseGMT:$true -
OutputPath "C:\Temp\EASReports"
```

That will work fine for a single log file, but what if you need to export multiple log files? Well, you can list all the log files in a directory using **Get-ChildItem,** which you can in turn pipe to the Export-ActiveSync command:

```
Get-ChildItem "C:\Windows\System32\LogFiles\W3SVC1" | Export-ActiveSyncLog -UseGMT:$true -
OutputPath "C:\Temp\EASReports"
```

This syntax will combine the data from each log file and give you produce a single set of CSV files covering the entire range of the input log files. Because I need to be able to report on usage over time this approach won't give me what I need.

Another way to process multiple log files is to produce a set of CSV files for each log file. However because the CSV files would typically all use the same names I also need to specify a prefix for the name of the output CSV files, which will ensure I get all the output. For that I use the -OutputPrefix parameter of the Export-ActiveSyncLog cmdlet.

This command will create CSV files prefixed with the name of the log file they were generated from:

```
Get-ChildItem "C:\Windows\System32\LogFiles\W3SVC1" | ForEach { Export-ActiveSyncLog -FileName
$_.FullName -OutputPath "C:\Temp\EASReports" -OutputPrefix $_.Name.Replace("."log,""_")
-UseGMT:$true}
```

Now that I have the CSV files for all my log files I can import the data into a database and run reports. For the database I have an SQL database which consists of a single table based on the Users.csv file, with the addition of an ID field as the primary key, and a date field to store the date of the log file.

Getting data from PowerShell into the database is a little bit more complicated. This PowerShell script will import all the Users.csv log files which were exported with the previous command into the SQL database.

```
# Script for importing Exchange ActiveSync Users.csv files into a SQL database

# Set up the parameters for connecting to the SQL database
$dbserver = "dbserver.company.com"
$dbname = "EASReports"
$dbuser = "dbusername"
$dbpass = "dbpassword"

# Create the ADO database object
$objConnection = New-Object -comobject ADODB.Connection

# Open the database connection
$objConnection.Open("PROVIDER=SQLOLEDB;DATA SOURCE=$dbserver;UID=$dbuser;PWD=$dbpass;DATABASE=$
dbname")

# Find all the Users.csv files and import them
Get-ChildItem "C:\Temp\EASReports\*Users.csv" | ForEach {
   # Get the date from the name of the file
   $Date = ($_.Name).SubString(2,6)
   $Year = "20" + $Date.SubString(0,2)
   $Month = $Date.SubString(2,2)
   $Day = $Date.SubString(4,2)
   $Date = Get-Date -Year $Year -Month $Month -Day $Day -Hour 0 -Minute 0 -Second 0

   # Import the CSV file
   $CSVFile = Import-Csv $_
```

```
    # Get the column names from the first line of the CSV file
    $CSVFileProperties = Get-Content "$_" -totalcount 1 | % {$_.split(,"")}

    # Loop through each entry in the CSV file
    ForEach ($Entry in $CSVFile) {

       # Ignore lines with an empty Device ID
       If ($Entry."Device ID" -ne "") {
          # Construct the SQL insert statement
          $SQLString = "INSERT INTO Users ("
          Foreach ($Prop in $CSVFileProperties) {
             $SQLString = $SQLString + "[$Prop],"
          }

          $SQLString = $SQLString + "[Date]) VALUES ("

          Foreach ($Prop in $CSVFileProperties) {
             $SQLString = $SQLString + "'" + $Entry."$Prop" + "',"
          }

          $SQLString = $SQLString + "'$Date')"

          # Add the record to the database
          $null = $objConnection.Execute($SQLString)
       }

    }
}

# Close the database connection
$objConnection.Close()
```

With the data in an SQL database I can then use Excel to connect to the database and analyze the data. The resulting output looks like this:

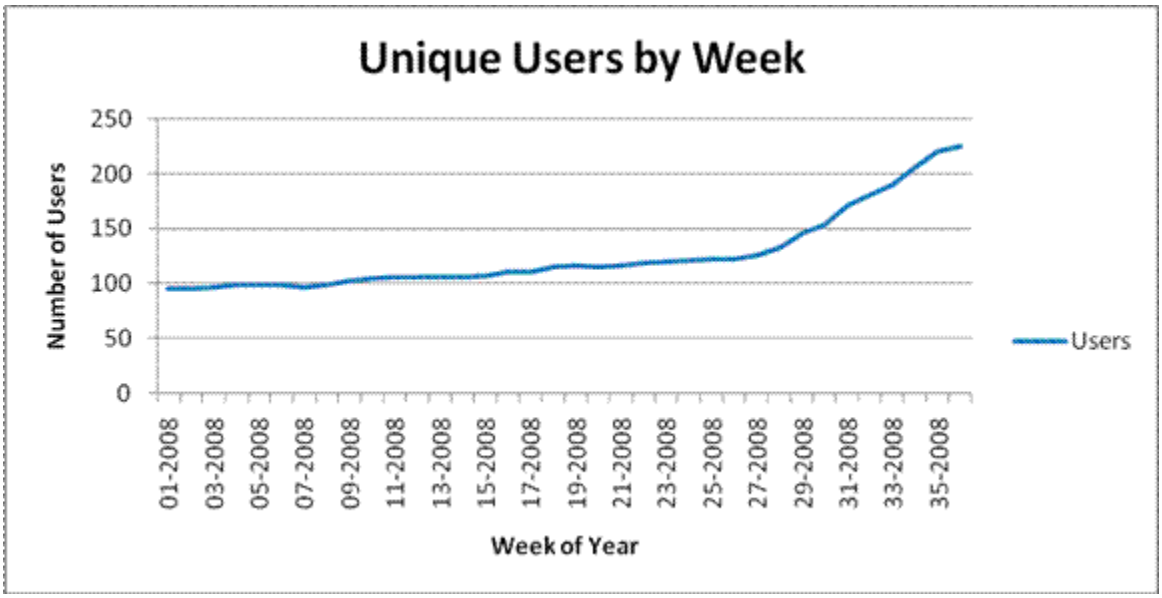

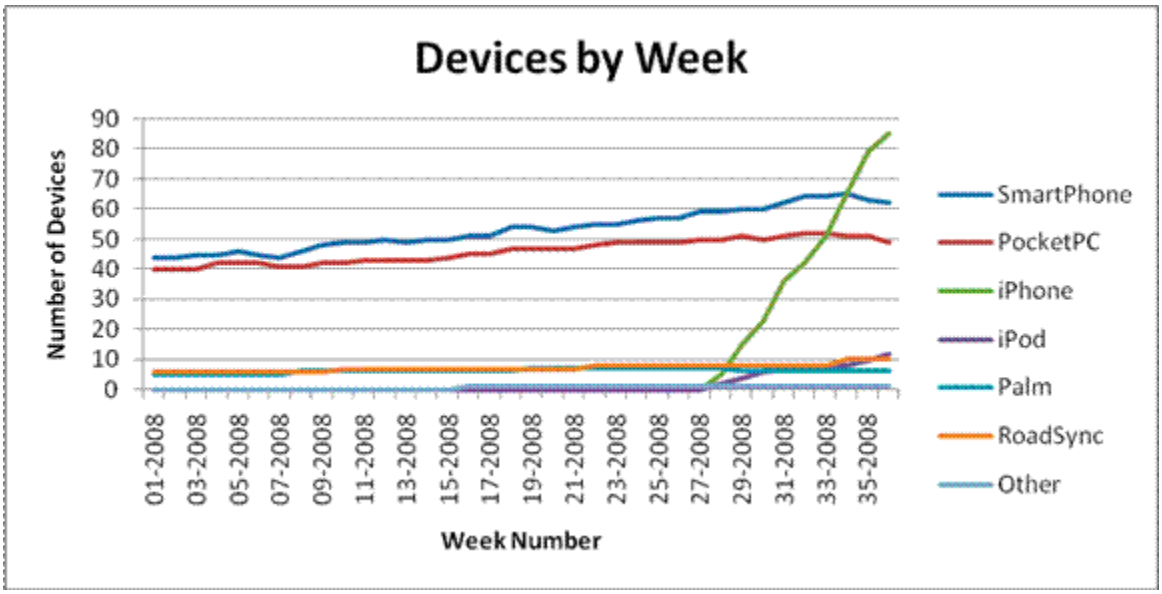

As I said at the beginning, it takes a small amount of effort to extract the data and get it into a format suitable for long-term reports, but once the pieces are in place it's a relatively simple task.

More information on the Export-ActiveSyncLog cmdlet can be found on the Microsoft Exchange TechNet website can be found on the Microsoft Exchange TechNet website.

Online Exchange Backups

18 October 2008

by Jaap Wesselius

> ***In the third of Jaap's popular series on Exchange Backups, (see Exchange Database Technologies, and Exchange Server Log File Replay), he explains Online backups. These have major advantages over offline backups since the backup application does all the logic, and work, for you.***

In my previous article, Exchange Server Log File Replay, I explained how to create an offline backup, how to check the database for consistency, how to purge the log files in a safe way and how to restore an offline backup. The major disadvantages are that the database has to be offline, it's a lot of work and you need a very good understanding what you're doing.

A much better solution is to use online backups. Online backups do all the work, check the database and purge the log files. And they do this all online, so there's no interruption for the end users.

NTBackup

Windows NT, Windows 2000 and Windows 2003 have a neat little backup utility called NTBackup. NTBackup is a lightweight version of a (very old) Veritas BackupExec version. But it's cheap and it does precisely what we can expect from a backup solution.

When installing Exchange Server on a Windows 2003 Server the NTBackup application is extended with two ESE DLL's which makes it possible to create online backups from the Exchange Server. The process is the same for all backwards versions of Windows and all backwards versions of Exchange Server. Unfortunately Windows Server 2008 is configured with a new backup application (it's a feature that has to be installed separately) that can create snap-shot backups of your Windows Server. It does not contain a plug-in for Exchange Server, so when you run Exchange Server 2007 on a Windows Server 2008 machine you have to buy a 3rd-party application to backup your Exchange databases.

In NTBackup and all other streaming backup applications there are four types of backups:

Full backup – makes a backup of the entire mailbox database and purges the log files.

Incremental backup – only the changes made since the last full backup are backed up. Since all changes are written to the log files only the log files since the last full backup are backed up. When finished they are purged.

Differential backup – only the changes made since the last full backup are backed up, but the log files are not purged.

Copy backup – this is the same as a full backup, but it does not interfere with your normal backup cycle, i.e. the header information is not updated and the log files are not purged.

Creating a full backup

NTBackup creates a backup at the ESE level. This means it accesses the database through the ESE engine and not on the file level. When opening NTBackup you have to select the Microsoft Information Store. Do not select the file "**mailbox database.edb**" from the disk. Although this will put your mailbox database in the backup set it will not do the necessary maintenance and it does not prepare for restoring your file.

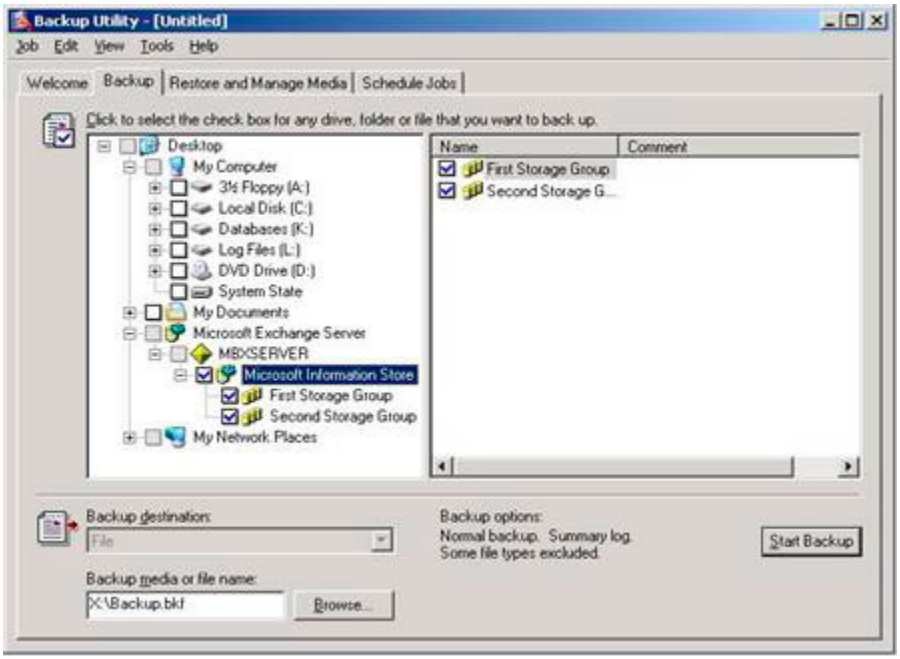

Figure 1. For backing up the Exchange databases select the "Microsoft Information Store."

When you start the backup the following things are happening:

- The current position of the checkpoint file is recorded. This location is needed for purging purposes; all log files older than the current log file can be deleted when the backup is finished. This location is recorded in the header of the database in the "Current Full Backup" section.

- NTBackup begins the backup operations and it starts reading the pages in the database. It starts with the first page and continues until it reaches the last page at the end of the database. During this operation all pages are checked for their checksum before they are streamed to the backup media. If a page fails its checksum test the backup operation is aborted and an error is written to the Windows Event log. New pages that are created during the online backup will still be flushed to the database, even when they are flushed to a portion of the database that already has been backed up. This is no problem since all transactions are still written to the log files and thus create a recovery path. During a restore the Exchange server will correct this during a so called "hard recovery'.

- When all database pages are written to the backup media the database is safe. All information that's written to the log files needs to be written to the backup media as well. To achieve this, a "log file rollover" is forced. This means that the current log file E00.log is closed (or E01, E02 etc., depending on the storage group), a new log file is created and the lGeneration number (check my first article on the lGeneration number) is increased. The log files from the point recorded in step 1 until the last log file created in step 3 are now written to the backup media. Because of the log file "roll over" you will never see the E00.log file in the backup set.

- All log files prior to the point recorded in step 1 are now purged from the log file disk.

- The "previous full backup" section of the database is now updated with the last information when the backup was running.

- NTBackup is now finished with the backup of the database.

When checking the database header after creating a streaming backup you will see something like this (with irrelevant information removed):

```
K:\sg1>eseutil /mh "mailbox database.edb"

Extensible Storage Engine Utilities for Microsoft(R) Exchange Server
Version 08.01
Copyright (C) Microsoft Corporation. All Rights Reserved.

Initiating FILE DUMP mode...
     Database: mailbox database.edb

Previous Full Backup:
    Log Gen: 34-35 (0x22-0x23)
      Mark: (0x23,1D,183)
      Mark: 09/16/2008 07:50:19

Previous Incremental Backup:
    Log Gen: 0-0 (0x0-0x0)
      Mark: (0x0,0,0)
      Mark: 00/00/1900 00:00:00

Operation completed successfully in 0.101 seconds.

K:\sg1>
```

Creating an incremental backup

An incremental backup can only be created if a previous full backup is performed on the Exchange database. The process of NTBackup creating an incremental backup is as follows:

- The backup session is initialized and the ESE engine is contacted. The location of the checkpoint file is logged in the Current Incremental Backup section.
- A log file roll-over is performed, forcing a new log file to be created.
- All log files up to the new log file are written to tape.
- All log files are purged from the disk.
- The Current Incremental Header section of the database is updated.
- NTBackup is now finished.

If you check the header information of the database again you will see both the header information of the full backup as well as from the incremental backup:

```
K:\sg1>eseutil /mh "mailbox database.edb"

Extensible Storage Engine Utilities for Microsoft(R) Exchange Server
Version 08.01
Copyright (C) Microsoft Corporation. All Rights Reserved.

Initiating FILE DUMP mode...
```

```
     Database: mailbox database.edb

Previous Full Backup:
    Log Gen: 34-35 (0x22-0x23)
      Mark: (0x23,1D,183)
      Mark: 09/16/2008 07:50:19

Previous Incremental Backup:
    Log Gen: 34-52 (0x22-0x34)
      Mark: (0x35,8,16)
      Mark: 09/16/2008 19:16:45

Operation completed successfully in 0.0 seconds.

K:\sg1>
```

Note

The process for a differential backup is identical to an incremental backup, except that the log files are not purged.

-1018 Errors

One of the tasks performed by a streaming backup is a checksum check on all pages being streamed to tape. As explained in my previous article an incorrect page can be in the database, but as long as the page isn't touched by the Exchange server you would never know. If the backup touches the page it sees that the checksum isn't correct. What will happen depends on the version of Exchange:

The release version of Exchange Server 2003 and earlier
The original release of Exchange Server 2003 (and earlier) do not contain any error correcting code for CRC errors. If an issue is detected the backup will fail and an error with Event ID 474 is logged in the eventlog. In the description you can read:
"The read operation will fail with error -1018 (0xfffffc06)."
As per Microsoft knowledgebase article 867626 you have to perform a database repair action or restore the database from a previous backup.

Exchange Server 2003 SP1 and later
Service pack 1 for Exchange Server 2003 contains error correcting code for checksum errors and thus can handle database pages that have an incorrect checksum. A streaming backup will check the pages and will notice any inconsistencies. Owing to the error correcting code in SP1 the backup application will continue, but will fix the page and write a notice in the eventlog with Event ID 399:
"Information Store (2672) First Storage Group: The database page read from the file "G:\SG1\priv1.edb" at offset 6324224 (0x0000000000608000) for 4096 (0x00001000) bytes failed verification. Bit 24032 was corrupted and has been corrected."

A lot of detailed information can be found in the Microsoft Exchange Server 2003 SDK Documentation which can be found here. Also Jerry Cochran's book "Mission-Critical Microsoft Exchange 2003: Designing and Building Reliable Exchange Servers" is a very valuable source of information as it describes backup technologies from a programming perspective. It is sold for example via Amazon.

Online restore

When performing an offline backup you have to take care of the database, the log files, and the checkpoint file and take all the necessary steps in the correct order. An online backup does all the dirty work.

When performing an online restore make sure that you mark the database for being overwritten. This is a property of the database and can be set using the Exchange Management Console.

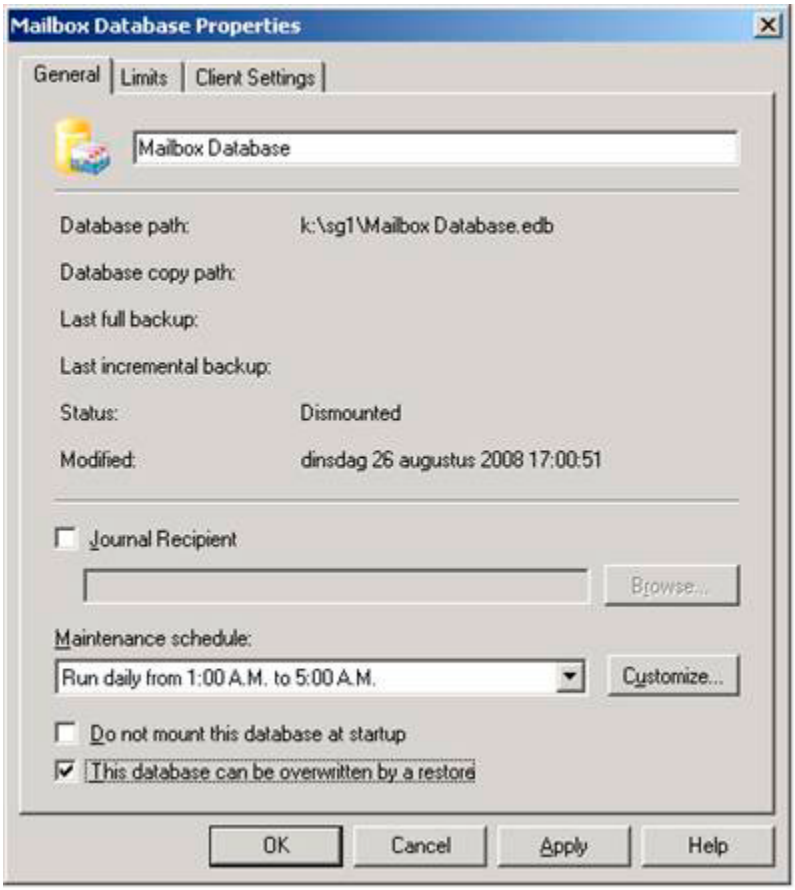

Figure 2. This database can be overwritten by a restore.

In NTBackup select the "Restore and Manage Media" tab and select the appropriate full backup set. After clicking on the restore button you are presented another, important Window.

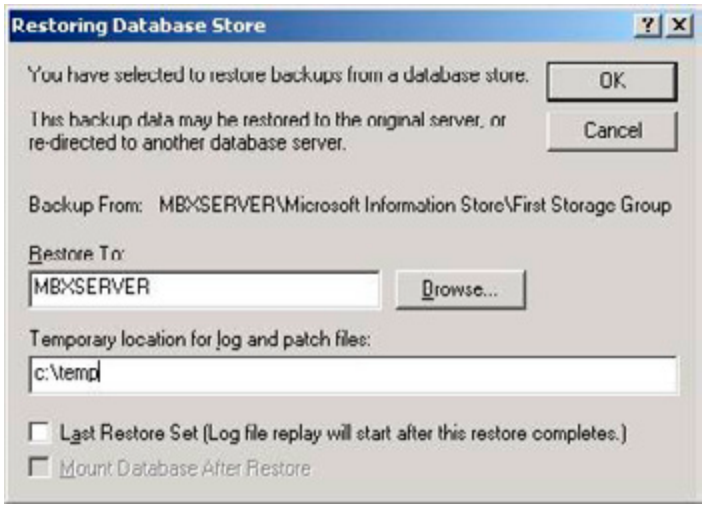

Figure 3. Is this really the last restore set? If so then checkmark the checkbox

If you want to restore an incremental backup DO NOT SELECT the "Last Restore Set." If you do, the database is being hard recovered immediately after the restore is finished and you do not have the option anymore to restore an incremental backup.

But also if this is the last restore set I'd like to not set the checkbox. This allows the possibility to check the database and the log files before a hard recovery is started. The log files that were written to the backup are restored in the Temporary location path c:\temp.

When the restore is finished you will see all the log files in the directory c:\temp\first storage group (or any other storage group that you've restored) and a file called restore.env. This file contains the information needed for hard recovery and can be read using the ESEUTIL tool:

```
C:\temp\First Storage Group>eseutil /cm

Extensible Storage Engine Utilities for Microsoft(R) Exchange Server
Version 08.01
Copyright (C) Microsoft Corporation. All Rights Reserved.

       Restore log file: C:\temp\First Storage Group

         Restore Path: C:\temp\First Storage Group
          Annotation: Microsoft Information Store
            Server: MBXSERVER
       Backup Instance: First Storage Group
       Target Instance:
 Restore Instance System Path:
  Restore Instance Log Path:

          Databases: 1 database(s)
            Database Name: Mailbox Database
                 GUID: 10364919-F177-46D7-B5BE46D7D1B7C03F
             Source Files: k:\sg1\Mailbox Database.edb
          Destination Files: k:\sg1\Mailbox Database.edb


       Log files range: E0000000022.log - E0000000034.logog
     Last Restore Time: <none>
       Recover Status: recoverNotStarted
       Recover Error: 0x0000000000
        Recover Time: Tue Sep 16 19:34:41 2008
```

```
Operation completed successfully in 0.0 seconds.

C:\temp\First Storage Group>
```

The process to fix the database with the log files that were restored from backup is called hard recovery. You can manually start the hard recovery using the ESEUTIL /CC command. This will replay all log files in the temporary directory into the database. The database itself is already in the correct location as you can see in the output above. When more log files exist in the normal production environment beyond the point of backup (and thus beyond the point of restore) they will be replayed into the production database as well. This will bring your database into a consistent state up to the most current point possible.

If you set the checkmark at "Last Restore Set" in Figure 3 this will happen automatically. The database will be hard recovered and all additional log files will be replayed as well. When finished the database can be mounted automatically as well.

Note

This process is the same for all backup applications from major vendors that support streaming backups with Exchange Server.

VSS or snapshot backups

With Exchange Server 2007 Microsoft is shifting its backup focus from the traditional online streaming backup to VSS or snapshot backups. Why do I still spend a serious amount of time on streaming backups? Because the underlying ideas are still very important and it gives an understanding what steps to perform in a VSS or snapshot backup.

Note

NTBackup cannot create VSS or snapshot backups from an Exchange Server database. It does however contain functionality to create a file level VSS backup. If you see a snapshot backup of your Exchange Server database in NTBackup it is very likely that the Exchange Server database is selected on the filesystem level and not on the Information Store level!

A snapshot is just a point-in-time and at this point-in-time an image is created. This image can be used to roll back to in case of a disaster. The Volume Shadow Copy Service in Windows Server 2003 and later provides an infrastructure to create these point-in-time images. These images are called Shadow Copies.

There are two kinds of Shadow Copies:

- **Clone** (Full Copy or Split Mirror) – a complete mirror is maintained until an application or administrator breaks the mirror. From this point on the original and clone are fully independent of each other. At this point it is effectively frozen in time.
- **Copy on Write** (Differential Copy) – a shadow copy is created that is a differential rather than a full copy of the original data. Using the Copy on Write a shadow copy of the original data is made before the original data is overwritten. Effectively the backup copy consists of the data in the shadow copy combined with the data on the original location. Both need to be available to reconstruct the original data.

The Volume Shadow Copy Infrastructure consists of the following components:

- **Requestor** – the software that invokes the VSS and creates, breaks or deletes the shadow copy. Typically the Requestor is the backup application.

- **Writer** – a software part that is provided by an application vendor, in our case this is provided with the Microsoft Exchange Server. A writer is responsible for providing a consistent point-in-time image by freezing or quiescing the Exchange Server at the point-in-time. Please note that an Exchange writer is provided for Exchange Server 2003 and higher.

- **Provider** – the interface to the point-in-time image. This can either be on a storage array (hardware provider) or in the Operating System (software provider). Windows Server 2003 provides a software provider with VSS functionality out-of-the-box.

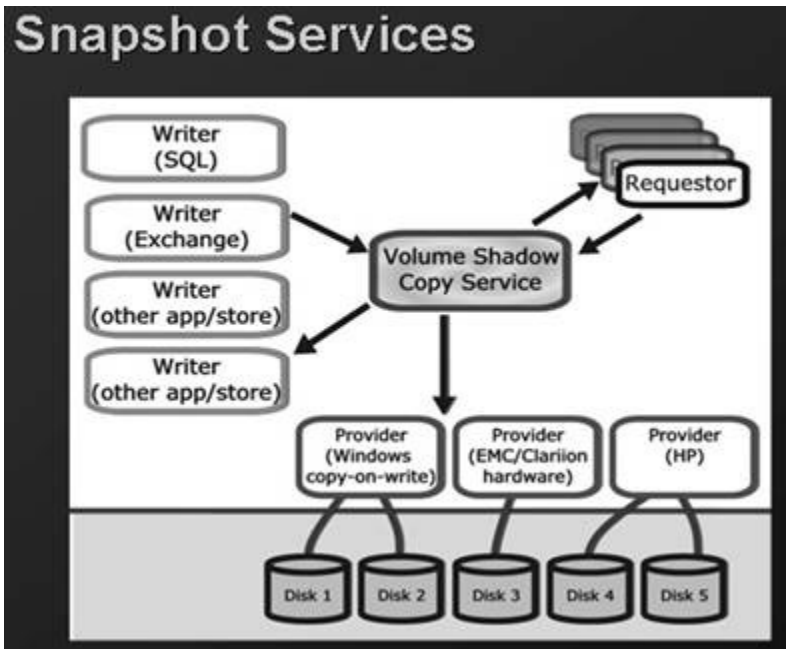

The following steps occur when a VSS backup is performed:

- The requestor, i.e. the backup application, sends a command to the Volume Shadow Copy Service to take a shadow copy of the Storage Groups.

- The VSS service sends a command to the Exchange writer to prepare for a snapshot backup.

- The VSS service sends a command to the appropriate storage provider to create a shadow copy of the Exchange Storage Group. This storage provider can be a hardware storage provider or the default Windows storage provider.

- The Exchange writer temporarily stops, or quiesces the Storage Group and puts them in read only mode and all data is flushed to the database. Also a log file roll-over is performed to make sure that all data will be in the backup set. This will hold a couple of seconds for the snapshot to be created (in the next step). All write I/O's will be queued.

- The shadow copy is now created.

- The VSS service releases the Exchange server to resume ordinary operations and all queued write I/O's are completed.

- The VSS service queries the Exchange writer to confirm that the write I/O's were successfully held during the shadow copy creation. If the writes were not successfully held it could mean a potentially inconsistent shadow copy, the shadow copy is deleted and the requestor is notified. The requestor can retry the shadow copy process or fail the operation.

- If successful, the requestor verifies the integrity of the backup set (the clone copy). If the clone copy integrity is good the requestor informs the Exchange Server that the backup was successful and that the log files can be purged.

Note

It is the responsibility of the backup application to perform a consistency check of the shadow copy. The Exchange writer does not perform this check. This is also the reason why you have to manually copy the ESE related files like ESE.DLL to the backup server.

Steps 1 through 6 usually take about 10 seconds, this is the time needed to create the actual snapshot. This is not the time to create a backup though. A backup application still has to create a backup to another disk or to tape, which still can take hours to complete depending on the size of the databases.

When the backup application has finished the header information of the Exchange database is updated as well. Using ESEUTIL /MH as explained in my previous article you can check the status of the database.

```
K:\sg1>eseutil /mh "mailbox database.edb"

Extensible Storage Engine Utilities for Microsoft(R) Exchange Server
Version 08.01
Copyright (C) Microsoft Corporation. All Rights Reserved.


Previous Full Backup:
    Log Gen: 13-14 (0xd-0xe) - OSSnapshot
      Mark: (0xF,8,16)6)6)
      Mark: 10/02/2008 14:27:30

Previous Incremental Backup:
    Log Gen: 13-16 (0xd-0x10) - OSSnapshot
      Mark: (0x11,8,16)
      Mark: 10/02/2008 14:34:25

Operation completed successfully in 0.171 seconds.

K:\sg1>
```

Note

This screenoutput has been edited for readability.

Microsoft does not currently offer a GUI-based VSS requestor for Exchange and using command-line tools puts you in a difficult support situation in regards to creating restores. Microsoft is working on a solution which can be read on the Microsoft Product Group teamblog:

For testing purposes Microsoft offers the Exchange Server 2007 Software Development Kit, the accompanying documentation can be found on the Microsoft website: the Microsoft website.

Microsoft also provides a VSS Software Development Kit. This SDK contains a very small test tool that is able to create VSS backups of the Exchange Storage Group. This command line tool is called BETest and can be used for test, development, troubleshooting or demonstrating VSS and Exchange 2007. More information regarding BETest can be found here. The Exchange product team released a blog on troubleshooting VSS issues which can be found here.

Using the BETest tool you can make very basic VSS backups. The backup is written to disk (on a location you can enter on the command line) and the log files are purged. Since the responsibility of performing a consistency check is at the backup application this check is not performed when using BETest.

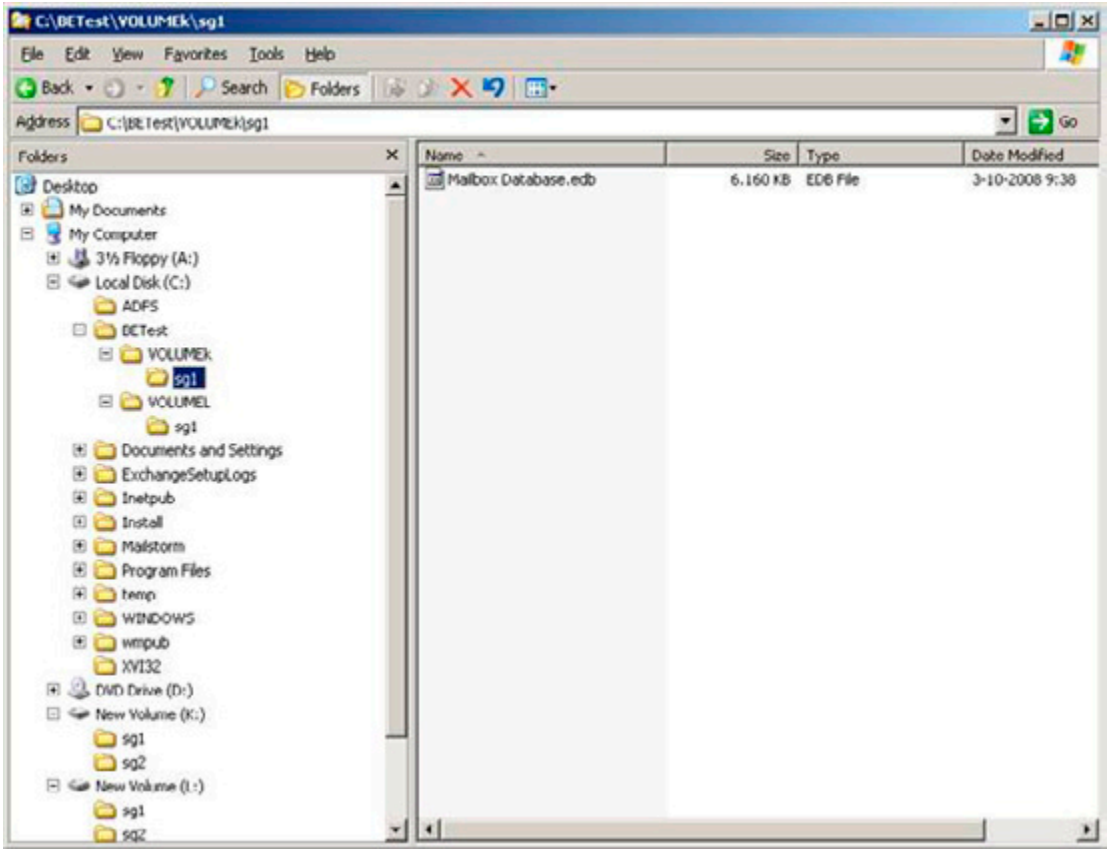

Figure 4. The backup on disk after using the BETest tool. The original database location was on disk k:\, the log files were on disk L:

When the database contains a faulty page you will never notice during the backup. Even the backup of the database on the disk will contain the faulty page. You have to manually perform a consistency check on the backup of the database using the ESEUTIL tool with the /K option:

```
C:\BETest\VOLUMEk\sg1>eseutil /k "mailbox database.edb"

Extensible Storage Engine Utilities for Microsoft(R) Exchange Server
Version 08.01
Copyright (C) Microsoft Corporation. All Rights Reserved.

Initiating CHECKSUM mode...
    Database: mailbox database.edb
 Temp. Database: TEMPCHKSUM2496.EDB

File: mailbox database.edb

           Checksum Status (% complete)

     0  10  20  30  40  50  60  70  80  90 100
     |----|----|----|----|----|----|----|----|----|----|
     WARNING: page 455 checksum verification failed but the corruption ( bit 35816 ) can be
corrected
..........................................................


770 pages seen
```

```
0 bad checksums
1 correctable checksums
297 uninitialized pages
0 wrong page numbers
0x5634 highest dbtime (pgno 0x1c8)

97 reads performed
6 MB read
1 seconds taken
6 MB/second
7866 milliseconds used
81 milliseconds per read
121 milliseconds for the slowest read
30 milliseconds for the fastest read

Operation completed successfully in 0.271 seconds.

C:\BETest\VOLUMEk\sg1>
```

The Event log shows all VSS steps performed by the BETest tool, but since no checking was performed nothing is logged in the Event log about a possible corruption!

Since this is a correctable error you can still access the mailbox and access the message that contains this faulty page. This is automatically corrected by Exchange Server. When this happens a message is logged in the Event log with Event ID 399:

```
"MSExchangeIS (3680) First Storage Group: The database page read from the file "k:\sg1\Mailbox
Database.edb" at offset 3735552 (0x0000000000390000) (database page 455 (0x1C7)) for 8192
(0x00002000) bytes failed verification. Bit 35816 was corrupted and has been corrected."
```

When using an Exchange Server 2007 Continuous Cluster Replication (CCR) solution a VSS backup solution is even more interesting. When the backup application is CCR aware it is possible to create a VSS backup against the passive copy of the CCR cluster. End users will never even notice a backup is created because all actions are performed against the passive copy. Any performance impact will only be noticed on the passive copy, where no users reside.

VSS Restore

A Volume Shadow Copy backup is created on a storage group level. Exchange Server 2003 is very rigid on restoring, it can restore to the same location, to the same server and only on a Storage Group level. Exchange Server 2007 is much more flexible, it can restore on the same location, but also to other servers or to a Recovery Storage Group for example.

Restoring a backup is a straight forward process and basically the same as an online restore. The database is restored, the log files in the backup set are also restored and hard recovery takes place to bring the database in a consistent state. If needed additional log files can be replayed as well to bring the database up-to-date to the last moment possible.

Replication and backups

One might ask if the replication technology is a good alternative for creating backups? The answer is simple and short: NO.

Since there is a copy of the database in an Exchange 2007 Cluster Continuous Replication solution, there is some more time available in a disaster recovery scenario. When a database crashes there's a copy of the database available to resume operations. But deleting messages or deleting mailboxes is a legitimate action from an Exchange point of view, and these actions will be replicated to the copy of the database as well. Also the offsite storage that's possible with backups (to tape) is a very important factor in creating backups.

Third-party application vendors

Microsoft does offer a VSS backup solution for Exchange Server via the System Center Data Protection Manager (DPM) 2007, but Microsoft does not offer a VSS backup solution for Exchange Server out-of-the-box like for example NTBackup. If you don't want to use DPM 2007 but you want to use a VSS backup solution for Exchange Server you have rely on a 3rd-party solution. Microsoft has several partners working with the Exchange team in Redmond to build really great products, each with its own feature set. You can find a complete list of backup partners here.

Conclusion

Creating online backups has major advantages over an offline backup solution since the backup application does all the logic for you. It also checks the database for any inconsistencies and if checksum errors are found they are automatically corrected before the data is sent to the backup set.

Recovery from an online backup is much easier than recovery from an offline backup. The database and the log files are automatically recovered from the backup set using the so called hard recovery. Additional log files that were created after the last backup set was created are automatically replayed, bringing the database up-to-date to the last moment.

A streaming backup is a default Windows 2003 solution, but Microsoft is shifting its focus from the streaming backup to the VSS (snapshot) backup due to the dramatic performance increase of VSS backups. Microsoft offers the System Center Data Protection Manager which can make backups, but it works very different from the standard VSS Backup solution that third-party vendors offer.

In these three articles (Exchange Database Technologies, Exchange Server Log File Replay, and this one) I have explained some basics about Microsoft Exchange database technologies, replaying log files, recovery techniques and backup and restore solutions.

It's your turn now to start working with this information in a lab environment and start thinking about a disaster recovery solution for your Exchange Server environment. Document all the steps needed when a disaster strikes and perform regular fire drills on disaster recovery. Only this will help you recover more quickly when something bad happens in your Exchange Server environment.

Optimizing Exchange Server 2007

24 November 2008

by Brien Posey

Brien Posey ponders an "off the cuff" remark that Exchange 2007 runs so well with a default configuration that you don't even have to worry about optimizing Exchange any more. He decides that there is actually plenty that can be done to help Exchange to perform even better.

A couple of weeks ago, I was in Las Vegas presenting several Exchange Server sessions at one of the many IT conferences that they have out there. After one of my sessions, I was having a conversation with one of the other conference attendees. During the conversation, he made an off the cuff remark that Exchange 2007 runs so well with a default configuration, that you don't even have to worry about optimizing Exchange any more.

Maybe it was because I was because I was still jet lagged, or because my mind was still on the session that I had just presented, but the remark didn't really register with me at the time. Later on though, I started thinking about his comment, and while I don't completely agree with it, I can kind of see where he was coming from.

I started working with Exchange at around the time when people were first moving from Exchange 4.0 to version 5.0. At the time, I was working for a large, enterprise class organization with roughly 25,000 mailboxes. Although those mailboxes were spread across several Exchange Servers, the server's performance left a lot to be desired. I recall spending a lot of time trying to figure out things that I could do to make the servers perform better.

Although there is no denying how poorly our Exchange Servers performed, I think that the server hardware had more to do with the problem than Exchange itself. For example, one night I was having some database problems with one of my Exchange Servers. Before I attempted any sort of database repair, I decided to make a backup copy of the database. In the interest of saving time, I copied the database from its usual location to another volume on the same server.

All of the server's volumes were using "high speed" RAID arrays, but it still took about three and a half hours to backup a 2 GB information store. My point is that Exchange has improved a lot since the days of Exchange 5.0, but the server hardware has improved even more dramatically. Today multi-core CPUs, and terabyte hard drives are the norm, but were more or less unheard of back in the day.

When the guy at the conference commented that you don't even have to worry about optimizing Exchange 2007, I suspect that perhaps he had dealt with legacy versions of Exchange on slow servers in the past. In contrast, Exchange 2007 uses a 64-bit architecture, which means that it can take full advantage of the CPU's full capabilities and that it is no longer bound by the 4 GB memory limitation imposed by 32-bit operating systems.

Although Exchange 2007 does perform better than its predecessors, I would not go so far as to say that optimization is no longer necessary. Think about it this way… If you've got an enterprise grade Exchange Server, but you've only got 20 mailboxes in your entire organization, then that server is going to deliver blazing performance. If you start adding mailboxes though, you are eventually going to get to the point at which the server's performance is going to start to suffer.

This illustrates two points. First, the Exchange 2007 experience is only as good as what the underlying hardware is capable of producing. Second, even if your server is running well, future growth may require you to optimize the server in an effort to maintain the same level of performance that you are enjoying now. Fortunately, there are some things that you can do to keep Exchange running smoothly. I will spend the remainder of this article discussing some of these techniques.

My Approach to Exchange 2007 Optimization

To the best of my knowledge, Microsoft has not published an Exchange 2007 optimization guide. They do offer some lists of post installation tasks, but most of the tasks on the list are related more to the server's configuration than to its performance. Although there isn't an "optimization guide" so to speak, the Exchange 2007 Deployment Guide lists a number of optimization techniques, and for all practical purposes serves as the optimization guide. Since I can't cover all of the techniques listed in the guide within a limited amount of space, I want to talk about some optimization techniques that have worked for me.

As I have already explained, the underlying hardware makes a big difference as to how well Exchange is going to perform. For the purposes of this article, I am going to assume that you have already got the appropriate hardware to meet your needs. If you have insufficient hardware capabilities for the workload that the server is carrying, then these techniques may or may not help you.

The Microsoft Exchange Best Practices Analyzer

If you were to ask me what the single most important (non hardware related) thing you could do to optimize Exchange Server, I would probably tell you to run the Microsoft Exchange Best Practices Analyzer (ExBPA). Although the ExBPA, shown in Figure A, was originally designed as a configuration analysis tool, and has been largely used by the IT community as a security tool, what it really does is scan your Exchange Server organization and make sure that it is configured according to Microsoft's recommended best practices. Many of Microsoft's best practices for Exchange Server are related to security, but there are plenty of performance and reliability related recommendations as well.

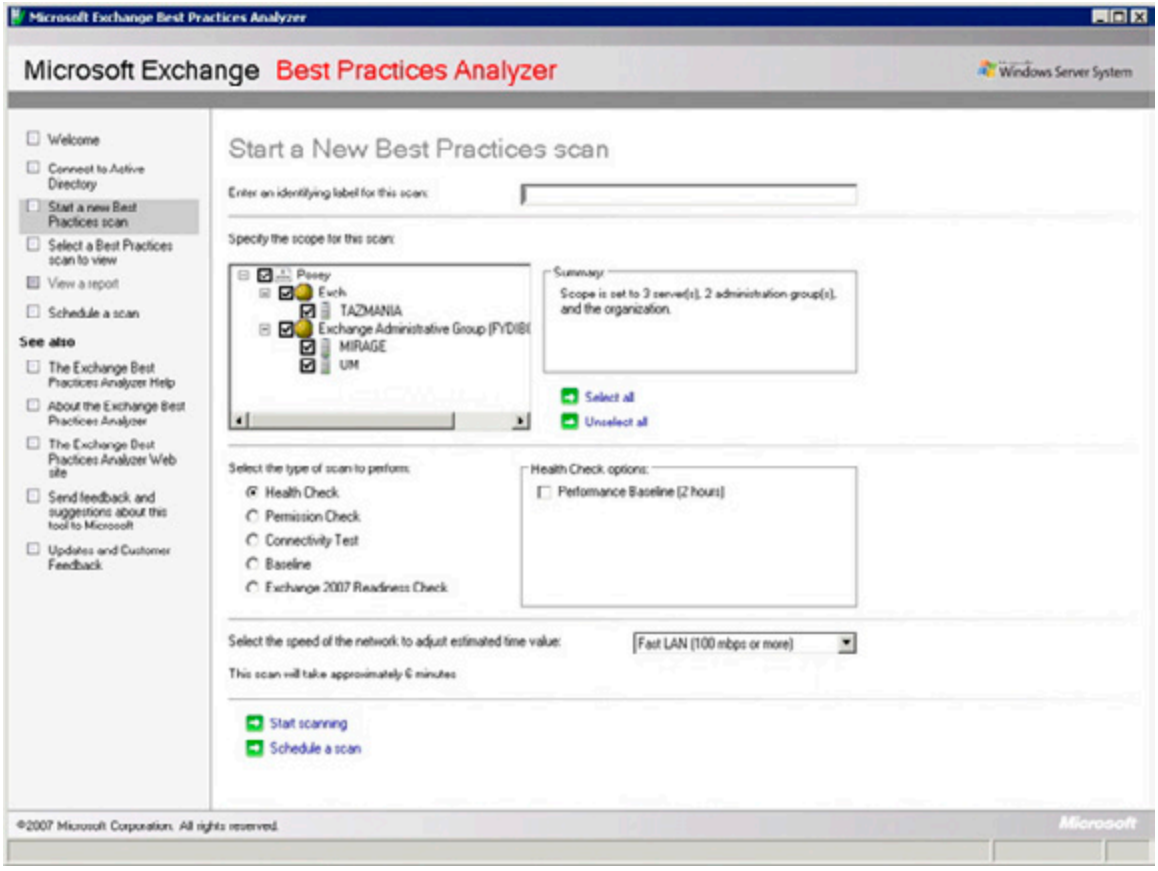

Figure A.

The Microsoft Exchange Best Practice Analyzer makes sure that your Exchange organization is configured according to Microsoft's best practices.

The ExBPA started life as a downloadable utility for Exchange 2003. Microsoft recommends this tool so strongly though, that they actually included it in the Exchange Management Console for Exchange Server 2007, at the very top of the list of tools in the Toolbox.

One thing that you need to keep in mind is that Microsoft is a dynamic organization. They are constantly doing research on the ways that their products are being used, and sometimes their recommended best practices end up changing as a result of that research. This means that optimizing Exchange 2007 is not a set it and forget it operation. You need to periodically run the ExBPA to see if any of Microsoft's recommendations for your organization have changed. Keep in mind though that if you are using Microsoft's System Center Operations Manager or System Center Essentials, they will automatically run ExBPA on a daily basis.

Disk I/O

If you are an Exchange veteran, then you might have been a little surprised when I listed running the ExBPA as the single most important step in optimizing an Exchange Server. In every previous version of Exchange, the most important thing that you could do from a performance standpoint was to optimize disk I/O.

There are a couple of reasons why I didn't list disk optimizing disk I/O as my top priority. First, if an Exchange Server has a serious problem in which the configuration of the disk subsystem is affecting performance, then, the ExBPA should detect and report the issue when you run a performance baseline scan. More importantly though, you can always count on your users to tell you when things are running slowly.

Another reason why disk I/O isn't my top priority is because Exchange Server 2007 uses a 64-bit architecture, which frees it from the 4 GB address space limitation imposed by 32-bit operating systems. Microsoft has used this larger memory model to design Exchange 2007 in a way that drives down read I/O requirements. This improves performance for database reads, although the performance of database writes have not really improved.

Of course that doesn't mean that the old rules for optimizing a disk subsystem no longer apply. It is still important to place the transaction logs and the databases onto separate volumes on high speed raid arrays for performance and fault tolerant reasons. Keep in mind that you can get away with using two separate volumes on one physical disk array, but doing so is not an ideal arrangement unless you divide your bank of disks into separate arrays so that each volume resides on separate physical disks.

If the same physical disks are used by multiple volumes, then fault tolerance becomes an issue unless the array is fault tolerant. That's because if there is a failure on the array, and both volumes share the same physical disks, then both volumes will be affected by the failure.

Even if the volume is fault tolerant, spanning multiple volumes across the same physical disks is sometimes bad for performance. In this type of situation, the volumes on the array are competing for disk I/O. If you were to separate the volumes onto separate arrays or onto separate parts of an array though, you can eliminate competition between volumes for disk I/O. Each volume has full reign of the disk resources that have been allocated to it.

Storage Groups

As was the case in previous versions of Exchange, Microsoft has designed Exchange 2007 so that multiple databases can exist within a single storage group, or you can dedicate a separate storage group to each individual storage group.

Microsoft recommends that you limit each storage group to hosting a single database. This does a couple of things for you. First, it allows you to use continuous replication, which I will talk about later on. Another thing that it does is that it helps to keep the volume containing the transaction logs for the storage group from becoming overwhelmed by an excessive number of transactions coming from multiple databases, assuming that you use a separate volume for the transaction logs from each storage group.

Using a dedicated storage group for each database means that the transaction log files for the storage group only contain transaction log entries for one database. This makes database level disaster recovery much easier, because you do not have to deal with transaction log data from other databases.

Mailbox Distribution

Since Exchange Server can accommodate multiple Exchange Server databases, it means that you have the option of either placing all of your mailboxes into a single database, or of distributing the mailboxes among multiple databases. Deciding which method to use is actually one of the more tricky points of optimizing Exchange Server's performance. There are really two different schools of thought on the subject.

Some people believe that as long as there aren't an excessive number of mailboxes, that it is best to place all of the mailboxes within a single store. The primary advantages of doing so are ease of management and simplified disaster recovery.

There is also however an argument to be made for distributing Exchange mailboxes among multiple stores (in multiple storage groups). Probably the best argument for doing so is that if a store level failure occurs then not all of your mailboxes will be affected. Only a subset of the user base will be affected by the failure. Furthermore, because the store is smaller than it would be if it contained every mailbox in the entire organization, the recovery process tends to be faster.

Another distinct advantage to using multiple mailbox stores is that assuming that each store is placed on separate disksfrom the other stores, and the transaction logs are also distributed onto separate disks, the I/O requirements are greatly decreased, because only a fraction of the total I/O is being directed at any one volume. Of course the flip side to this is that this type of configuration costs more to implement because of the additional hardware requirements.

As you can see, there are compelling arguments for both approaches, so which one should you use? It really just depends on how much data your server is hosting. Microsoft recommends that you cap the mailbox store size based on the types of backups that you are using. For those running streaming backups, Microsoft recommends limiting the size of a store to no more than 50 GB. For organizations using a VSS backups without a continuous replication solution in place, they suggest limiting your database size to 100 GB. This recommendation increases to 200 GB if you have a continuous replication solution in place

There are a couple of things to keep in mind about these numbers though. First, the limits above address manageability, not performance. Second, the database's size may impact performance, but it ultimately comes down to what your hardware can handle. Assuming that you have sufficient hardware, Exchange 2007 will allow you to create databases with a maximum size of 16 TB.

If you do decide to split an information store into multiple databases, there are some hardware requirements that you will need to consider beyond just the disk configuration. For starters, each database is going to consume some amount of CPU time, although the actual amount of additional CPU overhead varies considerably.

You are also going to have to increase the amount of memory in your server as you increase the number of storage groups. Microsoft's guidelines state that a mailbox server should have at least 2 GB of memory, plus a certain amount of memory for each mailbox user. The table below illustrates the per user memory requirements:

User Type	Definition	Amount of Additional Memory Required
Light	5 messages sent / 20 messages received per day	2 MB
Medium	10 messages sent / 40 messages received per day	3.5 MB
Heavy	Anything above average	5 MB

The 2 GB base memory, and the per user mailbox memory requirements make a couple of assumptions. First, it is assumed that the server is functioning solely as a mailbox server. If other server roles are installed, then the base memory requirement is increased to 3 GB, and the per user mailbox requirements remain the same.

These recommendations assume that the server has no more than four storage groups though. Each storage group consumes some memory, so Microsoft requires additional memory as the number of storage groups increase.

When Microsoft released SP1 for Exchange 2007, they greatly decreased the amount of memory required for larger numbers of storage groups. For example, a server with 50 storage groups requires 26 GB of memory in the RTM release of Exchange 2007, but when SP1 is installed, the requirement drops to 15 GB. The table below illustrates how the base memory requirement changes as the number of storage groups increases.

Number of Storage Groups	Exchange 2007 RTM Base Memory Requirements	Exchange 2007 Service Pack 1 Base Memory Requirements
1–4	2 GB	2 GB
5–8	4 GB	4 GB
9–12	6 GB	5 GB
13–16	8 GB	6 GB
17–20	10 GB	7 GB
21–24	12 GB	8 GB
25–28	14 GB	9 GB
29–32	16 GB	10 GB
33–36	18 GB	11 GB
37–40	20 GB	12 GB
44–44	22 GB	13 GB
45–48	24 GB	14 GB
49–50	26 GB	15 GB

One Last Disk Related Consideration

Exchange Server 2007 Enterprise Edition allows you to use up to 50 storage groups. I have only worked with a deployment of this size on one occasion, but made an interesting observation that I had never seen specifically pointed out in any of the documentation.

Normally, server volumes are referenced by drive letters. If you have a server with 50 stores, all on separate volumes, then you more than exhaust the available drive letters. As such, you will have to address the volumes as mount points rather than drive letters.

Backups

I don't know about you, but when I think about optimizing a server's performance, backups are not usually the first thing that comes to mind. Even so, my experience has been that you can get a big performance boost just by changing the way that nightly backups are made. Let me explain.

Most organizations still seem to be performing traditional backups, in which the Exchange Server's data is backed up to tape late at night. There are a couple of problems with this though. For starters, the backup process itself places a load on the Exchange Server, which often translates into a decrease in the server's performance until the backup is complete. This probably isn't a big deal if you work in an organization that has a nine to five work schedule, but if your organization is a 24-hour-a-day operation then a performance hit is less than desirable.

Another aspect of the nightly backup that many administrators don't consider is that a nightly backup almost always coincides with the nightly automated maintenance tasks. By default, each night from midnight to 4:00 AM, Exchange performs an automated maintenance cycle.

This automated maintenance cycle performs several different maintenance tasks, including an online database defragmentation. These maintenance tasks tend to be I/Ontensive, and the effect is compounded if a backup is running against a database at the same time that the maintenance tasks are running.

There are a few ways that you can minimize the impact of the maintenance cycle and the backup process. One recommendation that I would make would be to schedule the maintenance cycle so that it does not occur at an inopportune moment.

The maintenance cycle occurs at the database level. If you've got multiple databases, then by default the maintenance cycle will run on each database at the same time. Depending on how many databases you've got, you may be able to schedule the maintenance cycle so that it is only running against one database at a time. Likewise, you may also be able to work out a schedule that prevents the maintenance cycle and the backup from running against the same database at the same time.

If you want to see the current maintenance schedule for a database, open the Exchange Management Console, and navigate through the console tree to Server Configuration | Mailbox. Next, select your Exchange mailbox server in the details pane, followed by the store that you want to examine. Right click on the store and then select the Properties command from the resulting shortcut menu. You can view or modify the maintenance schedule from the resulting properties sheet's General tab.

Another way that you can mitigate the overhead caused by the backup process is to take advantage of Cluster Continuous Replication (CCR). Although CCR is no substitute for a true backup, CCR does use a process called log shipping to create a duplicate of a database on another mailbox server. It is possible to run your backups against this secondary cluster node rather than running it against your primary Exchange Server. That way, the primary Exchange Server is not impacted by the backup.

If you do decide to use CCR, then you will have to keep in mind that you are only allowed to have one database in each storage group. Otherwise, the option to use CCR is disabled.

The Windows Operating System

Another aspect of the optimization process that is often overlooked is the Windows operating system. Exchange rides on top of Windows, so if Windows performs poorly, then Exchange will too.

One of the best things that you can do to help Windows to perform better is to place the pagefile onto a dedicated hard drive (or better yet, a dedicated array). You should also make sure that the pagefile is sized correctly. Microsoft normally recommends that the Windows pagefile should be 1.5 times the size of the machine's physical memory. There is a different recommendation for Exchange 2007 though. Microsoft recommends that you set the pagefile to equal the size of your machine's RAM, plus 10 MB. If you use a larger pagefile, then the pagefile will eventually become fragmented, leading to poor performance.

Finally, make sure that your server is running all of the latest patches and drivers. You should also take the time to disable any services that you don't need. Every running service consumes a small amount of system resources, so disabling unneeded services can help you to reclaim these resources.

Conclusion

In this article, I have explained that while Exchange Server has improved over the years, there are still a lot of things that you can do to help Exchange to perform even better. This is especially important in larger organizations in which the server's finite resources are being shared by many different users.

Exchange: Recovery Storage Groups

06 January 2009

by Jaap Wesselius

> ***It can happen at any time: You get a request, as Admin, from your company, to provide the contents of somebody's mailbox from a backup set as part of an investigation. The Recovery Storage Group is usually the easiest way to do this. It may either mean using the Exchange Management Console or the Exchange Management Shell. Jaap explains all.***

Very recently I was at a customer and the Exchange administrators had received a request from the legal department to retrieve somebody's mailbox from several backup sets and create .PST files of the inbox. The .PST files had to be handed over to the legal department. Of course the particular user was not aware of these actions.

The Exchange administrator retrieved the Exchange database files from backup (130 GB!) but had absolutely no idea how to retrieve the mailbox data out of the files. The Recovery Storage Group in Exchange Server 2003 and Exchange Server 2007 can be very useful in a situation like this.

Backup and Restore

This particular customer made a full backup of their Exchange Server 2003 database every weekend and every night an incremental backup was created. This means that during a normal night the Exchange Server database is not backed up, but the Exchange Server log files are backed up and purged from the disk. Please refer to my article "Online Exchange Backups" on Simple-Talk.com for more information on these backup techniques.

If data from a mailbox out of a backup that's 14 days old needs to be retrieved you have to get that backup set, but you cannot just simply restore the database to its original location. In order to restore the database to its original location the database needs to be dismounted and the 14 days old backup has to be restored. This not only results in an outage but will also result in a loss of data if you're not careful, so this is not an acceptable option. Also you have to restore the full backup made during the weekend, and possibly one or more incremental backups, depending on the date stamp you have to restore.

A Recovery Storage Group is like any other Storage Group, except for the fact that databases mounted in a Recovery Storage Group do not contain live mailboxes but only disconnected mailboxes. A Recovery Storage Group is a Storage Group that can be used for recovery purposes; you can restore a database from a backup set in a Recovery Storage Group and recover data from it. A Recovery Storage Group and the database(s) it contains are invisible to the end-users and Outlook clients, only the administrator can access the Recovery Storage Group and the database(s) it contains.

Suppose there is a user in our Exchange Server 2003 environment named Joe Sixpack. Joe is under suspicion of the legal department and the legal department has requested a copy of his mailbox to check whether Joe has sent out some confidential information to the Internet. Joe's mailbox is located on an Exchange server named 2003SRVR in the default mailbox store in the First Storage Group. This Storage Group contains three databases in total. There's also a second Storage Group that hosts only one Mailbox Store.

To retrieve a copy of Joe's mailbox the IT department wants to use the Recovery Storage Group option in Exchange Server 2003.

Follow these steps to create a Recovery Storage Group in Exchange Server 2003:

- In Exchange System Manager, select and then right-click the Exchange Server 2003 object, select New and then select "Recovery Storage Group."
- Enter the location where the transaction log files and the checkpoint file should be located. By default it is in the "C:\Program Files\Exchsrvr\Recovery Storage Group" directory, but it can be any location you want. Please keep in mind that there should be sufficient storage available to host the complete database and all log files. If you want to perform recovery steps using ESEUTIL you have to take an additional storage requirement into account.
- An empty directory is created where the database and the log files will be stored.

When the Recovery Storage Group is created the database that needs to be restored must be selected.

- In the Exchange System Manager, right click the Recovery Storage Group and select "Add Database to Recover."
- All mailbox databases that are known in Active Directory are displayed, so Joe's Mailbox Store should be selected, this is in this example "Mailbox Store (2003SRVR)."
- After clicking OK the Mailbox Store properties are shown, leave these at their default values and click OK.

DO NOT MOUNT THE DATABASE

The database is now created in the Recovery Storage Group in Active Directory. The next step is to restore the database from the backup set.

No additional steps are needed here. When a Recovery Storage Group is created the behaviour of an online restore is changed. The database is not restored to its original location, but it is automatically redirected to the Recovery Storage Group location that's created earlier. Just restore the "Mailbox Store (2003SRVR)," it will be automatically placed in the right directory.

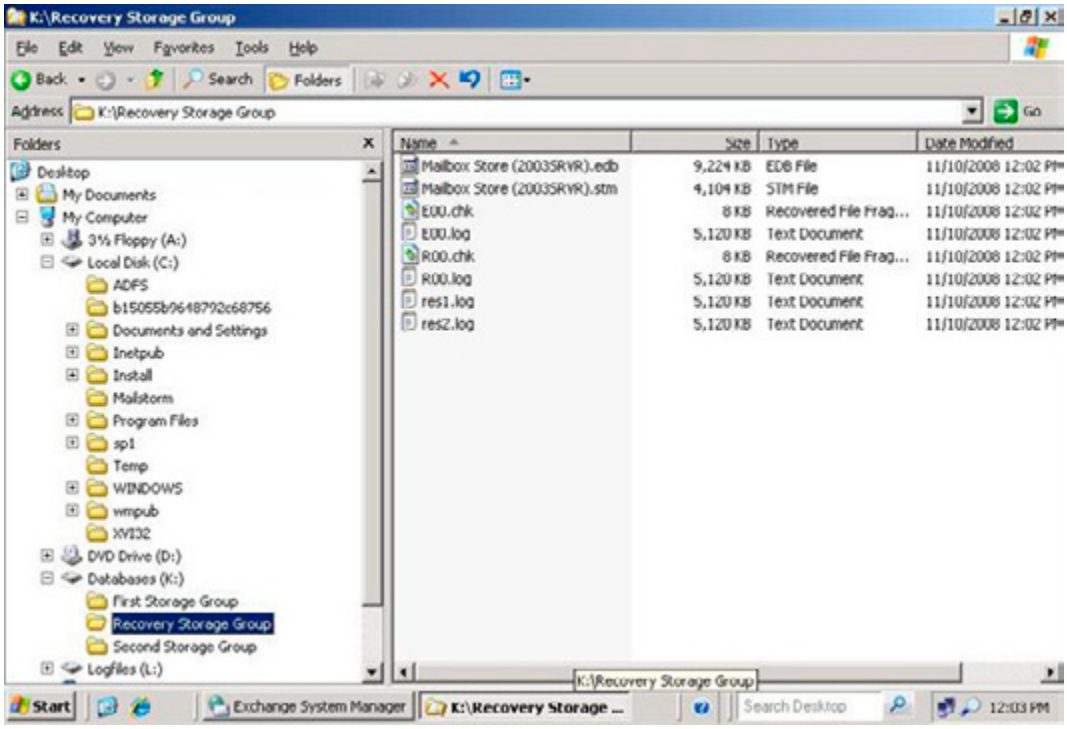

Figure 1. The Recovery Storage Group directory and its contents right after a restore operation.

In a normal Storage Group in Exchange Server 2003 the log files have a prefix like E00, E01, E02 or E03. In a Recovery Storage Group the prefix is R00. This way it is always clear when a particular directory contains data from a Recovery Storage Group.

After the Exchange database is restored from a backup set it can be mounted. Just right-click this database in the Recovery Storage Group in the Exchange System Manager and select "Mount Store."

Remember that the database in a Recovery Storage Group is actually an old version of the running database, so it also contains all mailboxes that are in the running database. In the Recovery Storage Group however these are disconnected mailboxes, i.e. the original user accounts in Active Directory are matched to the mailboxes in the running database.

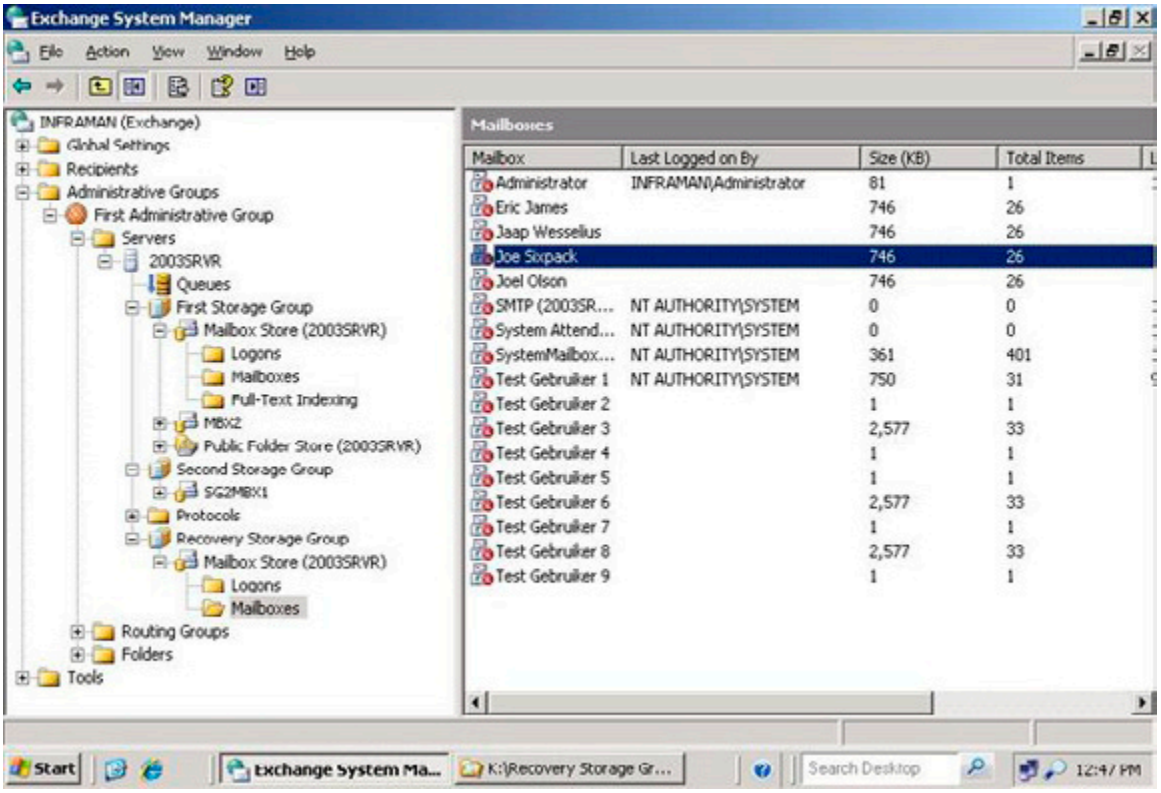

Figure 2. After mounting the database in the Recovery Storage Group contain disconnected mailboxes.

The last step is to create a .PST file from Joe Sixpack's mailbox. The Exchange System Manager contains some functionality to copy or merge this data into the running mailbox of the user, for creating a .PST file the free Microsoft tool EXMERGE have to be used. Exmerge can be downloaded from the Microsoft website: http://www.microsoft.com/downloads/details.aspx?familyid=429163EC-DCDF-47DC-96DA-1C12D67327D5&displaylang=en. More information about the usage of Exmerge can be found in Microsoft knowledgebase article 327304: http://support.microsoft.com/?kbid=327304

To retrieve the contents of Joe Sixpack's mailbox from the Recovery Storage Group follow these steps:

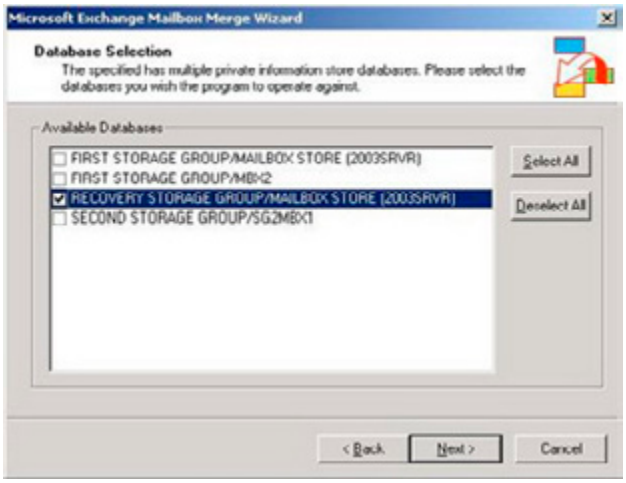

Figure 3. Select the database from the Recovery Storage Group.

- Start the Exmerge.exe utility from the local disk.
- In the Microsoft Exchange Mailbox Merge Wizard select "Extract or Import (Two Step Procedure)" and click Next.
- Select "Step 1: Extract data from an Exchange Server Mailbox" and click Next.
- In the Source Server window enter the servername and the domain controller name and click Next.
- In the Database Selection Windows select the Mailbox Store in the Recovery Storage Group. This is the "old" store where Joe's data need to be extracted from.
- In the next windows select Joe Sixpack's mailbox.

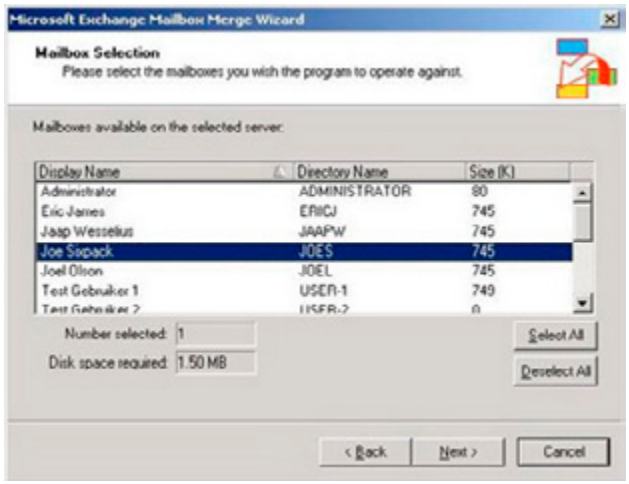

Figure 4. Select Joe Sixpack's mailbox to retrieve the data from.

- In the Next Windows select the locale, the default setting is "English (US)." The last mailbox login locale can also be selected by clicking the Checkbox.
- Select the Folder (directory) where the .PST file will be stored. Any directory can be used.

- Save the settings (default selection) and watch the export process going on.
- When the operation is completed successfully click Finish.

Joe Sixpack's mailbox that was in the backup set is now saved in the .PST file on local disk. This .PST file can be handed over to the legal department for further investigation while Joe Sixpack never noticed anything about the whole operation.

Recovery Storage Group in Exchange Server 2007

The Recovery Storage Group functionality is also implemented in Exchange Server 2007, but it is not available through the default Exchange Management Console interface. To perform all operations with a Recovery Storage Group the Exchange Management Shell needs to be used. You can use some Recovery Storage Group functionality using the toolbox in the Exchange Management Console; I will explain this later in this article.

Open the Exchange Management Shell and enter the New-StorageGroup command. Make sure that you use the –Recovery option, otherwise you'll create a regular Storage Group:

```
[PS] C:\>New-StorageGroup -Server MBXSERVER -LogFolderPath l:\rsg
-SystemFolderPath l:\rsg -Name RSG -Recovery

Name            Server      Replicated    Recovery
----            ------      ----------    --------
RSG             MBXSERVER     None       True
```

Now create the Database in using the New-MailboxDatabase cmdlet. This database has the same name as the original mailbox database that contains Joe Sixpack's mailbox:

```
[PS] C:\>New-MailboxDatabase -MailboxDatabaseToRecover "Mailbox Database"
-StorageGroup MBXSERVER\RSG -EdbFilePath "k:\rsg\mailbox database.edb"

Name          Server     StorageGroup     Recovery
----          ------     ------------     --------
Mailbox Database   MBXSERVER    RSG         True
```

The newly created database in the Recovery Storage Group should have the ability to be overwritten by a restore, so set this property:

```
[PS] C:\>Set-MailboxDatabase -Identity "MBXSERVER\RSG\Mailbox Database"
-AllowFileRestore:$TRUE
```

Before the database can be mounted, the database should be restored from the backup set. When a Recovery Storage Group is created the restore is automatically redirected to the Recovery Storage Group by the Information Store, just like in Exchange Server 2003. Restore the database using your backup application.

When the database is restored into the Recovery Storage Group directory on the disk the database can be mounted:

```
[PS] C:\>Mount-Database -Identity "MBXSERVER\RSG\Mailbox Database"
```

Now the Recovery Storage Group is created, the mailbox database is created, restored from the backup set and mounted. It is still invisible to the end users and no one except the administrator can access the mailbox database in the Recovery Storage Group. Please remember that it is only accessible through the Management Shell, so colleagues using the regular Exchange Management Console also do not notice anything.

There is a thing you have to know, though. In the Exchange Management Console you can go to the Toolbox and enter the "Database Recovery Management." From there there's also the possibility to manage the Recovery Storage Group. The following tasks are available in the "Manage Recovery Storage Group" area:

- Merge or copy mailbox contents.
- Mount or dismount databases in the Recovery Storage Group.
- Remove the Recovery Storage Group.
- Set up "Database can be overwritten by restore" flag.
- Swap databases for "dial tone" scenario.

When the database is mounted the database statistics can be requested of the database in the Recovery Storage Group with the Get-MailboxStatistics cmdlet. These are disconnected mailboxes and are unknown to Active Directory. As far as Active Directory is concerned the actual mailboxes are located in the "mailbox database.edb" in the First Storage Group. Joe Sixpack's mailbox is clearly visible.

```
[PS] C:\>Get-MailboxStatistics -Database "RSG\Mailbox Database"

DisplayName        ItemCount  StorageLimitStatus     LastLogonTime
-----------        ---------  ------------------     -------------
Test User 2        145                10-11-2008 14:18:29
Test User 4        122                10-11-2008 14:18:29
Test User 3        129                10-11-2008 14:18:29
Eric James        64
Test User 9        85                 10-11-2008 14:18:30
Administrator       1                 13-11-2008 15:03:25
Joel Olson        40
Test User 7        99                 10-11-2008 14:18:30
Test User 6        105                10-11-2008 14:18:30
Test User 5        111                10-11-2008 14:18:29
Test User 8        93                 10-11-2008 14:18:30
Test User 10       72                 10-11-2008 14:18:29
Joe Sixpack        56                 13-11-2008 17:12:21
Microsoft System Attendant 0          13-11-2008 16:34:10

SystemMailbox{10364919-F1 402
77-46D7-BEB5-3FC0B7D155A9
}
Test User 1        163                 10-11-2008 14:18:29

[PS] C:\>
```

Exchange Server 2007 SP1 has a new cmdlet to export data to a .PST file and as such the successor of EXMERGE that was used in Exchange Server 2003 and earlier. There are a few caveats with the Export-Mailbox cmdlet:

- To use this cmdlet Outlook 2003 SP2 (or higher) has to be installed on the server where the cmdlet is run. This can be a management workstation (or a management server) with only the Exchange Management Tools installed. It should also be a 32-bits server since the 64-bits Exchange Management Tools cannot cooperate with the 32-bits Outlook client software.
- To Export-Mailbox cmdlet runs only against a normal Exchange Server 2007 SP1 database and not against databases running in the Recovery Storage Group.

To overcome the second limitation a special "recovery mailbox" can be created in an ordinary Storage Group. When this mailbox is created the contents of Joe Sixpack's mailbox can be restored into this recovery mailbox:

```
[PS] C:\ >restore-mailbox -Identity "Recovery Mailbox" -RSGDatabase "RSG\Mailbox Database"
-RSGMailbox "joe sixpack" -TargetFolder Recover
```

When the Restore Mailbox operation is running a progress bar is shown.

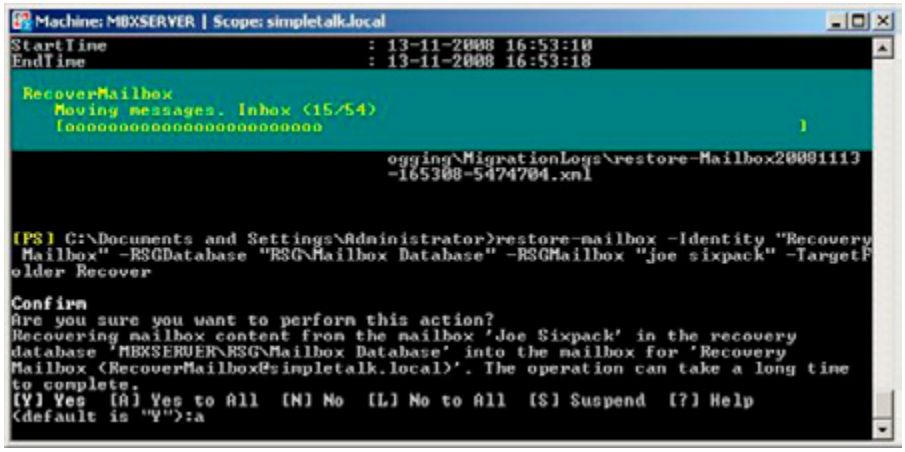

Figure 5. The progress bar shown during the Recover Maillbox operation.

When the Restore Mailbox operation is finished the contents of Joe Sixpack's mailbox is available in the Recovery Mailbox. When using Outlook to access this mailbox the contents is available in the \Recover folder in the mailbox. The \Recover folder in the Recovery Mailbox can be exported using Outlook, in the File menu select the "Import and Export option" and then select "Export to a file." Follow the wizard to complete the export to a .PST file.

When Outlook 2003 SP2 (or higher) is installed on the Management Server the Export Mailbox cmdlet can also be used:

```
Export-mailbox –identity "Recovery Mailbox" -PSTFolderPath c:\temp\joesix.pst –IncludeFolders "\
Recover"
```

This will export the contents of the \Recover folder, and therefore the recovered mailbox from Joe Sixpack to a .PST file in the c:\temp directory on the local disk.

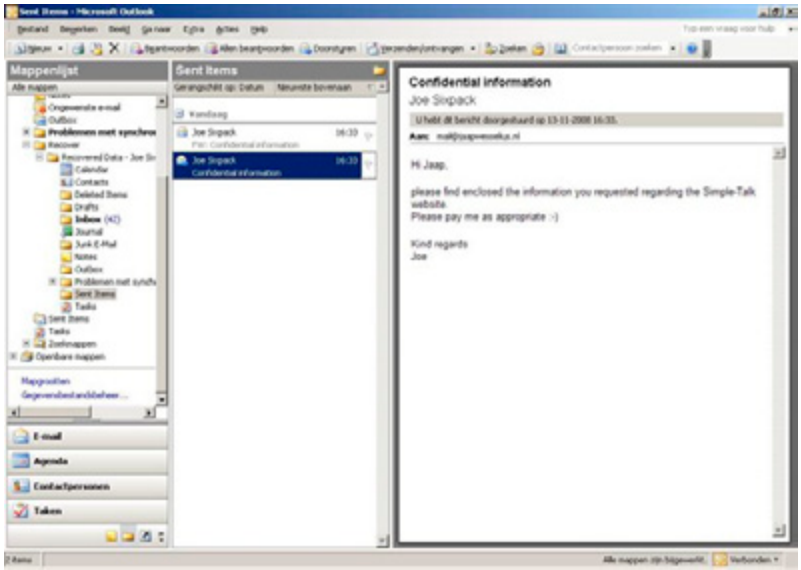

Figure 6. The Contents of the \Recovery folder and therefore of Joe Sixpack's mailbox. Joe Sixpack did, indeed, send out sensitive information.

Recover Mailbox Data

Another interesting feature of the Recovery Storage Group is to recover deleted mailbox data. Suppose a user named Joel accidentally deleted important items from his mailbox and he is unable to retrieve this using the "recover deleted items" option in Outlook.

To solve this problem please restore the database that holds Joel's mailbox data into the Recovery Storage Group. In Exchange Server 2003 in the Exchange System Manager go to the Recovery Storage Group and check Joel's mailbox. A disconnected mailbox should be visible. Right Click the mailbox and select "Exchange tasks."

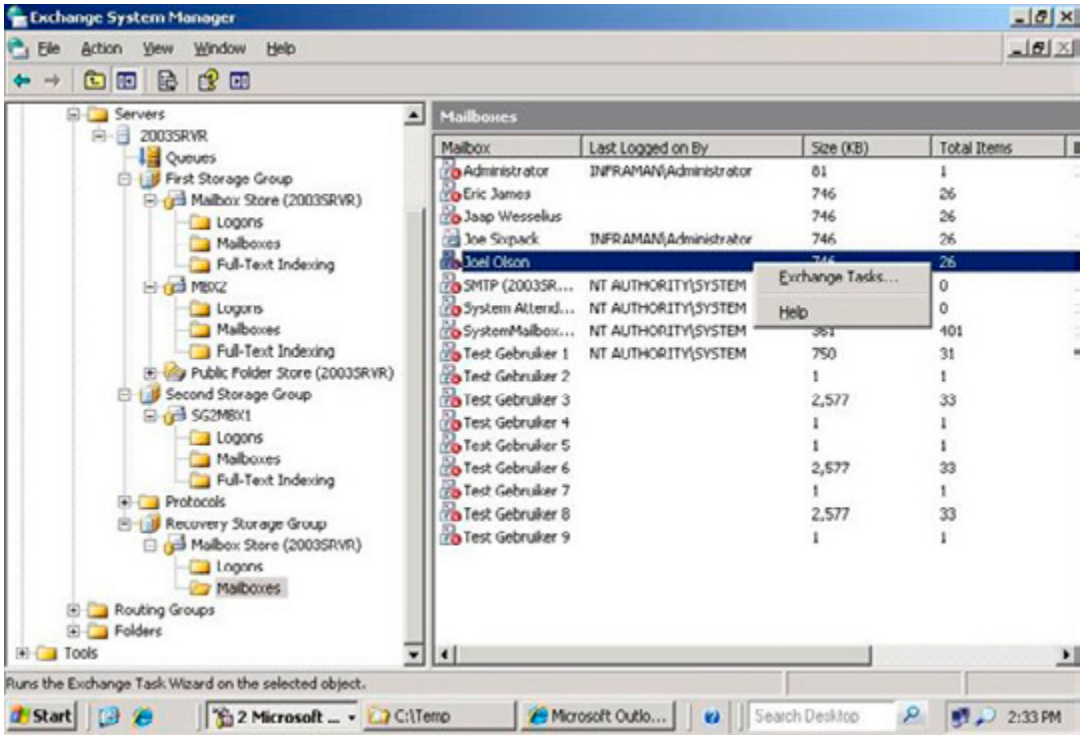

Figure 7. Select Exchange Tasks in the Recovery Storage Group to restore mailbox data.

After selecting Exchange Tasks please select "Recover Mailbox Data" (the only option available) and select the destination mailbox store. This should be the database that hold Joel's mailbox, in this case the "Mailbox Store (2003SRVR)" in the First Storage Group.

Now you have two options to continue:

- **Merge Data** – when you select "Merge Data" the data from the mailbox copy in the Recovery Storage Group will be merged into the original data. If any item exists it will be detected and not be overwritten. Items that no longer exist (i.e., that have been deleted) will be recreated during the merge process.

- **Copy Data** – when you select "Copy Data" the data from the mailbox copy in the Recovery Storage Group will be copied to a separate location, i.e. a special folder named "Recovered Data" followed with the data and time of the recovery process in the user's mailbox.

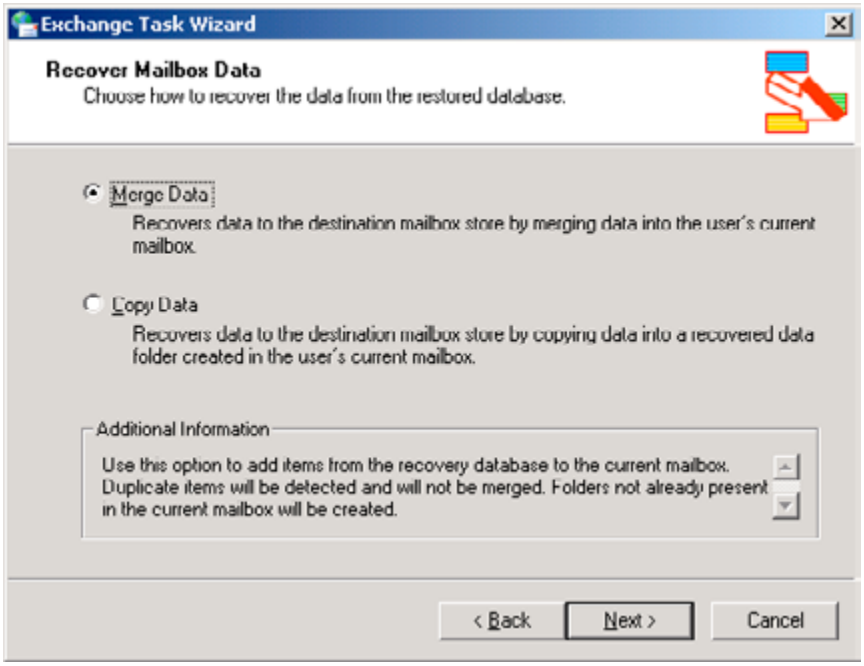

Figure 8. Merge the data into the mailbox or copy it to a "recovered items" folder in the user's mailbox.

When finished you have recovered the mailbox data from the backup that's in the Recovery Storage Group into the user's mailbox.

When using Exchange Server 2007 you have to use the Exchange Management Shell or the Database Recovery Management in the Tools option in the Exchange Management Console.

Again, restore the database that holds Joel's mailbox into the Recovery Storage Group. After restoring you can use the Restore-Mailbox cmdlet in the Exchange Management Console to restore the content to Joel's mailbox.

```
[PS] C:\ >Restore-Mailbox -Identity "Joel Olson" -RSGDatabase "MBXSERVER\RSG\Mailbox Database"

Confirm
Are you sure you want to perform this action?
Recovering mailbox content from the mailbox 'Joel Olson' in the recovery
database 'MBXSERVER\RSG\Mailbox Database' into the mailbox for 'Joel Olson
(Joelo@simpletalk.local)'. The operation can take a long time to complete.
[Y] Yes [A] Yes to All [N] No [L] No to All [S] Suspend [?] Help
(default is "Y"):a

[PS] C:\ >
```

You can also use the Exchange Management Console to recover mailbox items to Joel's mailbox.

- In the left-hand pane in the Exchange Management Console select the Tools option and in the middle pane select "Database Recovery Management" under "Disaster Recovery Tools".
- If needed enter the Exchange Server's name and the Domain Controller's name. Click Next to continue.
- After a couple of seconds some tasks are shown. Select the "Merge or Copy contents" task.
- The next window shows what database is mounted in the Recovery Storage Group; this is the Mailbox Database that was restored earlier. Click "Gather Merge Information" to continue.

- Select "Perform pre-migration tasks" to continue. If the database in the Recovery Storage Group contains a lot of mailboxes you can select "Show Advanced Options" first to set a filter to narrow the selection.
- In the "Select Mailboxes to Copy or Merge" windows select the mailbox that needs to be restored, in this example Joel Olson's mailbox.
- Click "Perform merge actions" to continue.
- When finished the mailbox data have been restored to their original location.

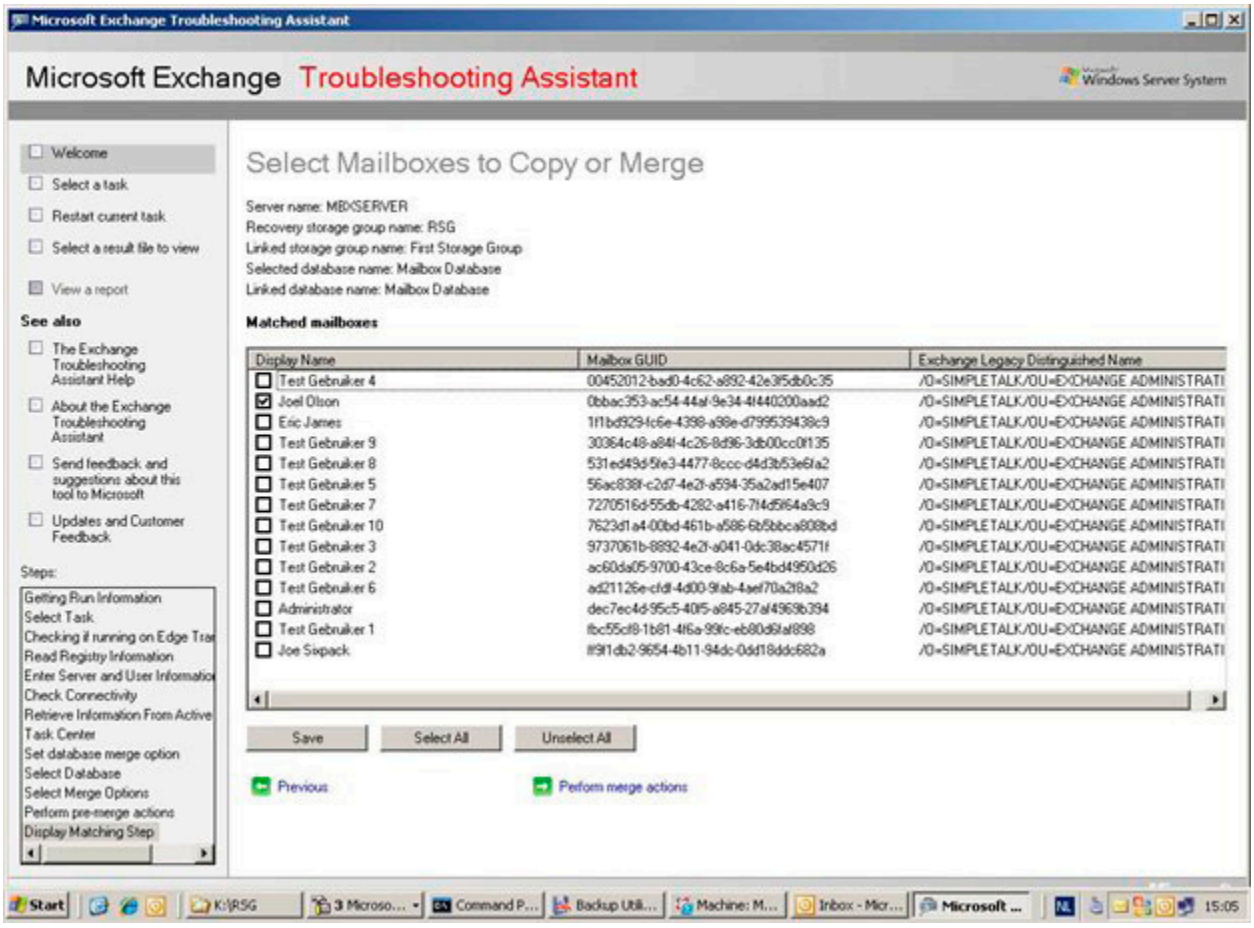

Figure 9. Select the mailbox that needs to be restored.

Conclusion

When you need to retrieve a mailbox or mailbox content from a backup set the Recovery Storage Group is a great tool to use. Depending on the task, you have to perform everything in the Exchange Management Console or you may need to perform some actions in the Exchange Management Shell.

In both Exchange Server 2003 as well as Exchange Server 2007 it is possible to create a .PST file from a restored mailbox, but in Exchange Server 2007 it's a bit more confusing because of the usage of the Exchange Management Shell.

Before you can efficiently use the Recovery Storage Group and its purposes I would advice to try it a couple of times, to see in what scenario's it can help you in your day-to-day Exchange Management tasks. Make sure that you document all steps very well. In this case you know exactly what steps to perform when a request is made to restore a single mailbox, either to its original location or to a .PST file.

Exchange E-mail Addresses and the Outlook Address Cache

12 January 2009

by Ben Lye

Because Exchange auto-complete cache uses X.500 addresses for e-mail sent to addresses within the Exchange organization, it will bounce back messages from a re-created mailbox even after you give the user account all the old SMTP addresses. This is because the old X.500 address in the auto-complete cache is missing, and this causes Exchange to reject the messages. Ben Lye explains how to solve this common problem.

A little while ago I had a case where, after all other troubleshooting had failed, I had to solve a mailbox corruption problem by exporting the mailbox content to a PST file, removing the existing mailbox, recreating a new mailbox, then finally importing the PST file back in. This solved the immediate problem of the corrupt mailbox, but created a new one – when Outlook users tried to e-mail the user either by replying to an existing message or by using Outlook's auto-completion of the user's e-mail address, the message would bounce back to the sender. This happened even though I had re-added all the SMTP addresses that the user previously had. E-mail from external senders was being received properly, and replies to new messages were OK.

This problem occurs because while the Outlook auto-complete cache stores SMTP addresses for e-mail sent to external addresses, it uses X.500 addresses for e-mail sent to addresses within the Exchange organisation. Even though we had given the user account all the old SMTP addresses, the old X.500 address which Outlook was sending to was missing, and this was causing Exchange to reject the messages.

The use of X.500 addresses goes back to before Exchange 2000, when previous versions of Exchange maintained their own LDAP directory. Since Exchange 2000 the mailbox's X.500 address has been stored in the **legacyExchangeDN** attribute in Active Directory. The **legacyExchangeDN** value is set when a mailbox is created, and includes the name of the Exchange administrative group where the mailbox belongs. **LegacyExchangeDN** values typically look like this:

```
/o=Organisation/ou=Administrative Group/cn= Recipients/cn=Username
```

Because the **legacyExchangeDN** value includes the administrative group name changes to admin group names will influence **legacyExchangeDN** values. For example when you upgrade from Exchange 2003 to Exchange 2007 your user-defined admin groups are replaced by a single admin group named "Exchange Administrative Group (FYDIBOHF23SPDLT)" – existing mailboxes are unaffected, but mailboxes created after the upgrade will use the new admin group name in their **legacyExchangeDN** values. (Incidentally, if you've ever wondered why the Exchange 2007 admin group has this name, or what it means, it's the text EXCHANGE12ROCKS, with all the characters shifted to the right by one!)

The current X.500 address of a mailbox can be retrieved from Active Directory using a tool such as ADSIEdit, or LDP.exe, or by using the Exchange Management Shell:

```
[PS] C:\>Get-Mailbox juser | fl LegacyExchangeDN

LegacyExchangeDN : /o=Example/ou=Exchange Administrative Group (FYDIBOHF23SPDLT)/cn=Recipients/
cn=juser

[PS] C:\>
```

The X.500 address of a message sender can be retrieved using a tool such as Microsoft Exchange Server MAPI Editor to open a message and get the **PR_SENDER_EMAIL ADDRESS** property:

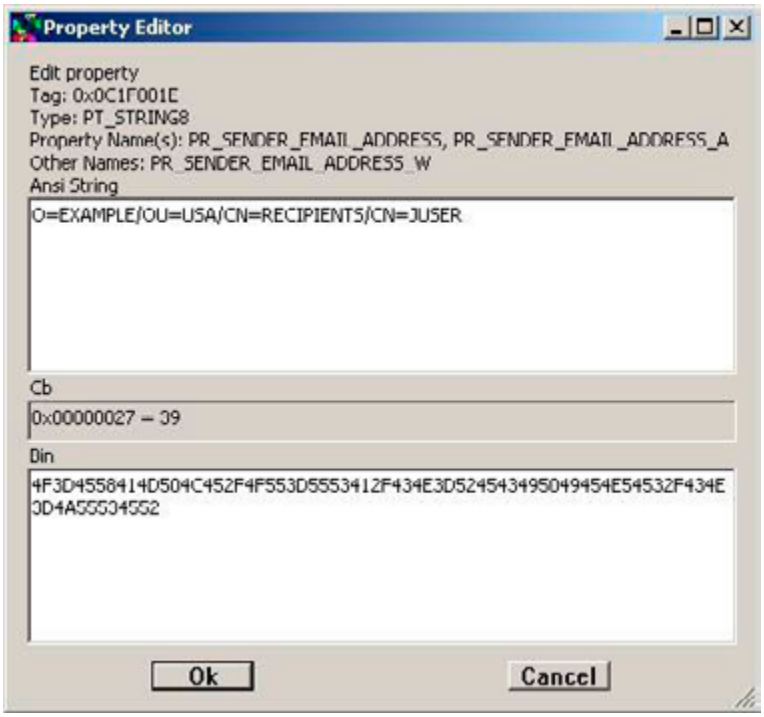

Alternatively, you can use a hex editor to open the Outlook auto-completion cache file and retrieve X.500 addresses from there. The cache is stored in a file in the user's profile, typically ...

```
%userprofile%\AppData\Roaming\Microsoft\Outlook\[Outlook profile name].NK2
```

... on Windows Vista, or ...

```
%userprofile%\Application Data\Microsoft\Outlook\[Outlook profile name].NK2
```

... on Windows 2000, XP or 2003. There are also other tools available on the Internet which will allow viewing and editing of the content of the auto-completion cache file, but they may not expose the X.500 addresses.

Diagnostic information for administrators:

```
Generating server: demo01.example.com

IMCEAEX-_O=COMPANY_OU=USA_cn=Recipients_cn=juser@company.com
#550 5.1.1 RESOLVER.ADR.ExRecipNotFound; not found ##
```

In my case, due to our upgrade to Exchange 2007, the user's **legacyExchangeDN** value had changed from this on the old mailbox (which had been created prior to the Exchange 2007 upgrade):

```
/o=Example/ou=USA/cn=Recipients/cn=juser
```

To this on the new mailbox:

```
/o=Example/ou=Exchange Administrative Group (FYDIBOHF23SPDLT)/cn=Recipients/cn=juser
```

Any new e-mail sent from Outlook using the previously cached X.500 address was being rejected because the old X.500 address no longer existed in the organisation.

The solution to the problem is actually quite simple – add the old **legacyExchangeDN** X.500 address to the new mailbox as a proxy address. You can add an X.500 proxy address through the Exchange Management Console, or the Exchange Management Shell.

To add the proxy address in the console, double-click the mailbox you need to add the proxy address to, go to the E-Mail Addresses property page, and add a new custom address:

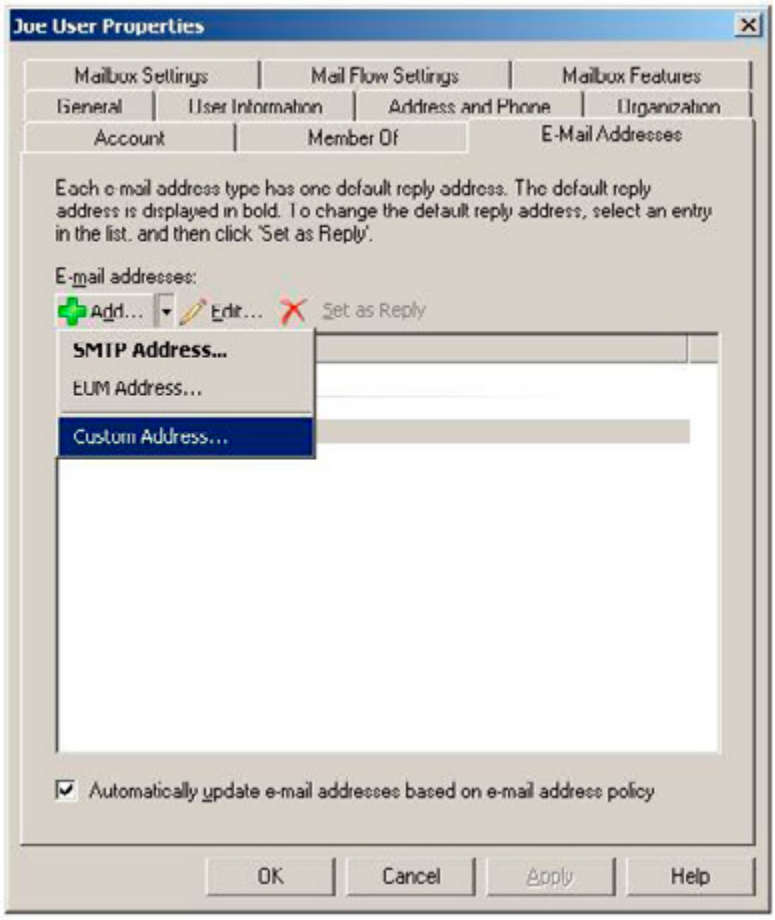

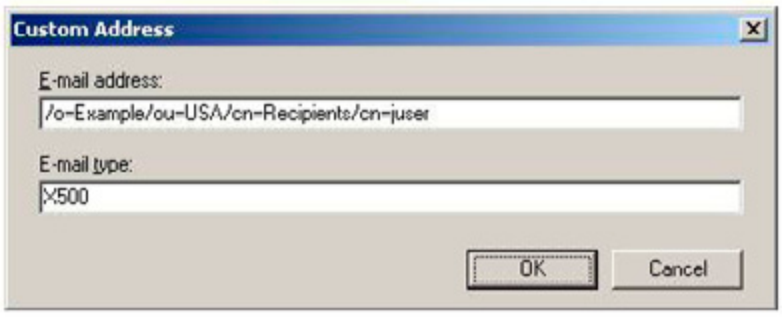

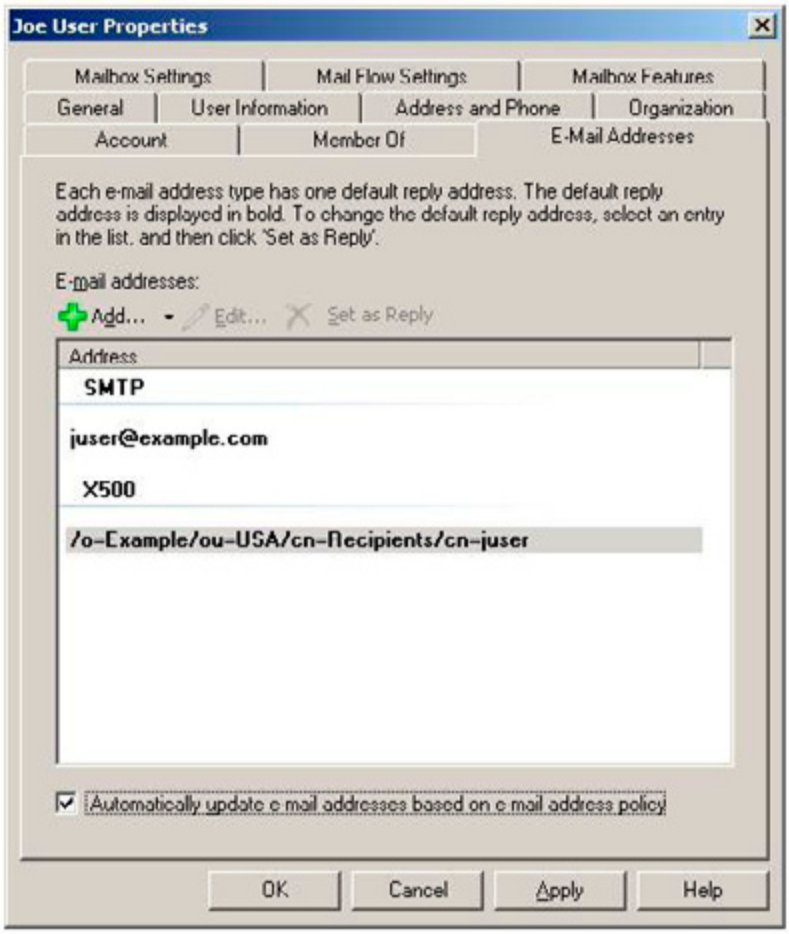

To add the proxy address in the shell we use the **Get-Mailbox** and **Set-Mailbox** cmdlets:

```
[PS] C:\>$ProxyAddresses = (Get-Mailbox juser).EmailAddresses
[PS] C:\>$ProxyAddresses += [Microsoft.Exchange.Data.CustomProxyAddress]("X500:/o=Example/ou=USA/
cn=Recipients/cn=juser")
[PS] C:\>Set-Mailbox juser -EmailAddresses $ProxyAddresses
```

Breaking these commands down:

```
[PS] C:\>$ProxyAddresses = (Get-Mailbox juser).EmailAddresses
```

...retrieves the existing proxy addresses for the mailbox and stores them in the $ProxyAddresses variable.

```
[PS] C:\>$ProxyAddresses += [Microsoft.Exchange.Data.CustomProxyAddress]("X500:/o=Example/ou=USA/
cn=Recipients/cn=juser")
```

...adds the new X.500 proxy address to the variable which contains the existing proxy addresses.

```
[PS] C:\>Set-Mailbox juser -EmailAddresses $ProxyAddresses
Updates the mailbox with the new set of proxy addresses
```

This technique can be used to solve this problem in a number of other scenarios where the **legacyExchangeDN** attribute has changed, and is not limited to mailboxes. For example, if someone leaves the Exchange organisation and you want their e-mail to go to an external e-mail address you would create a contact record with the necessary SMTP proxy addresses. If you also added the **legacyExchangeDN** of the old mailbox to the contact record as an X.500 proxy address Outlook users wouldn't get bounced messages if they used the old entry in their auto-complete caches.

Upgrading to Exchange Server 2007

20 February 2009

by Jaap Wesselius

> ***Jaap starts a series on the necessary steps to migrate an existing installation of Exchange to Exchange Server 2007. Of course it's simple! Everybody says so, but it is in the detail where you can get caught out.***

There are several reasons to migrate an existing Exchange Server 2003 setup to Exchange Server 2007. Beside the benefit of the the new features you get when Outlook 2007 is used in combination with Exchange Server, some of the most compelling business cases for doing so are

- the availability options of the mailbox servers
- the scalability of the 64-bit architecture
- the option to connect your PABX infrastructure to your messaging infrastructure.

Currently, I'm involved in consolidating and centralizing a 25.000-mailbox Exchange system and upgrading it to Exchange Server 2007. The customer intends that this new Exchange Server 2007 environment should be the foundation for their Unified Messaging infrastructure.

There are several papers available on the Internet about making the transition to Exchange Server 2007 that imply that it is a simple operation: But is it really that simple?

Exchange Server 2003 environment

Suppose we have an existing Exchange Server 2003 setup that consists of several mailbox servers in different places. There are also two Exchange servers configured as "front-end servers'. These servers are basically the same as the mailbox servers, but they are configured as protocol proxy servers.

SMTP, Outlook Web Access, and Activesync traffic originating from Windows Mobile devices, all enter the Exchange organization on the front-end servers. The front-end servers do not host mailboxes or public folders. To provide a higher scalability the front-end servers can be placed into a Network Load Balancing (NLB) cluster where the load from all clients can be spread across several front-end servers. These servers then use one common namespace, for example webmail.inframan.nl.

The Exchange Server 2003 environment is running on a Windows Server 2003 environment in a single forest, single domain configuration. The Active Directory domain is running in Windows Server 2000 native mode.

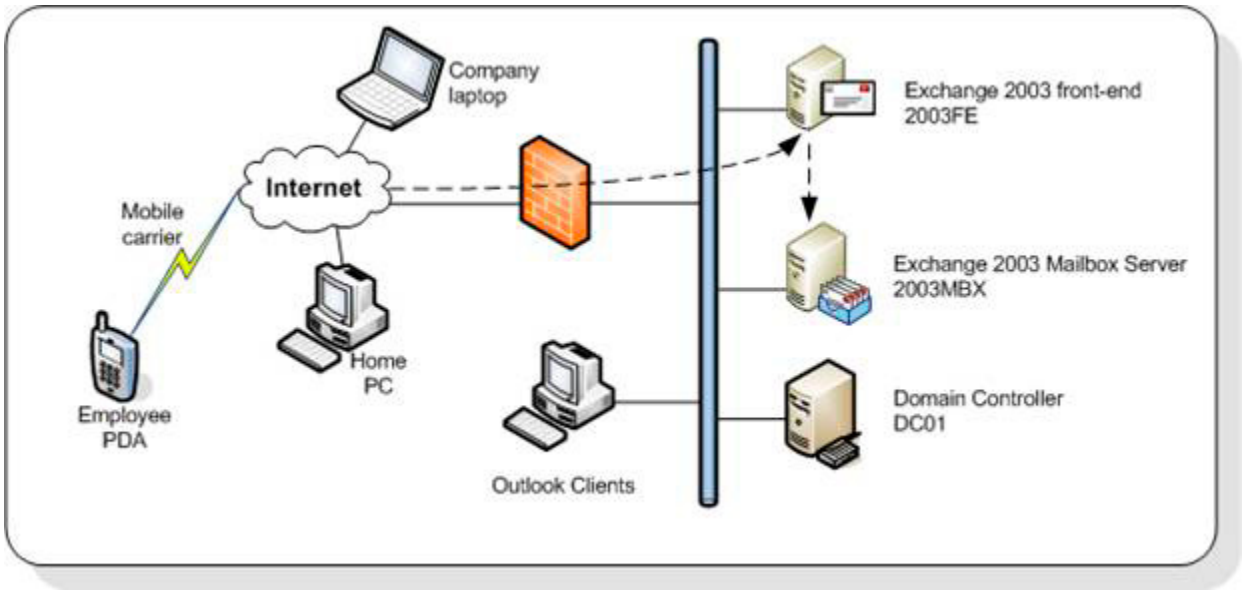

Figure 1. The existing Windows Server 2003 and Exchange Server 2003 infrastructure.

Internal Outlook 2003 clients connect directly to the Exchange Server 2003 mailbox server using the default MAPI protocol. The server that is hosting the user's mailbox is configured in the profile that Outlook 2003 is using. Laptop users can either use a VPN connection to connect to the internal network or use the "RPC over HTTPS" functionality in Outlook 2003. A couple of clients are also using Outlook 2007, but, because only Exchange Server 2003 is used in the company, any additional new functionality such as "Autodiscover" or the new "Exchange Web Services for free/busy or Out-of-Office functionality" are not available. Home PC's can use the Exchange Server 2003 Outlook Web Access. PDAs use the company's preferred mobile carrier to access the Exchange Server Activesync service.

Figure 2. Windows Mobile 6 working with an Exchange Server 2003.

To provide a secure connection from all outside clients, an SSL connection is used for all clients. This is a simple SSL certificate with a single name: https://webmail.inframan.nl. If you are using Windows Mobile for Activesync in your Exchange environment, please make sure that the certificate you're using is a supported one. Some vendors offer certificates that are working fine with websites and are fully

supported by all kinds of browsers, but have issues with Outlook RPC over HTTPS and Windows Mobile. But if you're using a known certificate, in this example from Digicert, RPC over HTTPS and Windows Mobile work great.

> ***Note***
>
> *If you want to test your connectivity for Windows Mobile you can download the Windows Mobile emulator from Microsoft. This emulator can be run on any Windows computer and as long as you have network connectivity you can use the emulator. Figure 2 is a screenshot from my laptop running at home, while the Exchange 2003 environment is running in a datacenter 100 miles from my home, just over an ADSL connection. You can download the Windows Mobile emulator at* Microsoft Device Emulator 3.0 – Standalone Release.

In the Exchange 2003 organization there's one Administrative Group. An Administrative Group is a management boundary. With Administrative Groups you can use "delegation of control." Suppose there are multiple messaging departments in your organization, and each department has control over its own Exchange Servers. In this case you can use multiple Administrative Groups, one for each department. Every messaging administrator in a department has full control over his own Exchange Servers, and not over other departments' Exchange Servers.

Upgrading to Exchange Server 2007

Exchange Server 2007 offers all kinds of new functionality:

- Support for 64-bit hardware for better scalability
- Autodiscover functionality for automatically configuring Outlook 2007 clients
- Availability Services and webbased Offline Address Book distribution as a replacement for the Free/Busy and Offline Address Book downloads in Public Folders for Outlook 2007 and higher
- Unified Messaging server role to connect your phone system (PABX) to your Exchange environment
- Multiple server roles to separate functionality and better scalability
- Exchange Server 2007 also uses a different administration model and uses a different routing model. This means that Exchange Server 2007 can cooperate with Exchange Server 2003, but it may require some significant changes to your environment.

The following steps need to be performed to prepare your Exchange Server 2003 environment for the implementation of Exchange Server 2007:

- The Domain Controllers and Global Catalog server need to be on Windows Server 2003 SP1 level or higher
- The Active Directory domain functional level needs to be "Windows 2000 native mode"
- No Exchange Server 5.5 may exist in your Exchange organization. To enforce this requirement the Exchange organization needs to be running in "native mode"
- The Active Directory Schema needs to be updated to support Exchange Server 2007
- The Active Directory organization needs to be upgraded to include Exchange Server 2007 permissions and system objects
- The Active Directory domain needs to be upgraded to include Exchange Server 2007 permissions and system objects.

When you have performed these steps it is a common best practice to run the Exchange Best Practices Analyzer (ExBPA) and perform an "Exchange 2007 Readiness Check." This will check the current infrastructure for its readiness for Exchange Server 2007. ExBPA can be downloaded from the Microsoft website or from Microsoft Exchange Analyzers. When ExBPA confirms the readiness you can proceed with the actual configuration changes and installation of Exchange Server 2007.

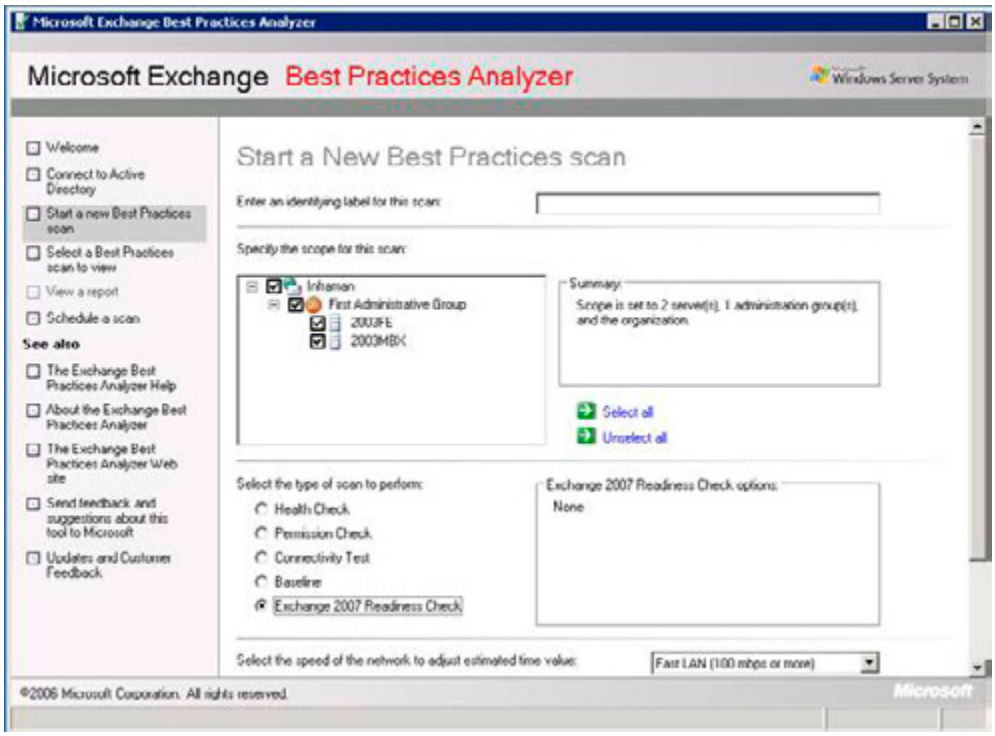

Figure 3. Exchange Best Practices Analyzer – Exchange 2007 Readiness Check.

Upgrading the Active Directory

The first step in changing the configuration for Exchange Server 2007 is upgrading the Active Directory schema to Exchange Server 2007 SP1. You can achieve this by running the following commands on a command prompt on the Active Directory schema master from the Exchange Server 2007 installation media:

```
Setup.com /PrepareLegacyExchangePermissions inframan.local
Setup.com /PrepareSchema
```

The first command with the **/PrepareLegacyExchangePermissions** ensures that the Recipient Update Service in Exchange Server 2003 continuous to run correctly after the schema change to Exchange Server 2007 by granting new permissions. This must be performed before the actual upgrade of the Schema, which is done with the second command.

To check what version your schema is or to check if the upgrade was successful you can check the Schema by using a tool like **ADSIEDIT** or **LDP.EXE** and check the `CN=ms-Exch-Schema-Version-Pt` object. Its property "rangeUpper" should have the value 11116 after the schema change. The property can have the following values:

Value	Corresponding Exchange version
4397	Exchange Server 2000 RTM
4406	Exchange Server 2000 service pack 3
6870	Exchange Server 2003 RTM
6936	Exchange Server 2003 service pack 2
10628	Exchange Server 2007
11116	Exchange Server 2007 service pack 1

Note

If you have multiple domain controllers in your Exchange Server environment you have wait for the Domain Controller replication to finish.

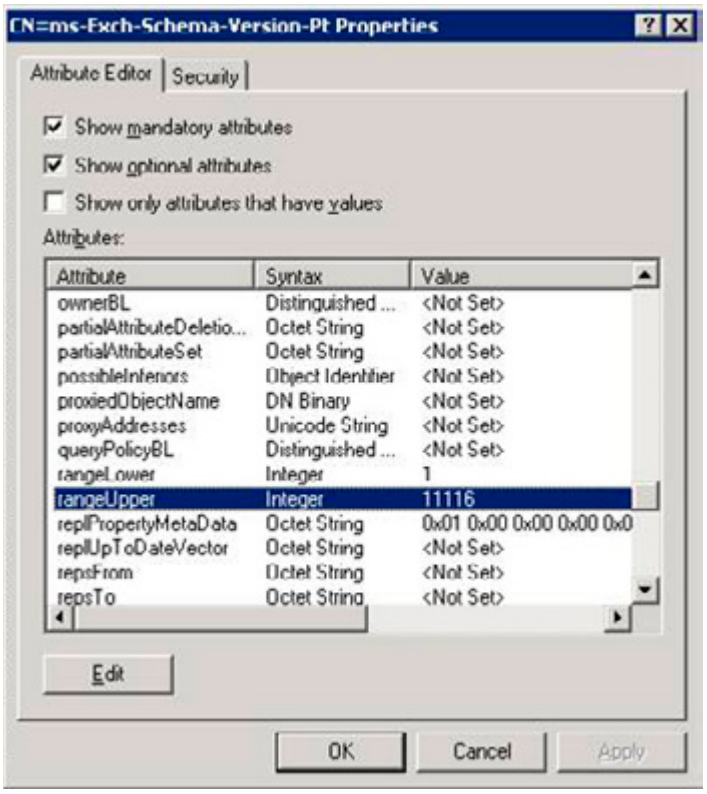

Figure 4. Check the schema version. This schema is on the Exchange Server 2007 SP1 level.

After upgrading the Schema the current Exchange Server 2003 organization can be upgraded to Exchange Server 2007. This is achieved by running the following command from the Exchange Server 2007 installation media:

```
Setup.com /PrepareAD /
```

Running this command wil configure the global Exchange objects in Active Directory (residing in the Configuration container of Active Directory), creates the Exchange Universal Security Groups in the root of the domain and it prepares the current inframan.local domain for Exchange Server 2007.

This command also creates the Exchange 2007 Administrative Group called "Exchange Administrative Group (FYDIBOHF23SPDLT)" and it creates the Exchange 2007 Routing Group called "Exchange Routing Group (DWBGZMFD01QNBJR)."

Note

For those wondering where FYDIBOHF23SPDLT and DWBGZMFD01QNBJR come from: take the string EXCHANGE12ROCKS and increase all individual letters with one (E becomes F, X becomes Y, etc) or decrease all individual letters (E becomes D, X becomes W, etc.).

To verify that this step completed successfully, make sure that there is a new organizational unit (OU) in the root domain called Microsoft Exchange Security Groups. This OU should contain the following new Exchange USGs:

- Exchange Organization Administrators
- Exchange Recipient Administrators
- Exchange View-Only Administrators
- Exchange Servers
- ExchangeLegacyInterop.

After performing this step the new Administrative Group will show up in the Exchange System Manager on an "old" Exchange Server 2003 machine (Figure 5).

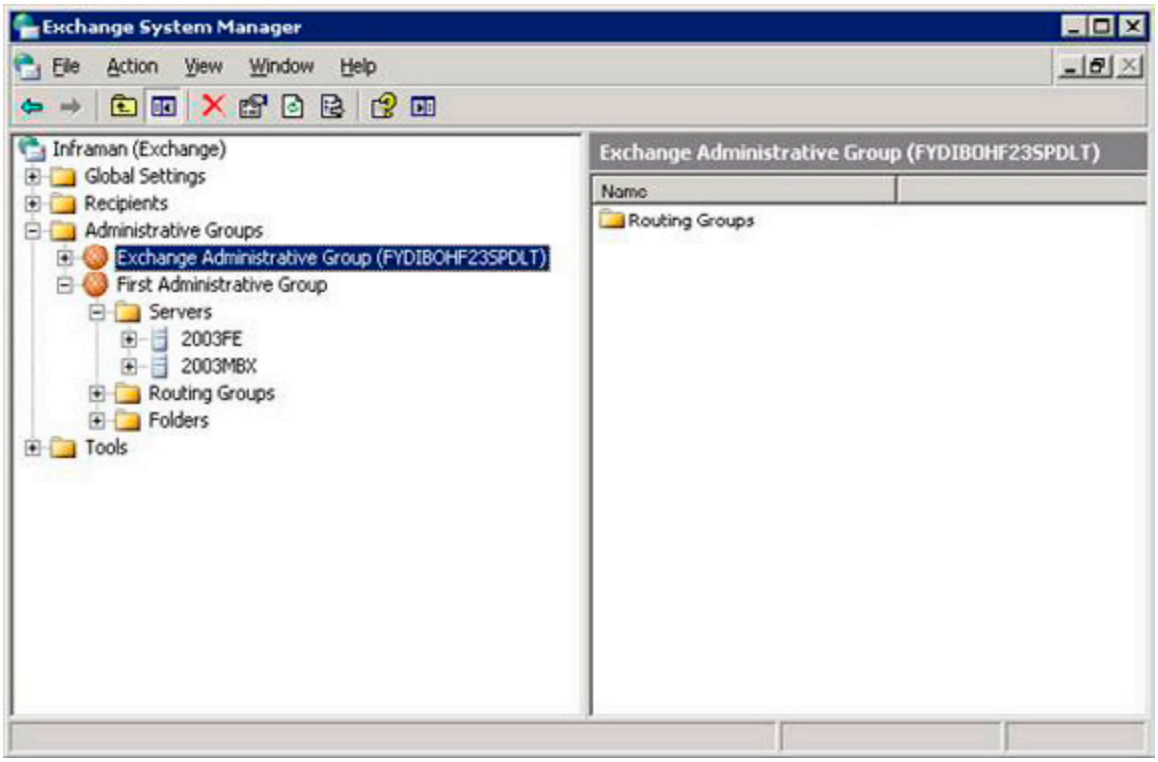

Figure 5. An additional Administrative Group appears after preparing the Exchange 2003 organization.

The last step in preparing your environment for the implementation of Exchange Server 2007 is to prepare the Active Directory domain or domains for Exchange Server 2007. The domain is prepared by running the following command from the Exchange Server 2007 installation media:

```
Setup.com /PrepareDomain
```

This step sets permission on the Exchange Server container in Active Directory and it creates a new Global Group called "Exchange install domain servers" in the domain where the command is run. It also assigns permissions for the Exchange Servers Universal Security Group (USG).

After performing these steps the Active Directory and Exchange Server environment is fully prepared for the installation of the first Exchange Server 2007 server.

Installing the first Exchange Server 2007 server

Installing the first Exchange Server 2007 server should be done carefully since Exchange Server 2007 is fully compatible with Exchange Server 2003, but not vice versa. This means that Exchange Server 2007 CAS and Hub Server can work with Exchange Server 2003 mailbox servers, but Exchange Server 2003 front-end servers cannot work with Exchange Server 2007 mailbox servers.

This automatically means that when installing multiple Exchange Server 2007 servers in the Exchange Server 2003 environment the first Exchange Server 2007 server that will be installed needs to be a Client Access Server and a Hub Transport Server. In our scenario we will also install a dedicated Mailbox Server role as depicted in Figure 6.

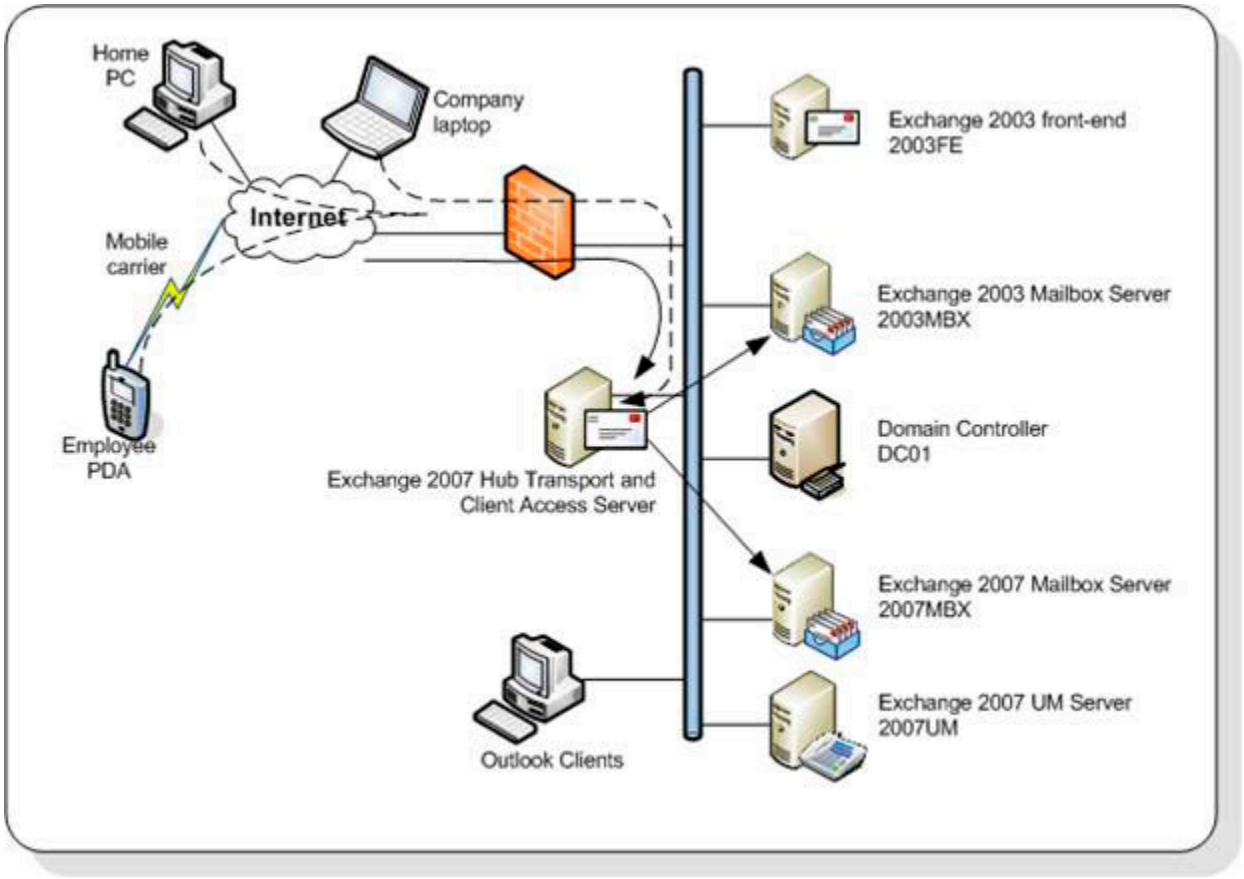

Figure 6. The Infrastructure with the Exchange Server 2007 coexistence.

The first server that will be installed will be a combined Hub Transport and a Client Access Server. The prerequisites for running an Exchange Server 2007 server are:

- An X64 version of Windows Server 2003 or Windows Server 2008. There is an X86 version of Exchange Server 2007 available, but this is only for test- and development purposes. The X86 version is not supported in a production environment.
- .NET Framework 2.0.
- Powershell 1.0.
- IIS 6.0 for the Client Access Server role. If only a Hub Transport Server is installed the IIS6 component isn't needed, although installing it is useful for management purposes.

During the installation the setup program will ask where to connect to the existing infrastructure. A little background: Exchange Server 2007 will be installed in a separate Administrative Group in the Exchange organization. Exchange Server 2007 does not use the routing infrastructure that Exchange Server 2000 and Exchange Server 2003 used for routing messages; instead it relies on Active Directory Sites and Services. This is the same routing infrastructure Windows uses for Active Directory replication traffic. The Exchange Server 2007 routing model and the Exchange Server 2003 routing model are not compatible, so Exchange Server 2007 has a legacy Routing Group, just for connecting with an existing Exchange Server 2003 Routing Group. When this question is presented the answer is used for creating the legacy Routing Group Connector between Exchange Server 2003 and Exchange Server 2007. Using this connector messages will be routed between Exchange Server 2003 and Exchange Server 2007.

After installation of the Exchange Server 2007 Hub Transport and Client Access Server inbound messages will arrive on the new Hub Transport Server. Since all mailboxes still reside on the Exchange Server 2003 mailbox server all messages will be relayed to this server via the legacy Routing Group Connector.

Outlook Web Access (OWA), RPC over HTTPS (now called Outlook Anywhere) and Exchange ActiveSync can also be transferred to the new Client Access Server. But... there a couple of caveats that you need to be aware of:

- As soon as you implement Exchange 2007 and start migrating mailboxes to the new mailbox server, Outlook 2007 will notice immediately. Outlook 2007 has new functionality called autodiscover and Exchange Web Services. When starting Outlook 2007 it will query Active Directory for a Client Access Server object called the Service Connection Point (SCP). You have to be very carefull about the order in which to install the various server roles.
- The services mentioned in point 1 are all HTTPS based services. Even when you don't use RPC over HTTPS (aka Outlook Anywhere) at all, Outlook 2007 will use these services over HTTPS.
- When clients are domain joined this should not be a problem since Outlook will use the default FQDN's as used during the installation (i.e. 2007CASHUB.inframan.local). If clients are not domain joined Outlook 2007 will try connecting using the domain part of an e-mail address, i.e. @inframan.nl. Since no certificates are installed yet and the external DNS is not registered yet this will fail and users will start noticing errors like "Outlook not able to download the Offline Address Book."

So, the first server that we are going to install is a combined Hub Transport Server and Client Access Server. There are two ways to install an Exchange Server 2007, via the command line or using the GUI setup. Both can be started from the root of the installation media. The advantage of the command line is that you can script it, which makes it possible to install a larger number of servers in an identical manner.

The Exchange Server can be installed by starting the setup application (setup.exe) from the Exchange Server 2007 installation media. This will check for the prerequisites to be installed and will show a graphical interface where the roles to be installed can be selected:

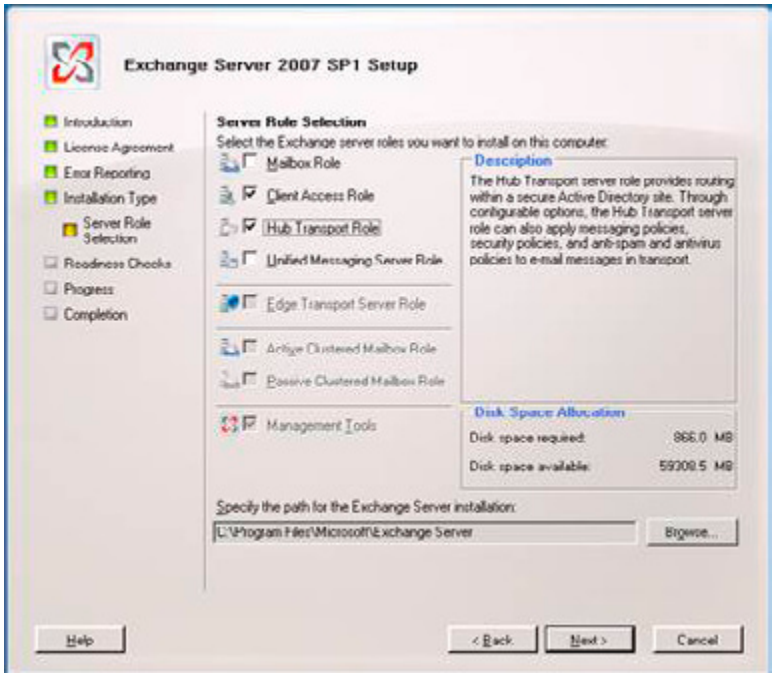

Figure 7. The setup program with only Hub Transport and Client Access selected.

When the installation is finished the first Exchange Server 2007 is installed and the server object is configured in the Active Directory. Outlook 2007 will immediately notice this and start using the Client Access Server.

Certificates

The next step is to configure certificates for the Client Access Server. The name of the certificate can be webmail.inframan.nl, just like the Exchange Server 2003 front-end server. But Outlook 2007 and Windows Mobile 6 can also use the autodiscover functionality. Outlook 2007 will setup an additional connection to the Client Access Server via autodiscover.inframan.nl. This is an HTTPS connection, so it needs a certificate. To use the same Client Access Server a so called "Unified Communications" or SAN (Subject Alternate Name) certificate needs to be used. This type of certificate can have multiple names. Besides the external names you should also register its internal name. When clients connect to the Client Access Server from the internal network the name of the Client Access Server can be resolved to its internal name, like 2007CASHUB.inframan.local.

The names that should be used in this case should be:

- Webmail.inframan.nl
- Autodiscover.inframan.nl
- 2007cashub.inframan.local
- Mail.inframan.nl (for SMTP).

A certificate in Exchange Server 2007 can be requested by using the New-ExchangeCertificate commandlet in the Exchange Management Shell on the Client Access Server:

```
New-ExchangeCertificate –GenerateRequest
`
  –SubjectName
"c=NL,o=Inframan,cn=webmail.inframan.nl"`
  –DomainName webmail.inframan.nl, autodiscover.inframan.nl,
`
  mail.inframan.nl, 2007cashub.inframan.local
  –FriendlyName webmail.inframan.nl
`
  –PrivateKeyExportable:
$TRUE`
  –path c:\cert_req.txt
```

This command will generate a certificate request file that can be submitted at your own certificate authority.

After approval by the DNS manager a certificate will be sent by your certificate authority that can be imported on the Client Access Server by using the Import-ExchangeCertificate commandlet in the Exchange Management Shell. The output of this commandlet can be piped into the Enable-ExchangeCertificate to enable the certificate after importing it:

```
Import-ExchangeCertificate -path c:\newcert.cer | Enable-ExchangeCertificate -services "iis, smtp"
```

Note

If needed this certificate can also be used for POP3, IMAP4 and Unified Messaging usage. In this case you can add these services on the command-line by typing –Services "IIS,SMTP,POP,IMAP,UM".

The Mailbox Server Role

After installing the Hub Transport Server and Client Access Server the Mailbox Server can be installed. When all prerequisites are met the installation can be started by entering the following command from the installation media on the 2007MBX server:

```
Setup.com /mode:install /roles:MB
```

This will automatically install only the Mailbox Server role on the particular server. The setup will automatically detect the existing Exchange 2003 environment and configure itself accordingly. When the setup is finished there will be a fully functional Exchange Server 2007 environment integrated in the Exchange Server 2003 environment.

Please note that Exchange Server 2003 as well as Exchange Server 2007 need to be managed with their own management tools. Exchange Server 2003 need to be managed with the Exchange System Manager, the Exchange Server 2007 need to be managed with the Exchange Management Console or the Exchange Management Shell.

In my next article I will explain a bit more on the coexistence phase with the two versions of Exchange, how to move mailboxes from Exchange Server 2003 to Exchange Server 2007 and how to decommission the Exchange Server 2003 environment.

Goodbye Exchange ExMerge, Hello Export-Mailbox

26 February 2009

by Ben Lye

> ***ExMerge was a great way of exporting a mailbox to an Exchange PST file, or for removing all occurences of an email virus, but it has now been replaced by two new Exchange Management Shell cmdlets, Export-Mailbox and Import-Mailbox which are equally as useful to have at hand.***

Most Exchange administrators who've worked with Exchange prior to Exchange 2007 will be familiar with ExMerge. It was a useful tool to have available if you needed to export a mailbox to a PST file, for example if a user left and wanted to take a copy of their mailbox with them, or you needed a PST file as an archive copy of one or more mailboxes. It was also useful if you needed to search your Exchange databases for a particular message and remove it, such as removing an email virus.

ExMerge was born in the Exchange 5.5 days as a product support tool, and was later re-released for Exchange 2003 as a web download. In Exchange 2007 ExMerge has been replaced by two new Exchange Management Shell cmdlets, Export-Mailbox and Import-Mailbox. The RTM release of Export-Mailbox was limited to moving messages only from one mailbox to another, but the SP1 release was enhanced to include the export-to-PST functionality which is familiar to ExMerge users, and that's the version I'm referring to in this article. There is no graphical interface for either Export-Mailbox or Import-Mailbox; they are only available in the shell.

It's important to note that Export-Mailbox is not intended for use as an Exchange migration tool. It's designed for moving mailbox content rather than an entire mailbox. If you need to migrate an entire mailbox you should use the Move-Mailbox cmdlet.

At a high level, Export-Mailbox can be used to

- Export mailbox content from a mailbox to another mailbox
- Export mailbox content from a mailbox to a PST file.

During mailbox exports content can be filtered by:

- Included or excluded folders
- Message sender keywords
- Message recipient keywords
- Message subject keywords
- Message and attachment content keywords
- Attachment file names
- Date range.

If you're going to use the filter options, make sure you have at least Update Rollup 4 for Exchange 2007 SP1, as this update included improvements to the filtering options.

When keyword filter options are specified Export-Mailbox will first export all the messages in each from the source mailbox to the destination mailbox. The folder in the destination mailbox is then searched and messages which do not match the keyword filters are deleted. On large mailboxes this can be a time consuming and resource intensive operation.

Other options are available to:

- merge content from the source to the destination (the top level folder is not time stamped)
- delete content from the source mailbox after it has been exported
- include associated messages such as rules, views, and forms
- increase the number of threads
- write output to an XML log file.

The source and destination mailbox must be located on one of these versions of Exchange:

- Exchange 2007
- Exchange Server 2003 SP2 or later
- Exchange 2000 SP3 or later.

Export-Mailbox will export any items in mailbox's dumpster by converting them back to regular messages.

For multiple mailbox exports, the output of the **Get-Recipient** or **Get-Mailbox** cmdlets can be piped to **Export-Mailbox.** When piping input to **Export-Mailbox,** the **MaxThreads** parameter can be used to increase the number of mailboxes processed simultaneously.

Copying mailbox content to a subfolder in another mailbox

This might be the case if you have a user who has left and you need to move some or all of their mailbox content into the mailbox of another user such as their co-worker, manager, or replacement, before the leaving user's mailbox is deleted.

This is the command to export the content from Joe's mailbox to a folder name "Joe" in Bob's mailbox:

```
Export-Mailbox -Identity joe@example.com -TargetMailbox bob@example.com -TargetFolder Joe
```

The exported content will look as below.

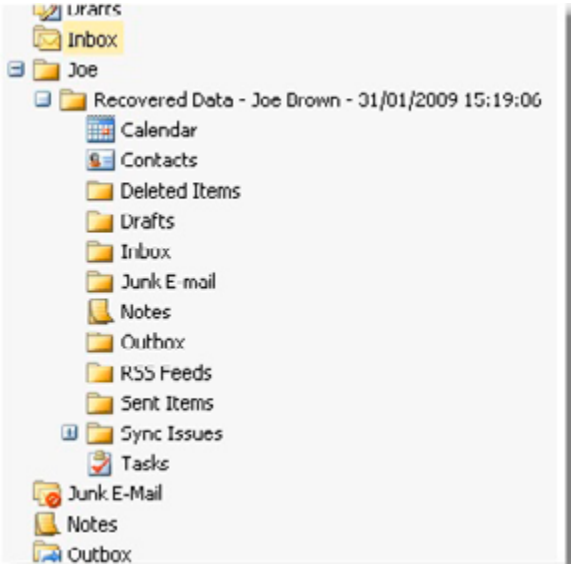

Optionally the *AllowMerge* parameter can be used to merge the content into the target folder. When using the *AllowMerge* parameter the top-level folder isn't time-stamped, this means that the source mailbox content can be exported to the destination multiple times:

```
Export-Mailbox -Identity joe@example.com -TargetMailbox bob@example.com -TargetFolder Joe -
AllowMerge
```

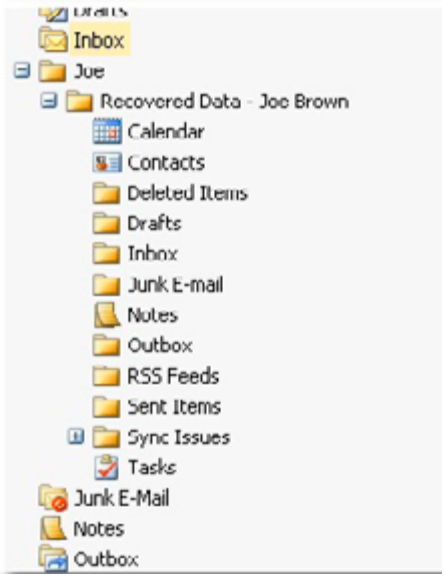

Specific folders can be included or excluded by using the **IncludeFolders** and **ExcludeFolders** parameters. For example, to include only the Inbox and Contacts folders:

```
Export-Mailbox -Identity joe@example.com -TargetMailbox bob@example.com -TargetFolder Joe -
IncludeFolders \Inbox, \Contacts
```

To export all content except for the Sent Items and Deleted Items folders:

```
Export-Mailbox -Identity joe@example.com -TargetMailbox bob@example.com -TargetFolder Joe -
ExcludeFolders "\Sent Items," "\Contacts"
```

Exporting mailbox content to a PST file

If a user is leaving and wants to take a copy of their mailbox you can use Export-Mailbox to move the data directly into a PST file.

> **Tip:** *To export mailbox content to a PST file you must run Export-Mailbox on a 32-bit Windows computer running the 32-bit version of the Exchange 2007 management tools, and Microsoft Outlook 2003 SP2 or later.*

To export a user named Joe's mailbox to the PST file C:\Temp\joe.pst:

```
Export-Mailbox -Identity joe@example.com –PSTFolderPath C:\Temp\joe.pst
```

To export all the mailboxes in the sales mailbox database to individual PST files named <alias.pst>:

```
Get-Mailbox –Database "Sales Mailbox DB" | Export-Mailbox –PSTFolderPath C:\PSTFiles
```

Searching for and removing content from a mailbox

If a virus has found its way into your Exchange organisation, or if a message has been delivered to a large number of mailboxes and you need to remove it, you can use Export-Mailbox to search mailboxes and remove the message.

> **Tip:** *Because Export-Mailbox first copies all content to the destination mailbox before performing the search the target mailbox can get quite large and will have a lot of IO. It's a good idea to use a mailbox created specifically for this task, especially if you are performing the search over many mailboxes and using the MaxThreads parameter to increase the default number of threads.*

To remove any message with the words "Company confidential" in the subject line from all mailboxes on the server EXCHANGE01, processing 10 mailboxes at a time:

```
Get-Mailbox -Server EXCHANGE01 | Export-Mailbox -TargetMailbox ExportMailbox -TargetFolder
MessageCleanup -SubjectKeywords "Company confidential" –DeleteContent –MaxThreads 10
```

To remove any messages from the sender NASTY.PERSON@EXAMPLE.COM from all mailboxes on the server EXCHANGE01, processing 10 mailboxes at a time:

```
>Get-Mailbox -Server EXCHANGE01 | Export-Mailbox -TargetMailbox ExportMailbox -TargetFolder
MessageCleanup -SenderKeywords nasty.person@example.com –DeleteContent –MaxThreads 10
```

The messages will be copied to the "MessageCleanup" folder in the target mailbox and deleted from the source mailboxes. The target folder will include a replica of the folder structure in each of the source mailboxes.

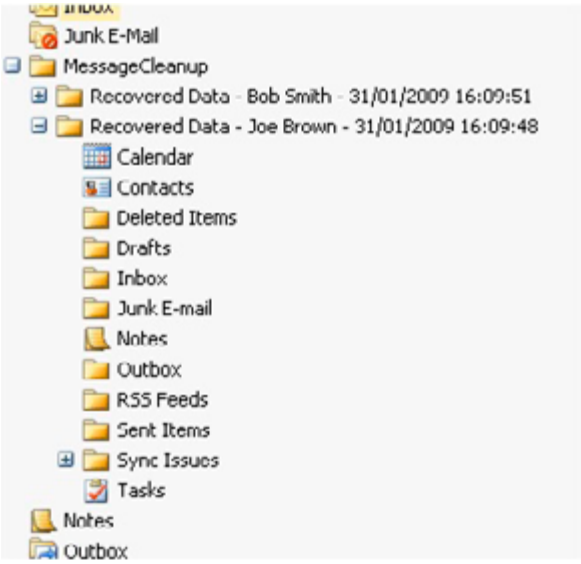

Troubleshooting Export-Mailbox

The most common problems with using Export-Mailbox are related to permissions. To use Export-Mailbox you need to be delegated the Exchange Server Administrator role, be a member of the local Administrators group for the target server, and have full access to the source and destination mailboxes. The source and destination mailboxes must be in the same Active Directory forest.

These are some common permissions-related errors:

Error	Cause
"The specified mailbox database [Mailbox Database Name] does not exist"	The user running the Export-Mailbox command needs to be delegated the Exchange Administrator role for the Exchange server.
"Error occurred in the step: Creating target folder in the target mailbox. An unknown error has occurred., error code: -2147221233"	The user running the Export-Mailbox does not have full access to the destination mailbox.
Error occurred in the step: Moving messages. Failed to copy messages to the destination mailbox store with error: MAPI or an unspecified service provider. ID no: 00000000-0000-00000000, error code: -1056749164"	The user running the Export-Mailbox does not have full access to the source mailbox.

ExMerge may be gone, but Export-Mailbox is equally as useful a tool to have at hand. The flexible filtering options make it possible to do more granular exports than were possible with ExMerge, and the command shell interface makes it easy to script.

More information on Export-Mailbox can be found on the Microsoft TechNet website:

http://technet.microsoft.com/en-us/library/aa998579.aspx

More information about the dumpster can also be found on the TechNet website:

http://technet.microsoft.com/en-us/library/aa997155.aspx

Determining MS Exchange Disk Performance

26 March 2009

by Michael B. Smith

With an Exchange Server, how do you measure performance? How, in particular do you do about measuring your disk subsystem's performance? Whereas the CPU Usage and Memory usage are both easy to measure, the measurement of disk performance is often described as a "black art." In fact, it isn't so difficult, as Michael Smith explains.

Performance – The Big Three

It is a real shame, actually. There are three major components that affect a computer's performance: the speed and power of its processor, how much memory that it has, and the performance of its disk subsystem. Figuring out those first two - whether you need more memory and/or processing power - is generally pretty easy. For an example, refer to Figure 1.

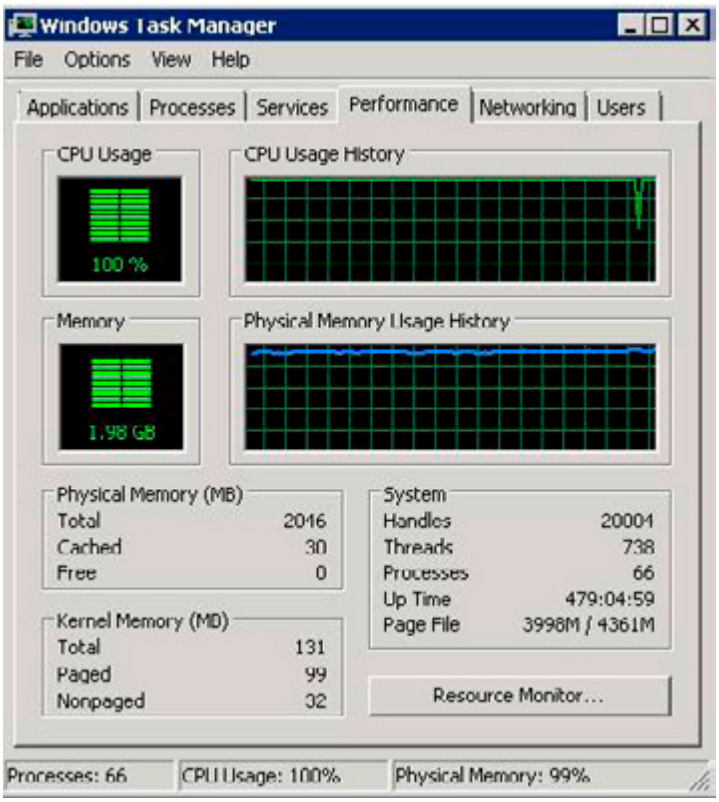

Figure 1: The First Sample Windows Task Manager Dialog.

In Figure 1, it's pretty easy to see that the processor on this computer is overburdened. For the entirety of the CPU Usage History graphical display, the utilization of the processor was almost always 100%. Similarly, when we look at the various memory indicators, it is also pretty plain that the memory on this computer is in short supply. It takes one more indicator to be certain of that, but we can reach that conclusion based on the following indicators:

- Physical memory is currently 99% utilized
- Free physical memory is at zero
- The usage history shows that physical memory has constantly been highly utilized
- The page file is 92% in use (3998 MB / 4361 MB).

Now, it is often true that processor utilization and memory utilization can be tightly related. For example, if a computer has so little memory that the operating system is constantly required to move data into and out of the swap file, this tends to keep the processor utilization higher than it would be otherwise. However, since disk is so much slower than memory, it's rare that such a situation can cause a processor to be 100% utilized.

Also, you should note that physical memory being 99% utilized is not, by itself, an indication of memory depletion. In the general case, Windows will try to cache system files and system data into memory, if memory is available to do so. This allows Windows to access those items much more quickly, since memory is so much faster than disk.

If Figure 2, you can see another picture of that same computer, just a short while later.

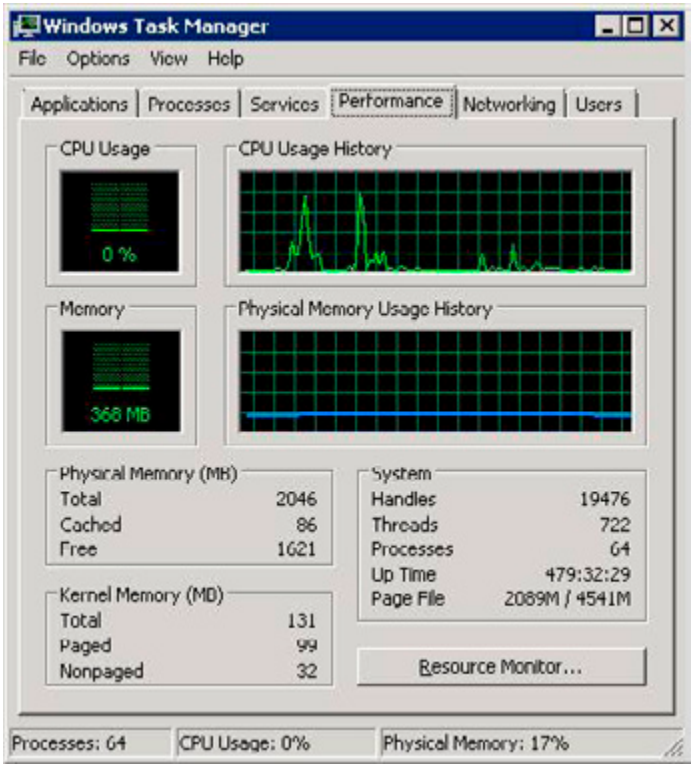

Figure 2: The Second Sample Windows Task Manager Dialog.

In Figure 2, you see a processor that is fairly idle with no memory pressure. The total physical memory has not changed, but in this second performance snapshot, the amount of free memory is up to 1,621 MB and the page file is only 46% in use.

Based on what we see in Task Manager, this computer should be performing very well, shouldn't it?

Of course, that is a trick question. In the first paragraph of this article I mentioned three key contributors to performance and so far, I've only discussed two. The third? Disk subsystem performance. Unfortunately, there is no sure-fire, single-graph view that will absolutely identify whether your disk subsystem is performing well or not. You have to do some investigation. That investigation is why some folks consider disk subsystem performance management something that is "more art than science." However, I would disagree. Diagnosing performance issues is actually very direct. Now, predicting end-user disk subsystem utilization - that is a black art!

LogicalDisk vs. PhysicalDisk

A key component of understanding how to measure the performance of an Exchange disk subsystem - or any disk subsystem - is to understand the difference between the LogicalDisk and PhysicalDisk performance objects. This is even more important than normal since these object may, or may not, measure the same things.

The easiest way I've learned to explain them is below. Stay with me through the entire four parts of the explanation to find your "Ah-hah!" moment.

- LogicalDisk – A logical disk is the unit of a disk subsystem with which Windows and users utilize a disk. When you open "Computer" (or "My Computer" for Windows 2003 and older versions of Windows) the hard disk drives shown there are logical disks.
- PhysicalDisk – A physical disk is the unit of a disk subsystem which the hardware presents to Windows.
- A logical disk may consist of multiple physical disks (think of RAID).
- A physical disk may host multiple logical disks (think Windows partitions).

If you put all of these together, this means that in the case where a physical disk contains only a single Windows volume, LogicalDisk and PhysicalDisk measure the same thing.

> ***Note***
>
> *The highest performance disk devices currently available (Fiber Channel dual-controller disks) provide about 180 input-output operations per second (IOPS). They are quite expensive. More common enterprise class disk devices provide about 150 IOPS (for a 15K RPM SAS disk). Workstation class disk devices have much lower performance, often around only 35–50 IOPS.*

Somewhat confusingly, disk aggregators (this includes RAID controllers, Storage Area Networks, Network Attached Storage, iSCSI, etc.) may present many physical disks as a single logical device to Windows. However, each of these devices (known as a logical unit number or LUN) may *again* actually represent multiple logical or physical disks. Thankfully, from a Windows performance perspective, those distinctions can be ignored, at least until a specific LUN is identified as having a performance issue. In that case, in order to acquire more specific data, you will have to use performance tools from the aggregator's provider as the disk aggregators can sometimes provide unpredictable results.

Is there a conclusion here? Yes, there is. The conclusion is that in the most common cases, the performance of a LogicalDisk object is in what you are most interested.

> ***Note***
>
> *Lots of SAN and NAS software provides a feature called "LUN stacking" or "disk stacking" which allows multiple LUNS to exist on a single physical disk. This just complicates your life. Avoid it. Just always remember that you have to be able to identify what you are measuring and the boundaries on that measurement. If you have multiple applications accessing a single physical disk, then your performance will always be non-deterministic and difficult to predict.*

What do I measure?

Now that you know which performance object is interesting, what do you do with it?

As you probably know, any performance object is composed of a number of individual *performance counters*. A counter contains a discrete value which may be any of:

- A value that has been increasing since the last time the counter was cleared (which normally happens at system reboot).
- A value that represents the delta (change) since the last measurement.
- A value that represents the percentage change since the last measurement.
- A value that represents an instantaneous measurement of some item.
- A value that is a percentage against some absolute (such as a time interval).

For the LogicalDisk performance object, there are 23 counters:

1. % Disk Read Time – the percent of wall-clock time spent processing read requests since the last sample.
2. % Disk Time – % Disk Read Time plus % Disk Write Time.
3. % Disk Write Time – the percent of wall-clock time spent processing write requests since the last sample.
4. % Free Space – the amount of unused space on the disk, expressed as a percentage of the total amount of space available on the disk.
5. % Idle Time – the percent of time spent idle since the last sample (that is, processing neither read nor write requests).
6. Avg. Disk Bytes/Read – the average number of bytes transferred from the disk in each read operation since the last sample.
7. Avg. Disk Bytes/Transfer – the average number of bytes transferred from or to the disk in each I/O operation since the last sample.
8. Avg. Disk Bytes/Write – the average number of bytes transferred from the disk in each write operation since the last sample.
9. Avg. Disk Queue Length – the average numbers of I/O requests (both read and write) queued for the selected disk since the last sample.
10. Avg. Disk Read Queue Length – the average number of read requests queued for the selected disk since the last sample.
11. Avg. Disk sec/Read – the average amount of time that it took for a read request to complete for the selected disk since the last sample.
12. Avg. Disk sec/Transfer – the average amount of time that it took for any I/O request to complete for the selected disk since the last sample.
13. Avg. Disk sec/Write – the average amount of time that it took for a write request to complete for the selected disk since the last sample.
14. Avg. Disk Write Queue Length – the average number of write requests queued for the selected disk since the last sample.
15. Current Disk Queue Length – the current number of I/O requests queued for the selected disk.
16. Disk Bytes/sec – the average rate, or speed, that bytes are transferred from or to the selected disk during any I/O operation.

17. Disk Read Bytes/sec – the average rate, or speed, that bytes are transferred from the selected disk during a read operation.

18. Disk Reads/sec – the average number of read requests that occur per second for the selected disk.

19. Disk Transfers/sec – the average number of I/O requests that occur per second for the selected disk.

20. Disk Write Bytes/sec – the average rate, or speed, that bytes are transferred to the selected disk during a write operation.

21. Disk Write/sec – the average number of write requests that occur per second for the selected disk.

22. Free Megabytes – the amount of unused space on the disk, expressed in megabytes.

23. Split IO/sec – if a read or write operation is too large to be satisfied in a single I/O operation, the I/O is split into multiple separate physical I/Os; this counter records how often that happens for the selected disk.

Speaking as a person who likes math, I think that all of these counters are interesting, and are worthwhile tracking over time. For example, if the *% Free Space* is decreasing on a regular basis, then you may need to plan to add more disks – or stop the expansion in some way. If your *Disk Write/sec* is trending upward, you may need to investigate why more write operations are occurring to that disk drive.

What Defines Good Performance?

For determining whether you have adequate performance, I would suggest that there are four main counters that you need to constantly monitor. Note that in this case, I'm suggesting that the counters should be checked every one to five minutes. Those counters are:

1. Average Disk Queue Length

2. Average Disk Read Queue Length

3. Average Disk Write Queue Length

4. Free Space (either % Free Space or Free Megabytes).

For [4], it should be obvious that if any write operations are occurring to a disk, if it fills up, performance will rapidly approach zero, the event log will starting filling with errors, and users will be calling your phone saying "Exchange is down" (or name your favorite application instead of Exchange). On the other hand, if a disk is read-only, such as a DVD-ROM database, there is no need to measure either [3] or [4]. Writes cannot occur. But that is a special case, and since we are primarily concerned with Exchange, it doesn't support any type of read-only volume.

You may find it interesting that I use the average values instead of the current values. That is because any system can be instantaneously overwhelmed. At any given moment of time, you may have 50 I/O's queued to a particular disk waiting to be satisfied; but that is unlikely to be the normal case. If that high queue value becomes the normal case, then the average values will trend quite high as well. That will server to provide you with an absolute indication of an issue which needs to be addressed.

If the average values are not high, then you had what is called a "usage spike" in which a device was temporarily overwhelmed. And that is why we have queues anyway

What is a "high value"? A high value occurs when the device cannot empty the queue. Because each I/O is variable in size, its arrival cannot be predicted, and the I/O takes a variable amount of time to complete; the point at which a device is considered saturated or overburdened with I/O requests actually occurs a little earlier than you may presume.

Now, I am not a statistician, nor do I play one on TV. However, like many of you, I took basic queuing theory in college. A device becomes saturated when its I/O queue exceeds 70% on an average basis. To put that another way: **if the average disk queue length of a device exceeds 0.70 then the performance of that disk is negatively impacting user performance**.

The overall *Average Disk Queue Length* is an average of the read and write average queue lengths. The read and write queues are maintained independently. A disk may (and in fact, usually will) have significantly different read and write profiles. While checking the Average Disk Queue Length is an important indicator of disk performance, if it indicates heavy usage, your first order of business is to determine whether read performance or write performance (or both) are causing the performance issue. Therefore my recommendation to continuously monitor all three counters.

So now you know: the best way to check your Exchange disk performance is to:

- Make sure there is room on the disk, and
- Monitor the average disk queue lengths.

The next thing you may want to know is – how can I predict how my disk subsystem **should** perform? But that is a topic for another article.

Upgrading to Exchange Server 2007: Part 2

26 March 2009

by Jaap Wesselius

Jaap completes his series on the necessary steps to migrate an existing installation of Exchange to Exchange Server 2007. He deals with the problems of installing the Hub Transport server, moving policies and mailboxes and decommissioning the old servers.

Upgrade to Exchange Server 2007 – Part II

In my previous article I explained how to upgrade an existing Exchange Server 2003 environment to a new Exchange Server 2007 environment. Two new servers were introduced, a combined Client Access Server / Hub Transport Server and a dedicated Mailbox server. The Client Access Server is supplied with a new Subject Alternate Name (SAN) certificate. Since the Exchange Server 2007 Mailbox server is in the same Exchange organization this server has a Public Folder database next to the Mailbox database. After introduction of the Exchange Server 2007 servers the environment will look like this:

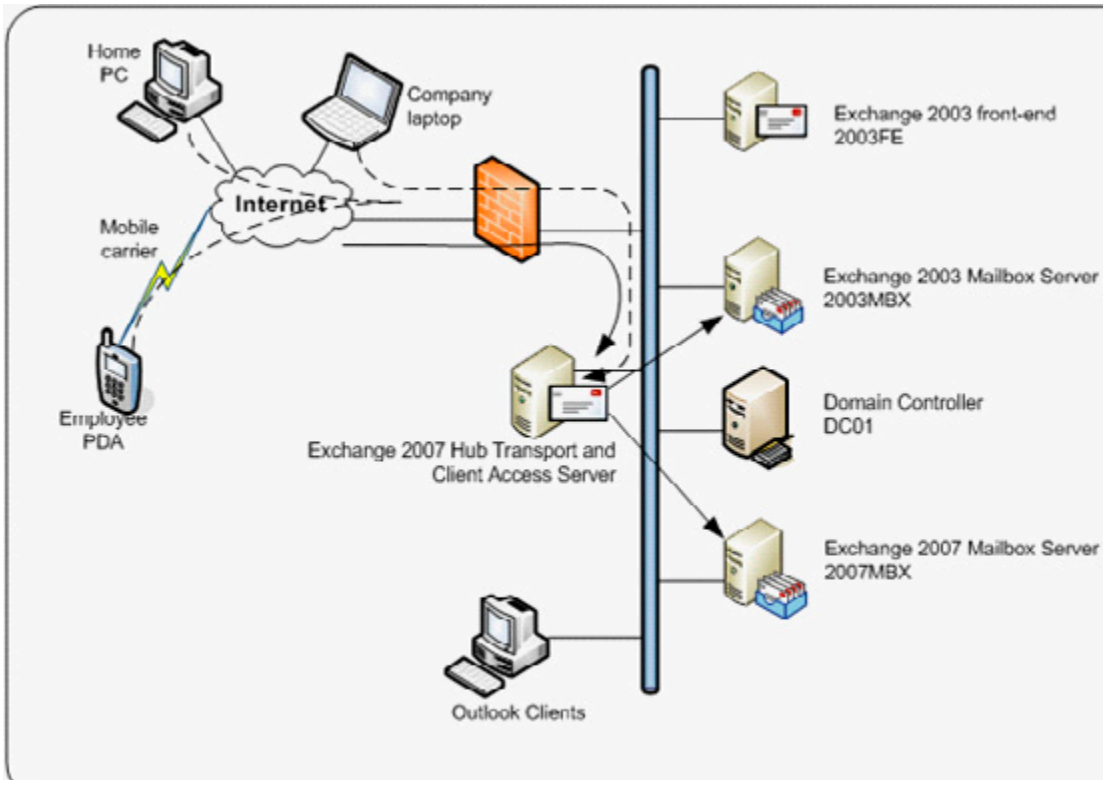

Figure 1. A combined Exchange Server 2003 / Exchange Server 2007 environment.

Since we have introduced Exchange Server 2007 into the environment and it is fully up and running, it's time to start moving resources from Exchange Server 2003 to Exchange Server 2007.

This includes:

- Changing the Internet e-mail acceptance.
- Changing the Internet facing Client Access.
- Replicate Public Folders and System Folders to Exchange Server 2007.
- Moving the Recipient policies.
- Moving Mailboxes.
- Remove the Exchange Server 2003 servers.

Changing the Internet e-mail acceptance

During the installation of the Exchange Server 2007 Hub Transport server a Routing Group connector is automatically created between the Hub Transport Server and the Exchange Server 2003 Front-End Server. Messages between both versions of Exchange server flow via the Routing Group Connector. Before the inbound SMTP flow (from the Internet) can be changed there are two configuration issues that need to be solved. Please note that the Routing Group Connector is only created when the Hub Transport Server is installed using the GUI. If the Hub Transport Server is installed using the command prompt the Routing Group Connector is *not* created automatically.

By default the Hub Transport Server does not accept anonymous connections. This means that inbound SMTP connections from the Internet are not accepted. This can be changed by adding "Anonymous users" to the Permissions Groups of the default receive connector on the Hub Transport Server.

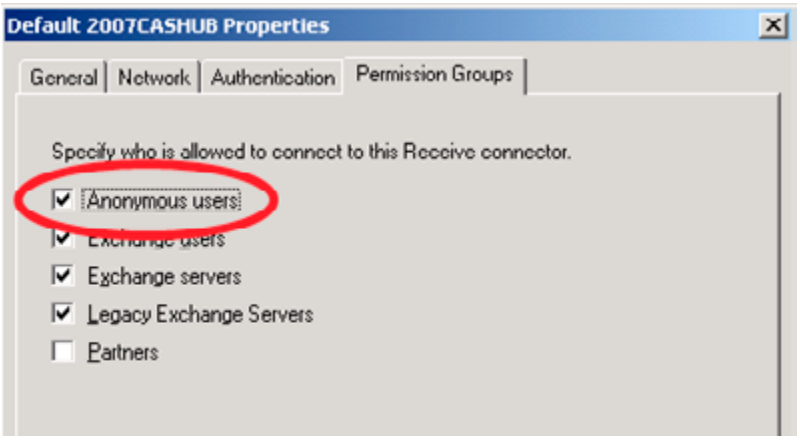

Figure 2. Add the anonymous users to the Permission Groups.

The reasons for this is in the potential implementation of an Exchange Server 2007 Edge Transport server. This server, that's typically installed in the network's Demilitarized Zone (DMZ) does accept anonymous connections by default. The connection between the Edge Transport and the Hub Transport server is not anonymous, but of the type "Exchange Servers."

Depending on your companies anti-spam policies and solution you might want to enable the anti-spam services on the Hub Transport Server. Microsoft supplies a script that enables all anti-spam functionality. The scripts is called *"install-AntiSpamAgents.ps1"* and is located in C:\Program Files\Microsoft\Exchange Server\Scripts. Open an Exchange Management Shell command window, go to this directory and execute the script to enable the anti-spam agents on this Hub Transport Server.

Note

The default installation is in C:\Program Files\Microsoft\Exchange Server Scripts. It is possible to install Exchange Server 2007 on another location. It is also possible to use a default variable called $exscripts, just open the Exchange Management Shell en type CD $Exscripts. Do not forget to restart the Transport Server after running the Install-AntiSpamAgents.ps1 script.

When opening the Exchange Management Console and selecting the Hub Transport under the Organization Configuration, a new tab will appear in the results pane with Anti-Spam. Here you can change the anti-spam settings like Content Filtering, IP Allow list, IP Block list, Recipient Filtering, etc.

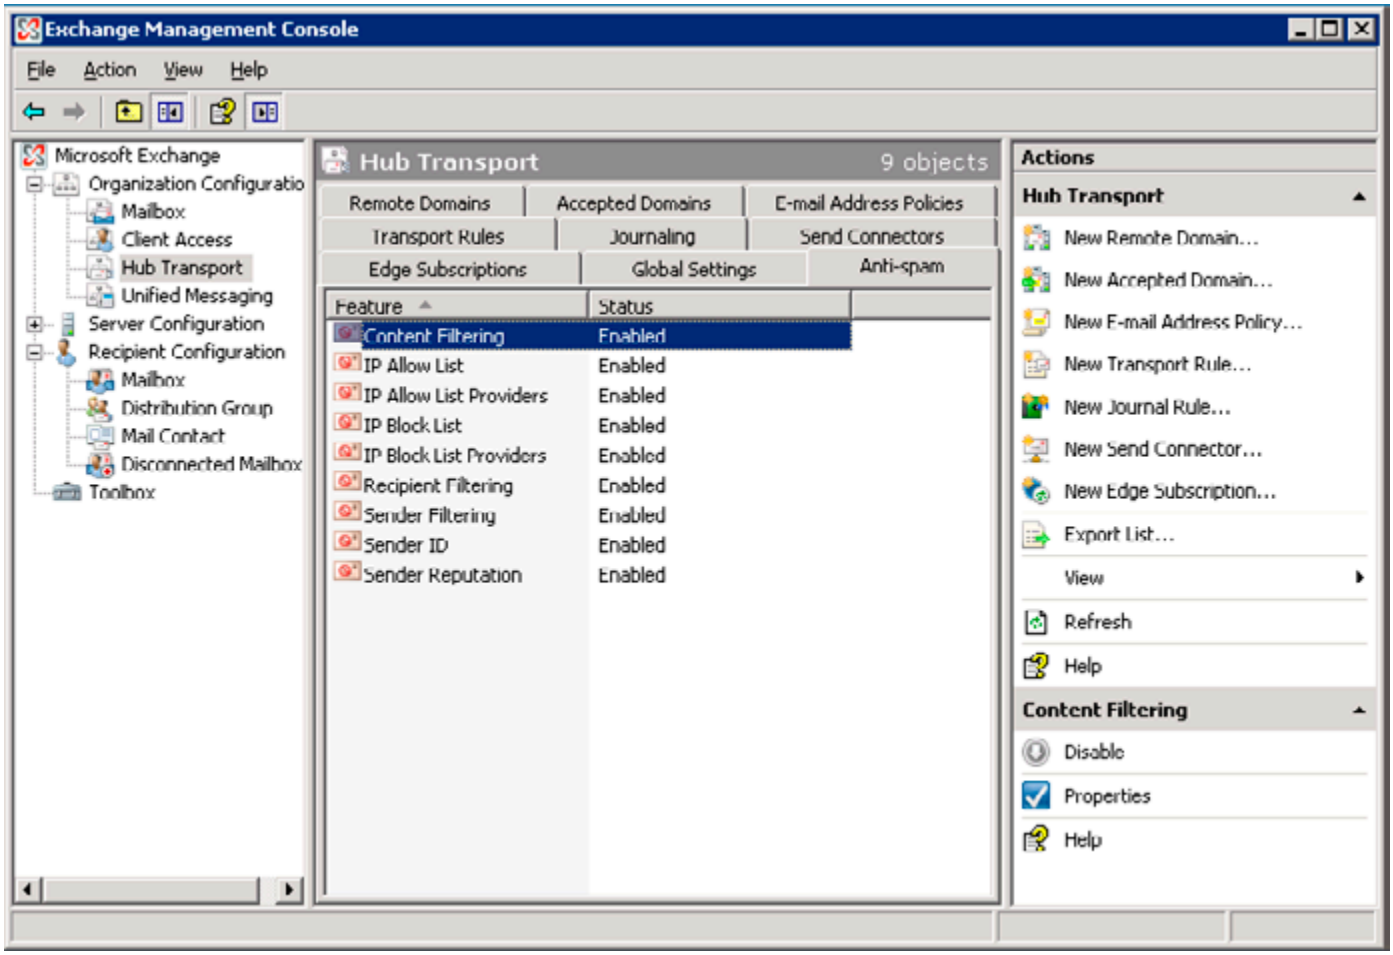

Figure 3. The new anti-spam tab after enabling the anti-spam options.

After making these changes the inbound SMTP mail flow can be changed. Depending on the infrastructure, you have to change the IP number of the MX record on the public DNS or change the port forwarding on your firewall. In either case you have to change it to the IP number of the Exchange Server 2007 Hub Transport Server. After the change you can check new messages in Outlook by looking at the Internet message header. It will show clearly the name and IP number of the new Hub Transport Server.

By default Exchange Server 2007 does not allow internal clients such as web servers, printers or applications to relay SMTP traffic. If you want to implement an anonymous connector for relaying SMTP traffic please visit the following Microsoft article: http://technet.microsoft.com/en-us/library/bb232021.aspx.

Changing the Internet facing Client Access

You always have to remember that an Exchange Server 2003 front-end server **is not** compatible with an Exchange Server 2007 mailbox server. This means that you have to move to the Exchange Server 2007 Client Access Server **before** you move any mailboxes to Exchange Server 2007.

In the previous article a new certificate for the Client Access Server was already requested and installed. The certificate that needs to be used is an "Unified Communications" Certificate, also known as a SAN (Subject Alternate Name) certificate or an Exchange 2007 Certificate. A certificate like this can contain multiple Fully Qualified Domain Names.

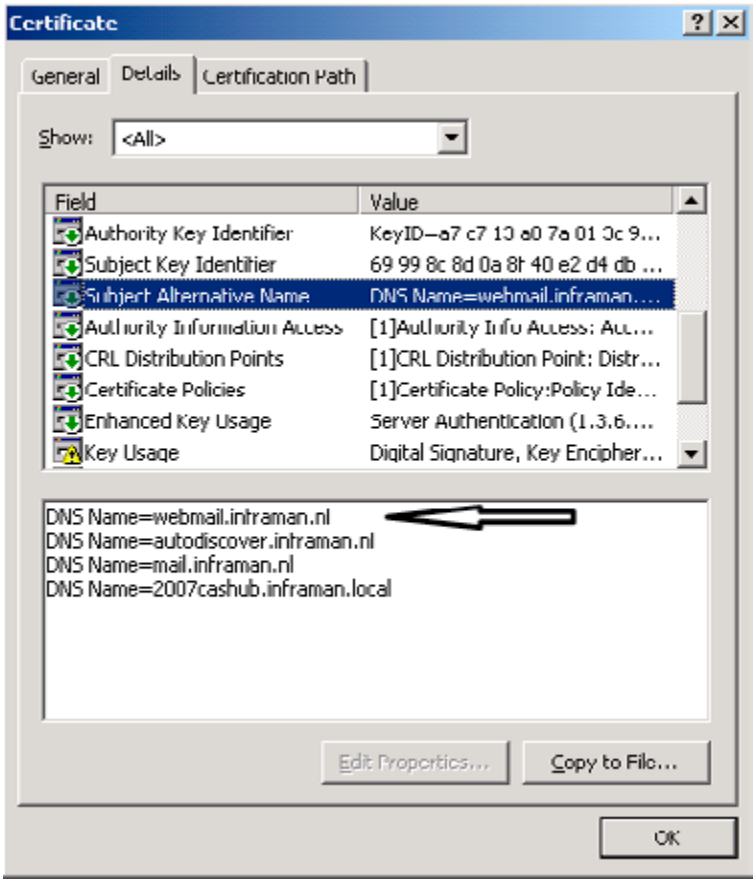

Figure 4. A Unified Communications Certificate from Digicert containing multiple names.

The certificate used for this environment is from Digicert (www.digicert.com) and is supported by most of the clients like Windows (Internet Explorer), Windows Mobile (for ActiveSync) and Outlook (for Outlook Anywhere).

Now that the certificate is right for usage with Exchange Server 2007 the following services have to be configured correctly on the Exchange Server 2007 Client Access Server:

- Outlook Web Access.
- ActiveSync.
- Outlook Anywhere.

To configure the client setting navigate to Client Access under the Servers Configuration in the console tree in the Exchange Management Console. When you open the properties of the OWA virtual directory on the Client Access Server, on the General tab you can set

the Internal URL as well as the External URL. On the Authentication tab you can select how users have to logon. Using forms based authentication you can choose between Domain \ UserName, User Principal Name (for example jaapw@inframan.nl) or using the UserName only. In this case you have to enter the logon domain.

The next step is to edit the properties of the Microsoft-Server-ActiveSync virtual directory. This virtual directory is used by Windows Mobile clients for accessing the e-mail information. Like the OWA Virtual directory you have to enter both the Internal URL as well as the External URL. The External URL, in this example https://webmail.inframan.nl/Microsoft-Server-ActiveSync should have a valid certificate that is recognized by the PDA.

There is an issue though with the certificates, even with real certificates. When requesting the certificate using the **New-ExchangeCertificate** commandlet you have to enter the domain name ("webmail.inframan.nl") in the **–SubjectName** option, but you also have to use it in the –DomainName option. This will result in the FQDN in both the SubjectName as well as the Subject Alternative Names (see Figure 3) and this will cost you an extra credit when requesting the certificate with your provider. It should also be the first name in the Subject Alternative Names field. If you fail to do so the Windows Mobile devices will contact the Client Access Server but will not recognize the certificate. This will lead to error 0x80072f06 on the Windows Mobile device.

Note

If you have an interim situation where PDA's access the Exchange Server 2007 Client Access Server but where the user's mailbox is still on Exchange Server 2003, then you have to change the authentication on the Microsoft-Server-ActiveSync Virtual Directory on the Exchange Server 2003 mailbox server. By default only "Basic Authentication" is selected on the Virtual Directory, but in the interim scenario you also have to select "Windows Integrated" authentication.

Outlook Anywhere, previously known as RPC over HTTP has to be enabled on the Exchange Server 2007 Client Access Server. In the actions pane (on the right side) of the Exchange Management Console you can select "Enable Outlook Anywhere," the button to enable this is somewhat hidden. Enter the external hostname (i.e. webmail.inframan.nl) and select the method of authentication. You have to remember this setting in case you want to manually configure the Outlook 2003 or 2007 clients. In the Outlook profile you also have to select the method of authentication. If these setting don't match the client's settings the client will not be authenticated and the user will not be able to connect to his or her mailbox.

If you have changed all the settings mentioned above according to your own situation and company policy you can change the firewall settings so that external clients can access the new Exchange Server 2007 Client Access Server. All requests will now be accepted by the Exchange Server 2007 Client Access Server and be forwarded to the old Exchange Server 2003 mailbox server successfully.

Replicate Public Folders System Folders to Exchange Server 2007

The next step in the migration process is to move the Public Folders from Exchange Server 2003 to Exchange Server 2007. If you are using only the System Folders in the Public Folders, i.e. the Free/Busy information and the Offline Address Book distribution you can follow the standard Microsoft approach by replication these folders from Exchange Server 2003 to Exchange Server 2007

The Exchange Server 2003 Public Folder Database has to replicate its information to the Exchange Server 2007 Public Folder Database and vice versa. In a coexistence scenario where mailboxes reside both in Exchange Server 2003 as well as in Exchange Server 2007 everybody will be able to see each others free/busy information.

To start replicating the Free/Busy information from Exchange Server 2003 logon to this server using the administrator credentials and open the Exchange System Manager. Navigate to the Public Folder Database and under Public Folders open the properties of the "Schedule+ Free Busy Information." Select the Replication tab and using the Add button add the Exchange 2007 Public Folder Database to the replication list. Repeat the same steps for the Offline Address Book folders to Add the Exchange 2007 Public Folder Database to the replication list.

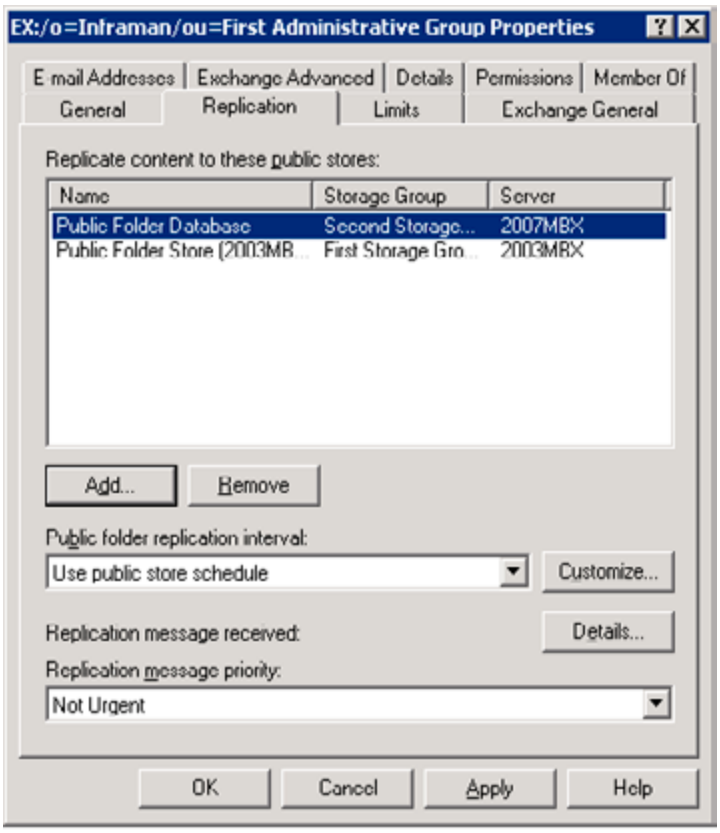

Figure 5. On the Exchange 2003 server add the Exchange Server 2007 Public Folder database to the replication list.

The Public Folder information has to be replicated from Exchange Server 2007 to the Exchange Server 2003 Public Folder database as well. To do this logon to the Exchange Server 2007 using the administrator credentials and open the Exchange Management Console. In the console tree select the Toolbox and in the results pane select the Public Folder Management Console. If no default server is selected then select the Exchange 2007 Mailbox server as the default server.

In the Public Folders tree navigate to the System Public Folders and navigate down the tree to the Schedule+ Free Busy Information. Open the properties of this folder, select the Replication tab and add the Exchange 2003 Public Folder to the replication list.

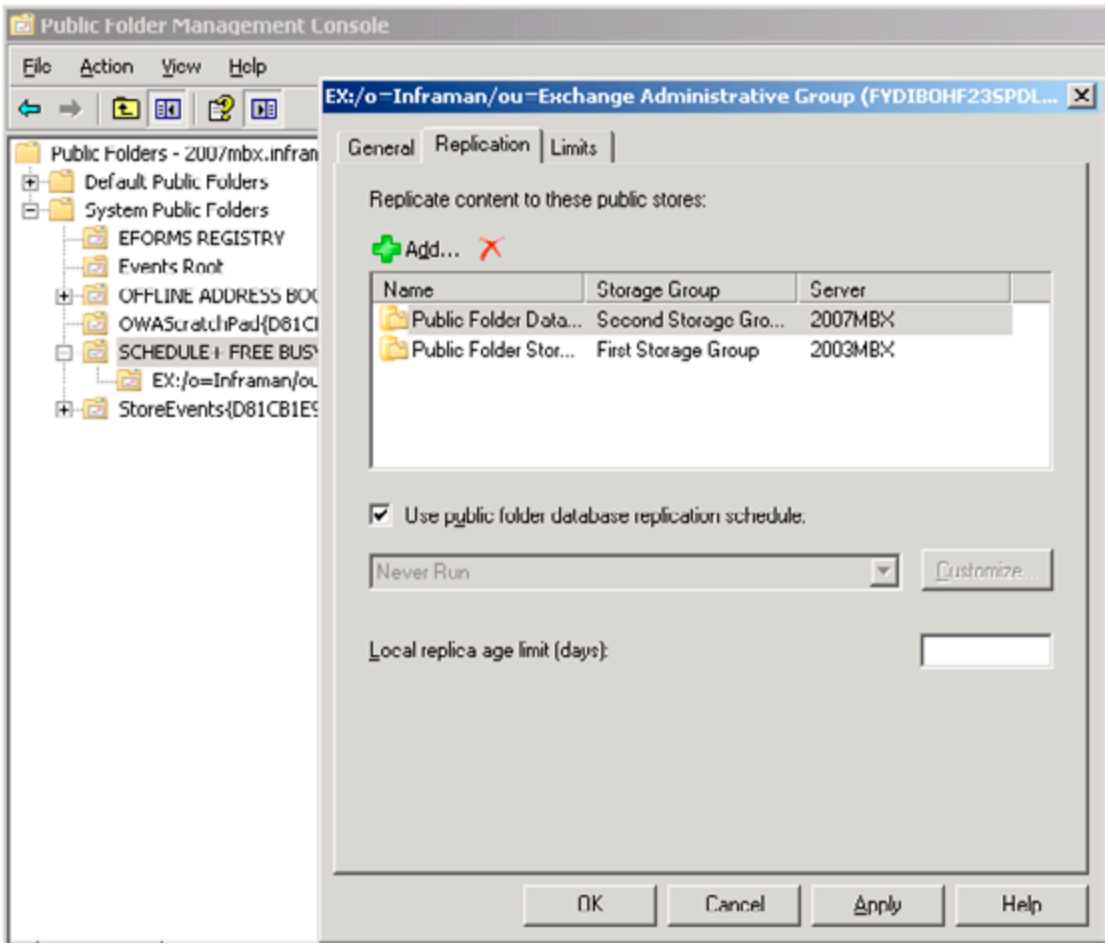

Figure 6. Add the Exchange 2003 Public Folder Database to the replication list.

During the replication you will see that the free/busy information from the "other" server will appear in the System Public Folder list. This can take some time to complete though. To improve the replication speed change the Public Folder replication interval to "always" and the "replication message priority" to high.

If you are using "normal" Public Folders you have to replicate the Public Folder hierarchy from Exchange Server 2003 to Exchange Server 2007 as well as the actual content of the Public Folders that your users are using. To achieve this you can use the Public Folder Migration Tool which is available from the Microsoft download site. You can use the Migration Tool to replicate Public Folders from Exchange Server 2003 to Exchange Server 2007. For more information regarding the Public Folder Migration tool please visit the following Microsoft knowledgebase article: http://support.microsoft.com/kb/822895/en-us.
The latest version of the Deployment Tools is available on the Microsoft download site:
Microsoft Exchange Server Deployment Tools.

Move the Recipient Policies

In Exchange Server 2003 there's a service called "Recipient Update Service" or RUS that's responsible for setting Exchange related information on users when they are mailbox enabled. In Exchange Server 2007 the RUS no longer exists and its functionality is now delivered by a "E-mail Address Policy." This policy is responsible for setting E-mail addresses on Exchange recipients like mailboxes. The Recipient Update Service needs to be upgraded to an Email Address Policy. Only after an upgrade you are able to manage this using the Exchange Management Console and if not upgraded an alert will be shown when trying to open an E-mail Address Policy using the Exchange Management Console:

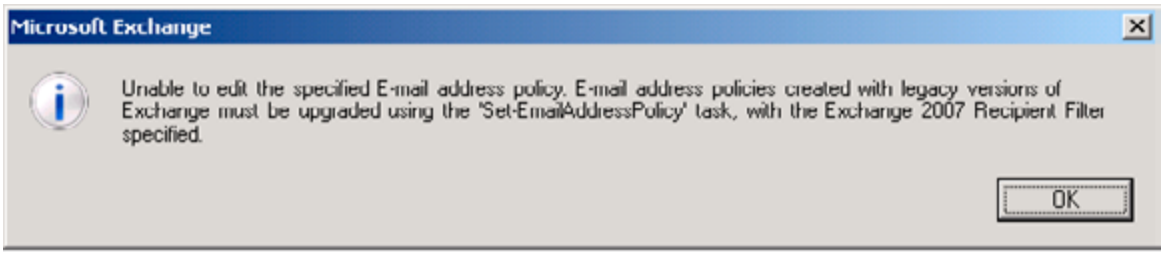

.Figure 7. Error message when a Recipient Policy is opened on an Exchange Server 2007 Management Console.

Before you continue with changing the Recipient Policies it's a good time to run the Exchange Best Practices Analyzer to see if your environment is in good shape. The following steps are one way only, there's no way back so you have to make sure everything is running fine!

You have to change all Recipient policies using the Exchange Management Shell. This is one example of an action that cannot be performed with the Exchange Management Console. You can retrieve a list of all Recipient Policies in your Exchange environment using this command in an Exchange Management Shell window:

```
Get-EmailAddressPolicy | where { $_.RecipientFilterType -eq "Legacy" }
```

This output can be used as input for the Set-EmailAddressPolicy commandlet using the PIPE functionality in the Exchange Management Shell. This results in the following command:

```
Get-EmailAddressPolicy | where { $_.RecipientFilterType -eq "Legacy" } | Set-EmailAddressPolicy –
IncludedRecipients AllRecipients
```

The last things that need to be converted are the Address Lists. Exchange Server 2003 Address Lists cannot be managed by the Exchange Server 2007 Management Console and vice versa. Like the Recipient Policies the Address Lists can only be converted using the Exchange Management Shell.

In an Exchange Management Shell windows enter the following commands:

```
Set-AddressList "All Users" -IncludedRecipients MailboxUsers
Set-AddressList "All Groups" -IncludedRecipients MailGroups
Set-AddressList "All Contacts" -IncludedRecipients MailContacts
Set-AddressList "Public Folders" -RecipientFilter { RecipientType -eq 'PublicFolder' }
Set-GlobalAddressList "Default Global Address List" -RecipientFilter {(Alias -ne $null -and
(ObjectClass -eq 'user' -or ObjectClass -eq 'contact' -or ObjectClass -eq 'msExchSystemMailbox'
-or ObjectClass -eq 'msExchDynamicDistributionList' -or ObjectClass -eq 'group' -or ObjectClass
-eq 'publicFolder'))}
```

For each command a confirmation needs to be given and the object will be upgraded. After the upgrade it won't be possible anymore to manage the Address Lists from an Exchange Server 2003 System Manager.

For more detailed information regarding the upgrade of Recipient Policies and Address Lists visit the blog of the Microsoft Exchange product team: http://msexchangeteam.com/archive/2007/01/11/432158.aspx - Address List and EAP filter upgrades with Exchange Server 2007.

Moving Mailboxes from 2003 to 2007

Now that everything is in place and working correctly it's time to start moving mailboxes from Exchange Server 2003 to Exchange Server 2007. This can only be done on the Exchange Server 2007 side using the Exchange Management Console or the Exchange Management Shell. The latter one can be very useful if you want to write scripts to move mailboxes from Exchange Server 2003 to Exchange Server 2007 in bulk.

In the Exchange Server 2007 Management Console select the mailboxes you want to move, right click your selection and choose "Move Mailbox." After selecting the appropriate Mailbox Database (if you have multiple) the Move Mailbox process starts. This is usually not a very fast process and for large mailboxes it can take a serious amount of time. Currently I'm working on a project where we are moving approximately 25.000 mailboxes with around 12TB of data. This will take weeks and weeks to finish.

It is also possible to move mailboxes using the Exchange Management Shell. With the Shell it is possible to make custom queries and use this output as the input for the actual move-mailbox commandlet. To move my own mailbox the following command is used:

```
Get-Mailbox -Identity "Jaap Wesselius" | Move-Mailbox -TargetDatabase "Mailbox Database"
```

This will select my mailbox (on the Exchange Server 2003 mailbox server) and move it to the Mailbox Database on the Exchange Server 2007 mailbox server.

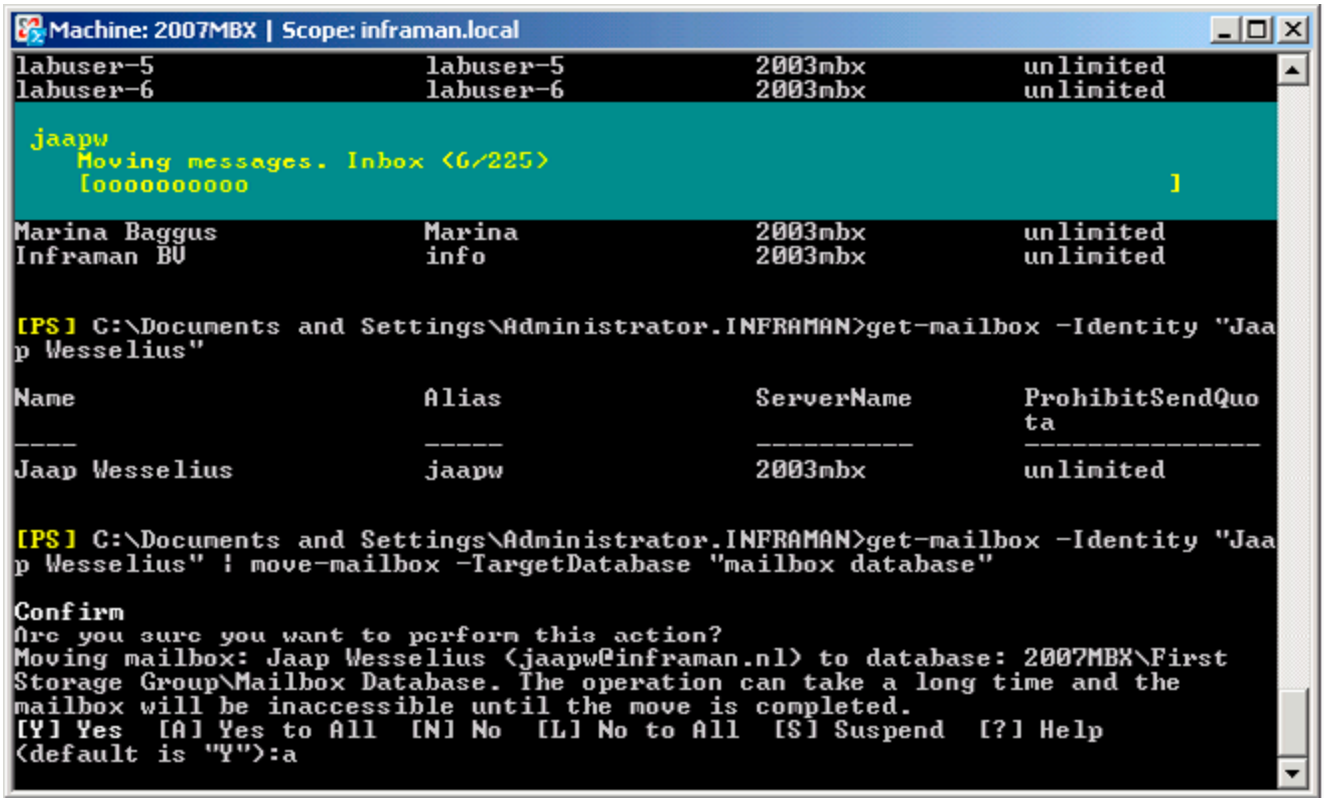

Figure 8. Moving my own mailbox using the Exchange Management Shell to Exchange Server 2007.

The last this that has to be moved to the new Exchange Server 2007 environment is the Offline Address Book generation. To do this open the Exchange Management Console on an Exchange Server 2007 server and navigate to the Mailbox section in the Organization Configuration. Select the Offline Address Book tab and right click on the Default Offline Address Book. The second option gives the possibility to select a new Offline Address Book generation server, which should be the Exchange Server 2007 mailbox server.

Note

For several steps in this article you have to wait for replication to finish. This can be for Active Directory replication as well as for Public Folder replication. It is also possible, like the Offline Address Book generation that a process occurs only once a day, for example in the middle of the night. If you're too fast with the several steps you can miss some actions which can result in erratic behavior.

Removing the Exchange Server 2003 servers

Before decommissioning the Exchange Server 2003 mailbox server the Public Folders have to be moved. Logon to the Exchange Server 2003 server and start the Exchange System Manager. Select the Public Folder database on this server and right click the object. Select "move all replicas" to have everything replicated to an Exchange Server 2007 Public Folder database in your Exchange Server organization. This can take several hours to complete! If not complete you are not able to remove the Public Folder database from the server.

The Public Folder tree itself should also be moved to the new Exchange Server 2007 Public Folder database. Logon to the Exchange Server 2003 server and open the Exchange Service Manager. Expand the Administrative Groups and right click the "Exchange Administrative Group (FYDIBOHF23SPDLT)," select "New" and select "Public Folders Container."

Then expand the old "First Administrative Group," expand "Folders" and move the Public Folders tree to the Public Folders container you created in the previous step.

The next step in our process it to remove the Routing Group Connector that connect both Exchange versions. This can only be done after the Public Folders are removed from Exchange Server 2003 since the replication process uses this connector! Make sure that the queues for this connector are empty so no messages gets lost.

You can remove the Routing Group Connector either with the Exchange Server Manager on Exchange Server 2003 or with the Exchange Management Shell (Remove-RoutingGroupConnector) on Exchange Server 2007. Removing the Routing Group Connector is also an example of something that can only be achieved using the Exchange Management Shell! Since all protocols were already targeted towards the Hub Transport Server or the Client Access Server we can uninstall the Exchange Server 2003 front-end server. Go to the Control Panel, select Add/Remove Programs and remove Exchange Server 2003. Please note that you need the installation media to finish the removal of the Exchange Server!

The Recipient Update Service is the next to remove from the Exchange Server 2003 server. Open the Exchange System Manager and in the Recipients Container select the Recipients Update Service (domain). Right click this Recipient Update Service and select "Delete." To remove the Enterprise Recipient Update Service it's not possible to use the Exchange System Manager. To remove this you have to use ADSIEdit.

Open ADSIEdit and open the Configuration Container in Active Directory. Navigate to the "`CN=Recipient Update Services,CN=Address Lists Container,CN=Inframan,CN=Microsoft Exchange,CN=Services,CN=Configuration,DC=inframan,DC=local`" container. There you'll find the Recipient Update Server (Enterprise) object. Right click this object and select "Delete."

> **Note**
>
> *Using ADSIEdit can be very disastrous if not being used properly. You can move the above mentioned object temporarily to the Exchange Server 2007 server until you're finished with the complete process and then deleted. If something goes wrong and you have it deleted you're in trouble.*

The Exchange Server 2003 mailbox server is now the last Exchange 2003 server in our organization and ready to be removed. Open Control Panel, select "Add/Remove Programs" and remove Exchange Server 2003 from this server. Again make sure that you have the installation media available during the remove process.

The last step is to remove the Write DACL right for the Exchange Servers group should be removed from the root of the domain. This can be achieved by running the following command on an Exchange Server 2007 management shell:

```
Remove-ADPermission "dc=Inframan,dc=local" –User "Inframan\Exchange Enterprise Servers" –
AccessRights WriteDACL -InheritedObjectType Group
```

All legacy Exchange Domain Servers and Exchange Enterprise Servers security groups can now be deleted from Active Directory. Please make sure that they are empty and not in use for other purposes!

Note

When you check Active Directory with ADSIEdit you'll notice that the old Exchange Server 2003 Administrative Group is still available, although empty. Do not remove this Administrative Group unless you're absolutely sure there's no object in Active Directory referencing this Administrative Group in the ExchangeLegacyDN attribute. For more information please check this Microsoft knowledgebase article: http://support.microsoft.com/kb/945602. Users who use Outlook 2003 cannot publish their free/busy data in Exchange Server 2007. My personal opinion would be just to leave it there and not touch it. Nobody will see this Administrative Group and will bother nothing else, don't touch it.

More information regarding the removal of the last legacy Exchange Server can be found on the Microsoft website: http://technet.microsoft.com/en-us/library/bb288905.aspx - How to Remove the Last Legacy Exchange Server from an Organization

Message Tracking in Exchange 2007

27 March 2009

by Ben Lye

'Where did my mail go?'. In order to answer this question, to troubleshoot mail problems and to analyse mail flow, the Exchange administrator can use message-tracking logs. Ben Lye elaborates on these essential logs and explains how you can use Powershell commands to search them for those emails that have gone adrift.

Exchange message tracking records the SMTP activity of messages being sent to and from Exchange servers running the Edge Transport or Hub Transport roles. Exchange administrators can use message tracking logs for mail flow analysis as well as troubleshooting and answering the ever-familiar "where did my mail go" question.

Configuring Message Tracking

By default message tracking is enabled on any Exchange server which has the one or more of the Edge Transport, Hub Transport, or Mailbox roles installed. The default settings are to store up to 30 days of log files in files of up to 10MB with a directory size limit of 250MB.

Message tracking settings can be retrieved using the **Get-TransportServer** cmdlet for Edge and Hub transport roles and the **Get-MailboxServer** cmdlet for Mailbox server roles.

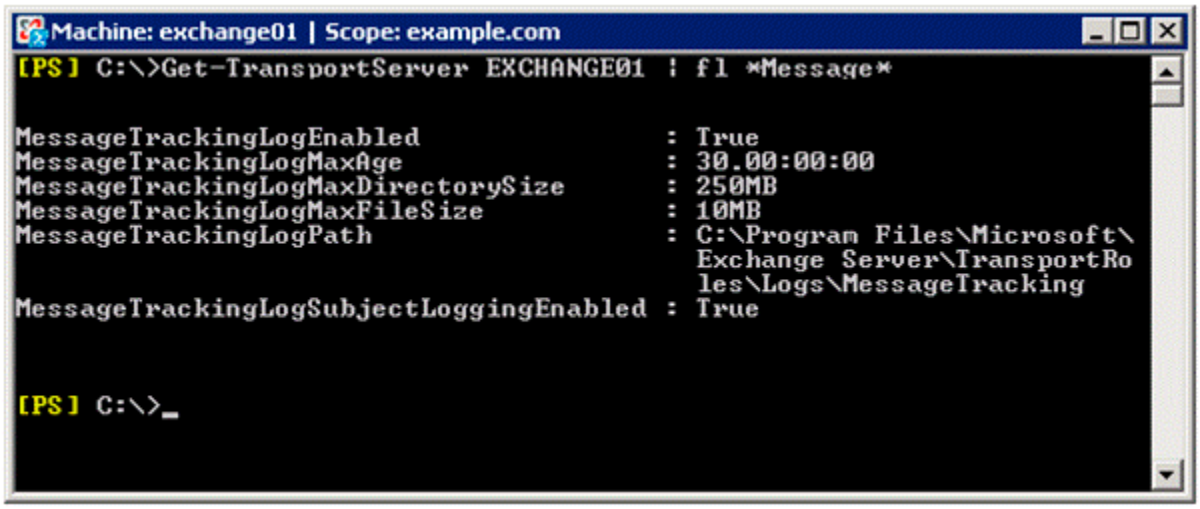

To modify the message tracking settings you can use the **Set-TransportServer** and **Set-MailboxServer** cmdlets. Using these cmdlets you can:

- Enable or disable message tracking (enabled by default)
- Enable or disable logging of message subject lines (enabled by default)
- Set the maximum age of message tracking log files (30 days by default)
- Set the maximum size of the log file directory (250MB by default)
- Set the maximum size of each log file (10MB by default)
- Change the path of the log file ("C:\Program Files\Microsoft\Exchange Server\TransportRoles\Logs\MessageTracking" by default).

If you change the path of the message tracking log directory, then new log files will be written to the new path straight away, but existing log files are not moved or copied from the old path to the new path.

Old log files are removed when either the maximum directory size has been reached, or the log file is past the maximum age. In the case of the maximum size being reached, the oldest log file is removed even though it may not have met the age limit. Because of this, if you are in a site with many users and where a lot of e-mail is sent, you may want need to increase the maximum directory size as you might find that the log files are being deleted well before the maximum age is reached.

You can use this command to increase the maximum directory size to 2GB and the maximum log file age to 90 days (adjust the values as appropriate for your environment):

```
[PS] C:\>Set-TransportServer EXCHANGE01 -MessageTrackingLogMaxDirectorySize 2GB –
MessageTrackingLogMaxAge 90.00:00:00
```

To configure Message Tracking you must be delegated the Exchange Organization Administrator role and be a member of the local Administrators group on the Exchange server.

Searching Message Tracking Logs

Once message tracking is configured, using either default or custom settings, you can use the message tracking data for testing, troubleshooting, or auditing mail flow.

Logs can be searched using with the Message Tracking Tool in the Exchange Management Console or the **Get-MessageTrackingLog** cmdlet in the Exchange Management Console. Both methods use the same set of search filters, and in fact the Message Tracking Tool uses the `Get-MessageTrackingLog` cmdlet to perform the search. `Get-MessageTrackingLog` gives the option of limiting the number of results returned, and the results can be converted into different formats.

Search results can be limited using the following filters:

Name	Description
Recipients	The complete e-mail address(es) of the message recipient(s). Multiple values can be entered using a comma delimiter.
Sender	The complete e-mail address of the message sender.
Server	The server on which to search
EventID	The specific event to search for – for example, "SEND" or "DELIVER"
MessageID	Unique ID of the e-mail message
InternalMessageID	Server-specific message ID
Subject	Subject line of the e-mail message
Reference	Additional information for some event types
Start	Starting date/time
End	Ending date/time

To perform a search using the Message Tracking Tool, launch the Exchange Management Console, navigate to the Toolbox pane, and double-click "Message Tracking." After a brief check for updates you'll be able to go to the Welcome Screen, where you can enter search parameters to begin looking for messages in the tracking logs. While you are constructing your search a box at the bottom of the tool shows you the `Get-MessageTrackingLog` command which will be used to perform the search.

To perform a search using the `Get-MessageTrackingLog` cmdlet, searching the server EXCHANGE01 for messages sent from **john@example.com** to **bill@example.net,** sent between 12/3/2009 and 13/3/2009:

```
[PS] C:\>Get-MessageTrackingLog -Server EXCHANGE01 -EventID SEND -Sender john@example.com
-Recipients bill@example.net -Start 12/3/2009 -End 13/3/2009 -ResultSize 100
```

To perform the same search and return only the first 100 matching records:

```
[PS] C:\>Get-MessageTrackingLog -Server EXCHANGE01 -EventID SEND -Sender john@example.com
-Recipients bill@example.net -Start 12/3/2009 -End 13/3/2009 -ResultSize 100
```

If you are using Exchange 2007 SP1 you must be delegated the Exchange View-Only Administrator role to use the `Get-MessageTrackingLog` cmdlet. If you are using Exchange 2007 RTM you need to be delegated the Exchange Server Administrator role and be a member of the local Administrators group on the target server.

Working With the Search Results

Once you have a search which returns the results you need, you may want to convert those results into other formats, perhaps to use for reports or to provide information to others. PowerShell includes built-in cmdlets for re-formatting output data, and those can be used in conjunction with the `Get-MessageTrackingLog` cmdlet. For the "Recipients', "RecipentStatus" and "Reference" properties it's necessary to convert the data so that it appears in the output files.

To convert the results to CSV format you can pipe the search command to the **Export-CSV** cmdlet. This command will create a CSV file called C:\Temp\SearchResults.csv, exporting all the available fields:

```
[PS] C:\>Get-MessageTrackingLog -Server EXCHANGE01 -EventID SEND -Sender john@example.
com -Recipients bill@example.net -Start 12/3/2009 -End 13/3/2009 | Select Timestamp,
ClientIp, ClientHostname, ServerIp, ServerHostname, SourceContext, ConnectorId, Source,
EventId, InternalMessageId, MessageId, {$_.Recipients}, {$_.RecipientStatus}, TotalBytes,
RecipientCount, RelatedRecipientAddress, {$_.Reference}, MessageSubject, Sender, ReturnPath,
MessageInfo | Export-CSV C:\Temp\SearchResults.csv
```

This command will create a CSV file including only the timestamp, event ID, sender, recipients, and subject line:

```
[PS] C:\>Get-MessageTrackingLog -Server EXCHANGE01 -EventID SEND -Sender john@example.com
-Recipients bill@example.net -Start 12/3/2009 -End 13/3/2009 | Select Timestamp, EventID,
Sender, {$_.Recipients}, MessageSubject | Export-CSV C:\Temp\SearchResults.csv
```

Alternatively, to convert the results to HTML you can pipe the search command to the **ConvertTo-HTML** cmdlet. Use this command to export the results to an HTML file showing the timestamp, event ID, sender, recipients, and subject line:

```
[PS] C:\>Get-MessageTrackingLog -Server EXHUB-00-UK -EventID SEND -Sender john@example.
com -Recipients bill@example.net -Start 12/3/2009 -End 13/3/2009 | ConvertTo-Html Timestamp,
EventID, Sender, {$_.Recipients}, MessageSubject | Set-Content C:\Temp\logs.html
```

Advanced Searches

PowerShell scripts can be used to do some interesting manipulation of the message tracking log data. Here are a few examples of what can be done without much effort.

Searching across multiple servers

`Get-MessageTrackingLog` only searches the message tracking logs of one server. To search the logs on multiple machines we need to use a few lines of PowerShell code.

First, get the names of all the Hub Transport servers:

```
[PS] C:\>$hubs = Get-TransportServer
```

Then pipe them into a `Get-MessageTrackingLog` command, in this case looking for all email with the subject line "Important news" sent on March 13th.

```
[PS] C:\>$hubs | Get-MessageTrackingLog -MessageSubject "Important news" -Start "13/03/2009
00:00:00" -End "13/03/2009 23:59:59"
```

This will return the message tracking information from all the hub transport servers in the Exchange organisation. As with regular message tracking log searches, it's possible to output this data to a reader-friendly HTML file.

```
[PS] C:\>$hubs | Get-MessageTrackingLog -MessageSubject "Important news" -Start "13/03/2009
00:00:00" -End "13/03/2009 23:59:59" | ConvertTo-Html ServerHostname, Timestamp, EventID,
Sender, {$_.Recipients}, MessageSubject | Set-Content C:\Temp\logs.html
```

Reporting on e-mail messages sent and received yesterday

Using PowerShell scripts it's possible to use the message tracking logs to create reports. This example will get the messages sent and received on the previous day for a group of mailboxes in a specific database.

```
# Get the start date for the tracking log search
$Start = (Get-Date -Hour 00 -Minute 00 -Second 00).AddDays(-1)
# Get the end date for the tracking log search
$End = (Get-Date -Hour 23 -Minute 59 -Second 59).AddDays(-1)
# Declare an array to store the results
$Results = @()
# Get the SEND events from the message tracking logs
$Sent = Get-MessageTrackingLog -Server EXCHANGE01 -EventID SEND -Start $Start -End $End
-resultsize unlimited
# Get the RECEIVE events the message tracking logs
$Received = Get-MessageTrackingLog -Server EXCHANGE01 -EventID RECEIVE -Start $Start -End $End
-resultsize unlimited
# Get the mailboxes we want to report on
$Mailboxes = Get-Mailbox -Database "EXCHANGE01\SG1\DB1"
# Set up the counters for the progress bar
$Total = $Mailboxes.Count
$Count = 1
# Sort the mailboxes and pipe them to a For-Each loop
$Mailboxes | Sort-Object -Property DisplayName | ForEach-Object {
# Update the progress bar
$PercentComplete = $Count / $Total * 100
Write-Progress -Activity "Message Tracking Log Search" -Status "Processing mailboxes"
-percentComplete $PercentComplete
# Declare a custom object to store the data
$Stats = "" | Select-Object Name,Sent,Received
# Get the email address for the mailbox
$Email = $_.WindowsEmailAddress.ToString()
# Set the Name property of our object to the mailbox's display name
$Stats.Name = $_.DisplayName
# Set the Sent property to the number of messages sent
$Stats.Sent = ($Sent | Where-Object { ($_.EventId -eq "SEND") -and ($_.Sender -eq $email)
}).Count
# Set the Received property to the number of messages received
$Stats.Received = ($Received | Where-Object { ($_.EventId -eq "RECEIVE") -and ($_.Recipients
-match $email) }).Count
# Add the statistics for this mailbox to our results array
$Results += $Stats
# Increment the progress bar counter
$Count += 1
}
# Output the results
$Results
```

The script works by finding all mailboxes in the DB1 database on the Exchange server EXCHANGE01, and searching the message tracking logs to find mail any RECEIVE and SEND events. The Get-Mailbox command can be easily modified to find a different group of mailboxes or changed to return distribution groups or contacts. The script could also be modified to search across multiple servers.

More information on configuring and managing message tracking and searching message tracking log files can be found on Microsoft TechNet:

http://technet.microsoft.com/en-us/library/aa997984.aspx

http://technet.microsoft.com/en-us/library/bb124375.aspx

http://technet.microsoft.com/en-us/library/bb124926.aspx

Third Party High Availability and Disaster Recovery Tools for Microsoft Exchange Server

30 March 2009

by Hilal Lone

Clustering Exchange 2007 to ensure high-availability has its disadvantages. It is expensive and complicated. Is it possible, as an alternative, to recover rapidly from a downed Exchange Server using one ot the third-party Exchange Server Disaster-Recovery tools?

The worst nightmare of any Exchange Server administrator is when the Exchange Server databases will not mount or there is a hardware failure resulting in mail flow stopping. Even seasoned Exchange Server admins break into a sweat when they hear that a database is corrupted, or that the server will not send or receive any email. I have been through this many times and I believe that most Exchange Administrators have as well. Even when we take all reasonable precautions, and do everything by the book, there is always a chance that we will get a call in the middle of the night to say that somehow the mail server is no longer accessible.

In Exchange 5.5, Exchange 2000 and Exchange 2003 days, the only hope of getting back up and online when a whole server failed was to use the backups. This was tedious, complex and had a low success rate. This method of recovery would take hours for an average sized database, during which time everyone from CEO to the guy at the reception desk will be expressing veiled (and not so veiled!) disappointment with your performance. Just yesterday you were their best buddy and now you are the target of their frustrations. So not only do you have to be up to date with the backups and follow them religiously, but also when we are attempting a recovery we are always under pressure. So it becomes very important that not only should we be confident about the recovery process, but also we need to be able to get it done with as little downtime as possible. Fortunately, now there are other options which allow us to do a soft or hard recovery without much fuss and with minimum downtime.

During the last few years, third party Exchange Server high availability and disaster recovery tools have flooded the messaging world. They provide alternatives to traditional tools and to clustering. The only hope of high availability before Exchange 2007 was to put a cluster in place, with a minimum of two nodes. And even then, the complex licensing system that Microsoft employs makes it really hard to sell the idea to management to get the required finances. Besides the investment in hardware, the required investment in software and storage was tremendous. One of my colleagues said that it was better to have a day without email than sending the clustering proposal to senior management. Not that the third party tools are cheap, but they cost considerably less and have a flat licensing system.

Limitations of Native Exchange Server Recovery Tools

Email has become the ultimate communication method in the corporate world. There have been instances when people working in adjacent cubicles email each other their lunch plans instead of just standing up and delivering the message verbally. How many times has it happened that the Information Store had to be brought down for maintenance and you had to re-schedule it again and again. I bet it never went on schedule, all because of the fact that nobody wants to be parted with their email for even a couple of hours. And how about un-scheduled downtime, say for example some firewall genius modified an access rule on the firewall and mail is no longer available to Blackberrys or to push mail devices. Naturally, mail would not be accessible via OWA as well.

How many trouble tickets and panic calls did you have to handle? I mean when this happens people always dig up the fact that this had happened before as well and I should have learnt my lesson. All those people in administration, HR, Engineering and Sales do not want the email to be down even for a few minutes. So email has to be up all the time without any break in service. If unfortunately, the mailbox store goes corrupt and now has to be restored from backup, how much downtime are we looking at? For a 100 GB store with fiber channel storage connected to it, would require at least on an average 6 hours to restore and we all know the success rate of such operation.

Let us take a look at the native tools at our disposal for Exchange Server store or server recovery. If it is just a matter of mailbox store corruption, we can use *Eseutil* or *IsInteg* to diagnose the problem and try to fix it. Even the simple process of dumping the checkpoint file (*.chk) contents would require at least 20-30 minutes. And the time is invaluable, and if a soft recovery is possible, then it would take us another 45-60 minutes to get the database back online and the mail services restored. So the limitation of Native Exchange Server recovery tools and procedures is the time they take. Even then success is not guaranteed.

Advantages and Disadvantages of Third Party Exchange Server DR tools:

Although third party tools such as CA XOSoft WanSync, Acronis Exchange Server Recovery and so on provide administrators much easier interfaces and environments to work with, they come with their own baggage as well.

While working on these technologies, it came as a surprise to me that although they all have the same objective, no two products work the same way. Each product has a different feature set with varying amount of administrative overhead.

Benefits of the Third party utilities:

- Most of them will provide automatic failover similar to what is provided by an active-passive failover cluster.
- Continuous Data Replication happens to standby server, hence reduces email loss in case of disaster.
- Freedom to switch servers as and when required.
- No hardware dependency. Standby Server does not need to match the hardware specifics of Active Server.
- Eliminates the need for streaming backups.
- User initiated or automatic backward replication after master server is back online.
- Scheduled syncs; avoid rush or peak hour sync and schedule them in non-peak hours.
- Provide Site resilience by synchronizing data over WAN.

- Near automatic mail route redirection (based on DNS, Routing Groups or SMTP connectors).
- No Client re-configuration required.
- Intuitive GUI Administration consoles.

Disadvantages of Third Party Utilities

- Not widely documented or implemented.
- Very limited resource availability for these products, if something does not work troubleshooting it would be extremely difficult.
- These tools require lot of nursing and close monitoring.
- Increase in LAN or WAN traffic.
- You get shuffled between Microsoft Technical Support and Vendor Technical Support. In the end you get to figure it out yourself.

About Standby Continuous Replication (SCR) and Clustered Continuous Replication (CCR)

If you have already migrated to Exchange Server 2007 and have the appropriate hardware and software available, then please go for it. Particularly, SCR would provide most of the features of third party DR tools at no extra cost and it comes bundled with Exchange Server 2007 SP1, but in my experience about 50–60% of Exchange deployments are still based on Exchange server 2003 SP2. CCR and SCR both have their limitations which are listed below:

- Only Mailbox Servers can be clustered, so remaining server roles have to be recovered manually.
- CCR has to be on the same subnet, so it is not site resilient. Hence cannot be considered a right DR solution.
- CCR, LCR and SCC are high availability solutions not DR solutions.
- Limited to only one database per storage group.

Guidelines for Selecting DR Tool

Due to so many choices available right now for Exchange Server DR Solutions, hunting for the right kind of DR solution can be daunting. Some solutions offer simplicity, but are less effective while others require lot of looking after, but are very effective. Some are really expensive while some are considerably cheaper. So what to look out for when choosing the right kind of DR solution? Following are some recommendations which should help you decide the best possible DR tool:

- Choose a solution that offers continuous replication, not just scheduled or streaming replication.
- Must have automatic and manual switchover / failover options.
- Must be application aware, that is, if MSExchangeMTA service is down, it should try to restart the service instead of rolling over to the standby server.

- Must have more than one Network traffic redirection possibilities, for example a solution should offer us not only DNS redirection, but also some other option like changing IP Address on Standby server to match master server in case of a disaster.
- Must be aware of IIS, DNS and AD services related issues. So in case of disaster, these services can be automatically switched over to the standby server.
- DR solution must have near zero Recovery Point Objective (RPO) so that users should not notice any significant amount of downtime.
- Must be site resilient, that is, it should offer replication to a remote site which is not part of the same subnet as that of the Master Exchange server.
- Should offer active notifications and administrative alerts for monitoring replication hygiene
- Should offer fire drill options, that is, automatically test the standby server periodically, to ensure that it can be recovered if the master server goes down.
- Must offer Master Server recovery after the switchover to Standby Server, i.e., offer reverse replication so that master server can be functional again.

Best Practices

Now that you have decided to deploy DR tool of your choice, here are some recommendations to make deployment smooth, operation easy and recovery faster:

- Despite everything, never forget to take a full backup on a daily basis.
- Always have more than one internal DNS server with proper forwarders configured.
- Make sure that DNS response time is within acceptable limits.
- Periodically test LDAP query response time.
- Monitor bandwidth, both on LAN and WAN, closely.
- On Master Exchange server, have only the Default Web Site configured. Do not configure any additional Web servers on that machine.
- If connected on the same LAN, put master and replica servers on separate LAN switches.
- Always manage the DR scenario from the Replica Server.
- Configure notifications for events related to DR scenario.
- Configure notifications for the services related to DR tool, on Exchange Server monitoring option.
- Do not modify the scenario too often as it could lead to database corruption on the standby server.
- In all probability the documentation about the product will not be available online, so give the vendor a call and have them send the documentation to you (it worked for me).
- Again, never forget that full Backup.

Summary

Exchange Server Administrators, like everyone else, need normal 6-8 hours of sleep daily and it may be that they want to catch a movie sometime. It does not have to be always finding that lost email or tracking where it went. Most of the tasks of an Exchange Server admin revolve around creating, deleting, managing and delegating. Backing up and monitoring the Exchange Server is also one of the core tasks, but it is not the only one. So sometimes we usually do not pay much attention to how we are going to recover the server in case of a disaster.

I can confidently say that 8 out of 10 administrators rarely test their backups. So when a disaster strikes we are often caught unaware. And actually doing a fire drill on Exchange Recovery Scenario is hardly ever done, maybe once a year ,but you never know when that *e00* log goes missing and you have to bring your Exchange Server back online before CEO makes a visit to your cubicle (and believe me - that is not pretty). On average, messaging infrastructure design changes every three years, and when it does, make absolutely sure that you have proper Disaster recovery plans in place. Investing in third party tools provides you flexibility to manage your DR scenario the way you like it and it can reduce the anxiety of recovering Exchange Server from a disaster. Third-party DR Tools might not be perfect, but if used carefully, they can provide you a solid DR plan.

Resources

DoubleTake Exchange Disaster Recovery Solution.

CA XOSoft WanSync HA and AR.

Acronis Exchange Recovery.

Exchange Server 2007 SP1 Standby Continuous Replication.

Exchange Server 2010 – The First Public Beta Version

15 April 2009

by Jaap Wesselius

Jaap takes a first look at the new Beta 1 of Exchange Server 2010, which is now available for download, and likes what he sees, especially with Outlook Live, mailbox replication, and the "self-healing" of the database. Much will change before the release, but there are already a lot of new features that are worth a look.

Microsoft has been working quite some time on a new version of Exchange Server and very recently Microsoft has released its first public beta of this new version. It's time to have a first look.

Exchange Server 2010 will be the next version of Exchange server and the successor of Exchange Server 2007. It is targeted for release by the end of this year, but since we're only at the first beta this date is not really fixed of course. We might see a public Beta 2 and a Release Candidate version as well. Beta 1 is by far not feature complete and it can contain bugs and other undocumented features. For sure it will change later this year, so features we're now missing can be added later on. Or vice versa, features we see now can be removed in later versions.

A quick look at Exchange Server 2010 might lead you to believe that it's just another version of Exchange Server 2007 but that's not entirely true. Of course it builds on top of Exchange Server 2007, but there are major improvements and new technologies in this product.

So, what's new?

- One of the first things is that the mailbox replication has changed dramatically, and personally I'm pretty excited about this. Microsoft has taken the Cluster Continuous Replication (CCR) and the Standby Continuous Replication (SCR) features and combined these to create a new feature called "database copies."
- One of the issues with CCR is the added complexity by the Windows Clustering and to lighten up the administrator's life Exchange Server 2010 no longer needs a fully fledged Windows Cluster Server. Under the hood it uses some parts of Windows Clustering but that's completely taken care of by Exchange Server 2010 itself.
- To create highly available mailbox environment, multiple mailbox server can be configured in a "Database Availability Group" or DAG. In a DAG, multiple copies of a mailbox database exist. If one database fails another server automatically takes over, a user will not notice anything about this.
- The concept of multiple databases in a Storage Group is removed in Exchange Server 2010. Also the name "Storage Group" isn't used anymore in Exchange Server 2010. The database technology, which is still based on the Extensible Storage Engine, or ESE still uses the "mailbox database.edb" format as well as the log files (E0000000001.log etc) and the checkpoint file.
- Local Continuous Replication and Standby Continuous Replication have been removed in Exchange Server 2010.
- The database schema has changed, or flattened. The database schema is less complex than in previous version of Exchange server making it possible to reduce the disk I/O compared to Exchange Server 2007 with up to 50% (although we cannot confirm this by individual testing).

- Public Folders are still there and Public Folders are still fully supported in Exchange Server 2010. Even better, there are improvements in Public Folders like enhanced reporting possibilities.
- In Exchange Server 2007 and earlier, MAPI clients connected directly to the mailbox server, while all other clients connected to the Client Access Server. In Exchange Server 2010 MAPI clients now connect to the Client Access Server. No clients connect directly to the Mailbox Server in Exchange Server 2010.
- Enhanced move-mailbox functionality.
- A very enhanced version of Outlook Web Access. One of the design goals was to create a cross browser experience for users. Users on an Apple Macbook with a Safari browser get the same user experience as users using a Windows Vista client with Internet Explorer! A lot of features that were already available in Outlook 2007 are now also available in Outlook Live. Webmail is getting better and better every release of Exchange server.
- Exchange Server 2010 has enhanced disclaimers, which means you can create HTML formatted disclaimers, containing hyperlinks, images and even Active Directory attributes!
- Exchange Server 2010 runs on PowerShell V2 and Windows Remote Management, making it possible to administer remote Exchange Server 2010 servers.

Furthermore there are a lot of changes in the administration of Exchange Server 2010, the routing model, compliancy features etc. Too many to mention in an article like this.

Installing Exchange Server 2010

Installing Exchange Server 2010 is pretty easy, but only on Windows Server 2008. Windows Server 2008 R2 should follow shortly, but for Beta 1 there are still some challenges. Windows should also be a 64-bit version of Windows. It is unclear if a 32-bits version for testing will be available, but like Exchange Server 2007 this one is not supported in a production environment. Other requirements are .NET Framework 3.5, Windows Remote Management 2.0 and PowerShell 2.0.

When finished installing Exchange Server 2010 the Management Console is shown like in Exchange Server 2007 and it looks familiar:

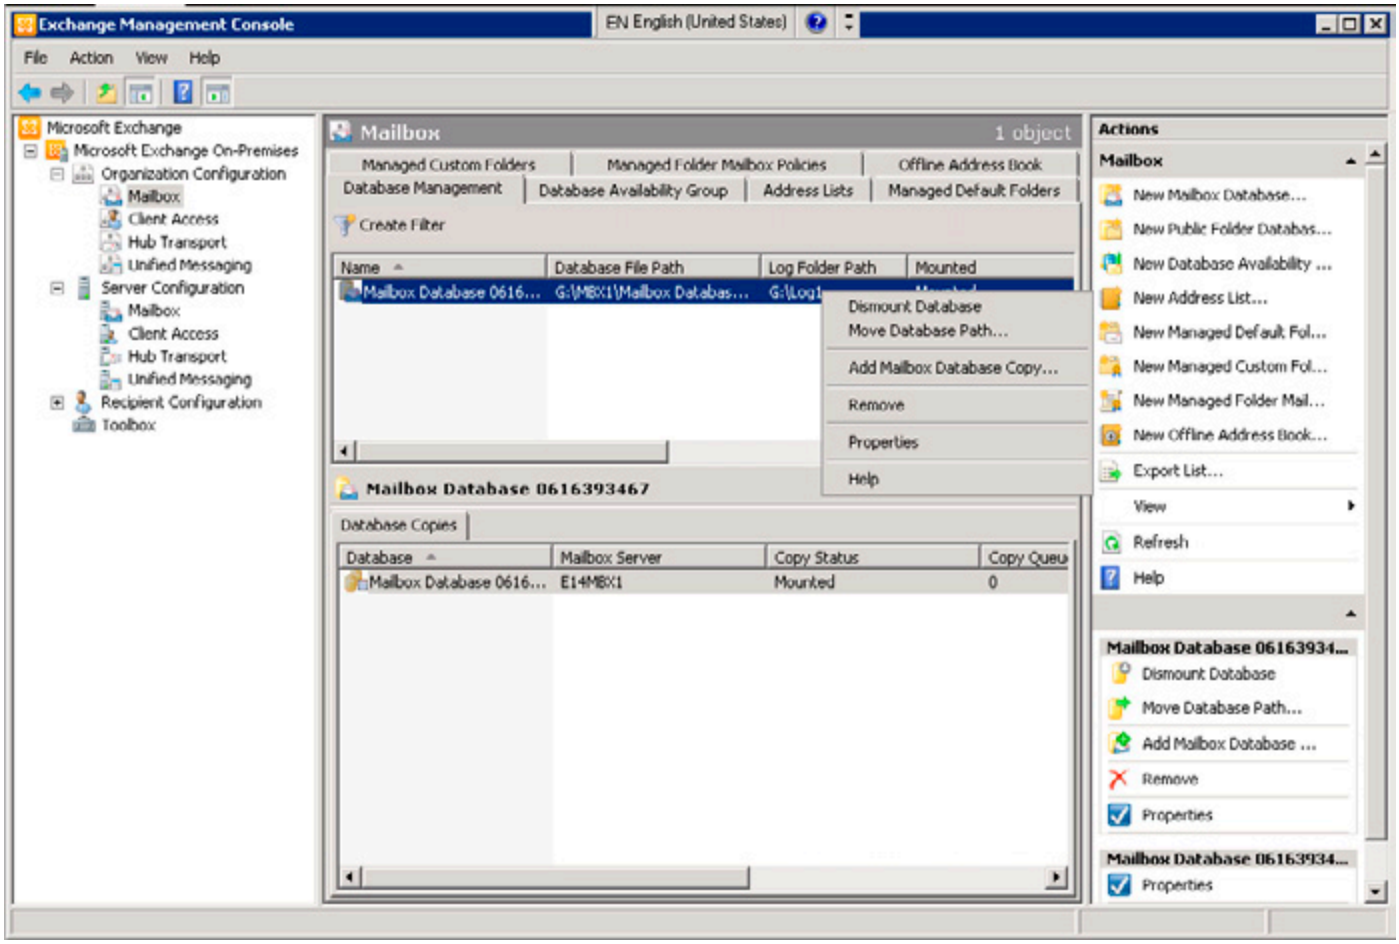

Figure 1. The Exchange Management Console of Exchange Server 2010.

As you can see in Figure 1, the Exchange Management Console looks familiar. But, because of the new high availability features and the flattened database model the database is no longer tied to a particular server but to the Exchange organization. When you want to mount or dismount a database you have to go to the Organization Configuration in the Exchange Management Console and no longer to the Server Configuration. Be aware of this, otherwise it can take you some time before you figure out what's wrong.

Storage Groups no longer exist in Exchange Server 2010, so all cmdlets regarding Storage Groups are removed. Exchange Server 2010 still uses the ESE database, the accompanying log files and checkpoint files, so all Storage Group commandlet options that are still valid for the log file and checkpoint file configuration have been moved to the Database commandlets.

Another neat feature in the new Management Console is the "Send Mail" option. When you are working on recipients and need to send a (test) mail to this recipient you can just right click the recipient and select "Send Mail." No need to send test messages from Outlook or Outlook Live anymore.

As said earlier Microsoft has introduced a concept called "database copies" in Exchange Server 2010. You can install a second Exchange server into the organization and the Exchange setup program takes care of everything. In Exchange Server 2007 only the mailbox role could be installed on a Windows Failover Cluster, in Exchange Server 2010 this is no longer the case. All server roles (except for the Edge Transport role) can be installed on a high availability cluster.

When you've installed a second server holding the Mailbox Server role you can create a copy of the database. Right click on the database and select "Add Mailbox Database Copy," select the 2nd server and you're done.

Names of Mailbox Databases should be unique in the organization and you have to setup a fairly clear naming convention for your Mailbox Databases. If you do not you will certainly get confused with the databases and their copies.

But wait, there's more… since there are multiple copies of a Mailbox Database Microsoft has introduced a "self healing" technique. Exchange knows every copy of the database, and all databases are identical. If a page of a database gets corrupt Exchange can retrieve this page from one of the copies of the database and insert that page in the original database.

In Exchange Server 2010 the move-mailbox functionality is enhanced dramatically. It is now possible to asynchronously move mailboxes. The mailbox is not actually being moved, but it is being synchronized with the new location. The user still accesses, and uses the mailbox on its old location. The move is performed by a new service called the "Mailbox Replication Service" (MRS), running on the Exchange Server 2010 Client Access Server. Like a previous move-mailbox the synchronization can take hours to complete, depending on the amount of data that needs to be synchronized. Once complete, the actual move can take place, but since the data is already in place the move itself will take seconds. Online mailbox moves are only available between Exchange Server 2010 mailboxes and from Exchange Server 2007 SP2 mailboxes.

From an Outlook perspective… in the past Outlook clients connected directly to the back-end server (Exchange Server 2003 and earlier) or to the Exchange Server 2007 mailbox server. Internet clients connected to the Front-end server or to the Exchange Server 2007 Client Access Server. In Exchange Server 2010 the MAPI access also moved to the Client Access Server. A new service is introduced called "MAPI on the Middle Tier" (MOMT), but this name will change before Exchange Server 2010 is officially released. What is the advantage of MAPI clients connecting to the Client Access Server? Suppose something happens to the mailbox database and a fail-over takes place. In the past the Outlook clients were disconnected, the mailbox database was transferred to the other node of the cluster and the clients reconnected. This can take somewhere between 15 seconds and a couple of minutes, depending on the load of the server.
In Exchange Server 2010 when a database fails the Outlook clients stay connected to the Client Access Server and the mailbox is "moved" to the other server. Not really moved, but the Client Access Server just retrieves the information from another copy of the database. This will result is a transparent user experience; he or she will never know what mailbox server the data is coming from, nor do they experience any outage of the mail environment!

Clients….

One of the major improvements on the client side is Outlook Live, previously known as Outlook Web Access. A design goal was to create a cross browser experience so that non-IE users get the same user experience. First test: take an Apple computer, start a Safari browser and open Outlook Live. Wow… that works like a charm:

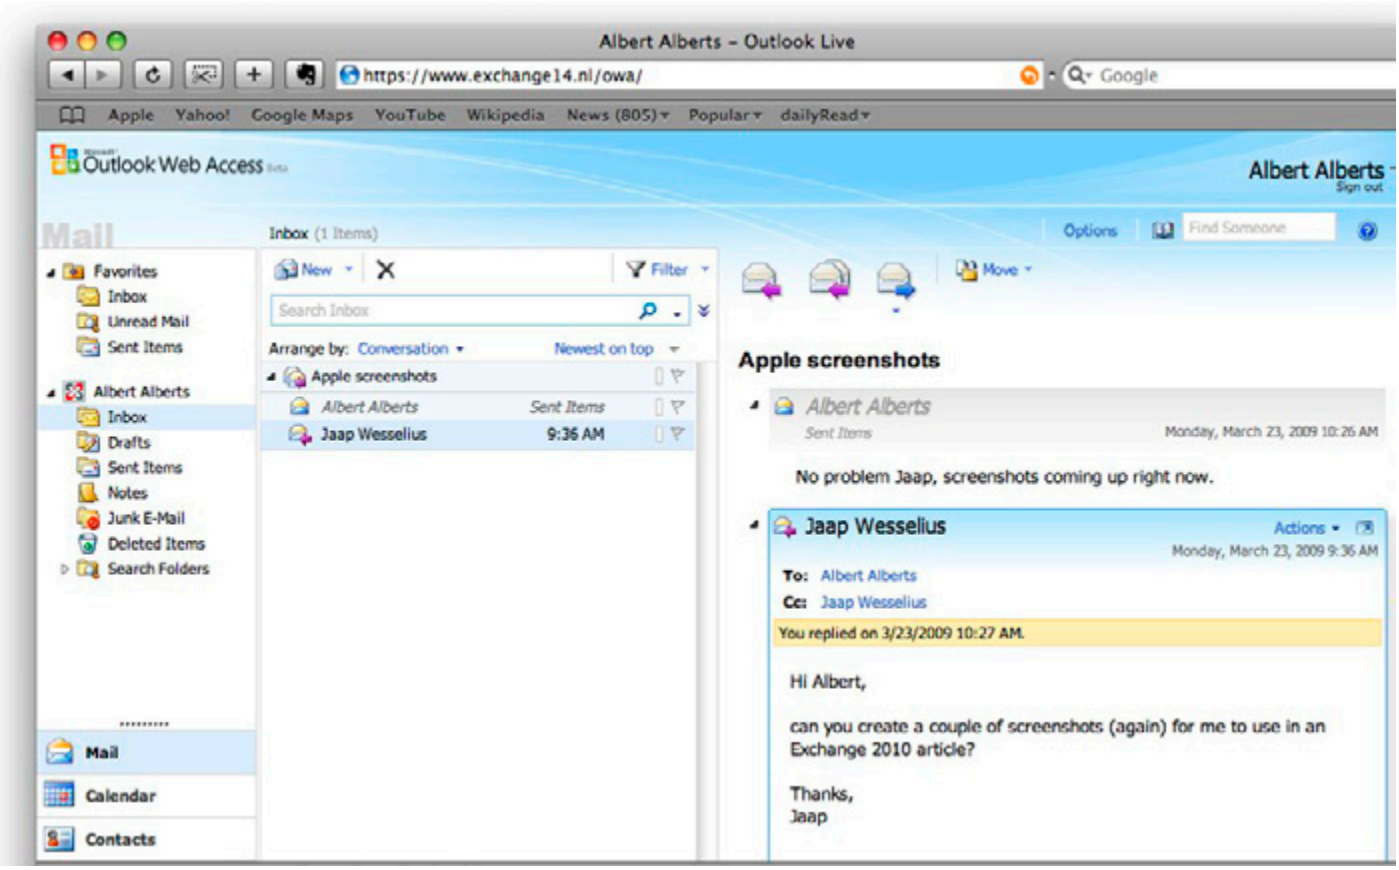

Figure 2. A Safari browser on an Apple Macbook gives a great user experience!

Fairly new in Exchange Server 2010 is the end-user administration option. End users have a lot of extra possibilities regarding the control of their personal information. They can change (basic) user properties in their personal property set like Name, Location, Phone Number etc., but they can also perform some basic administration regarding Distribution Groups.

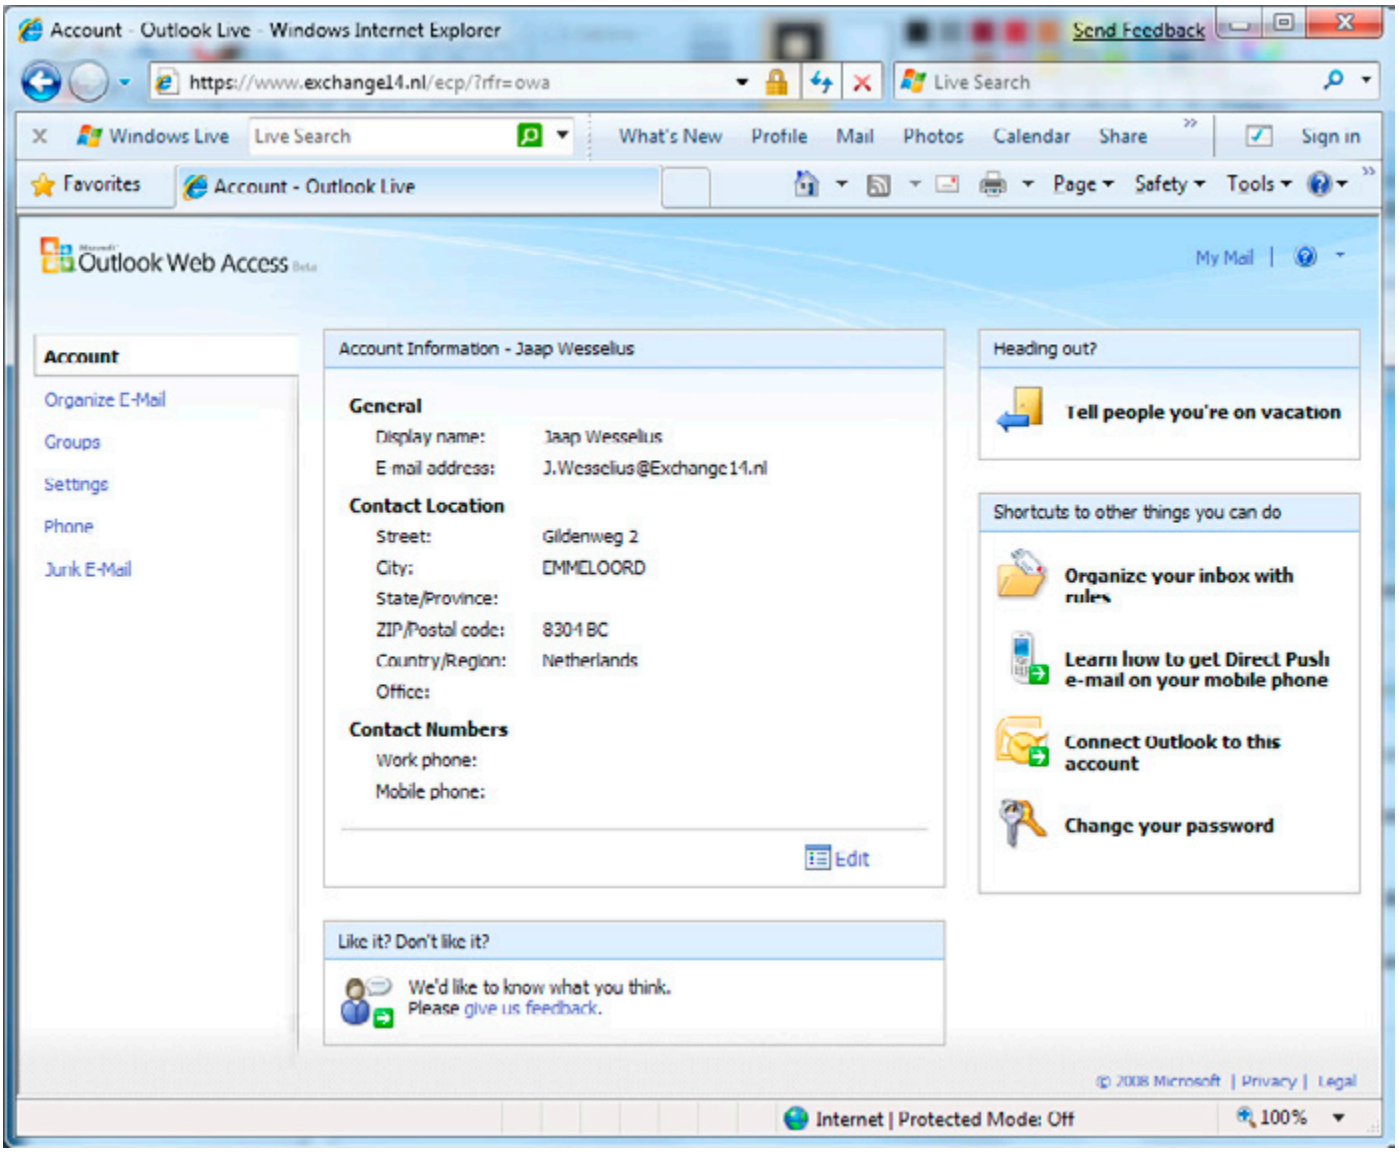

Figure 3. The options page for end users, a complete HTML management interface.

See the "Edit" button? Click here and you can change settings like Contact Location, Contact Numbers and General Information. On the right hand side in the actions pane there are quick links to the most important features for users, like the Out-of-Office assistant or the rules assistant.

And the Groups option in the navigation pane, right here users can create their own distribution groups and manage their own group membership. Don't worry, group owners can restrict ownership. And there's a difference between public and private distribution groups.

The default view in Outlook Live is now in conversation mode in the results pane (=middle pane). In the right pane a quick view of the last message is visible, and below that just some quick notes of earlier messages of this conversation.

Other improvements in Outlook Live are:

- Search Folders.
- Message filtering.
- Side by side view for calendars.

- Attach messages to messages.

Enhanced right-click capabilities.

- Integration with Office Communicator.
- Send and Receive text messages (SMS) from Outlook Live.
- Outlook Live Mailbox polices

But these are interesting topics for a future article

So, What Conclusion?

It's way too early to draw a conclusion about Exchange Server 2010. The only thing I can say is that I'm very enthusiastic about what I've seen so far. The database replication resulting in multiple copies of the data, combined with the self healing possibilities…. That's how the "old" Continuous Cluster Replication should have been. The scalability, the high-availability options, the new Outlook Live, it's all very promising. But, it is still beta 1 and no matter how enthusiastic one is, it's a bad idea to bring this into production. It's even not supported. And before Microsoft hits RTM (Release to Manufacturing) it has a long way to go, and a lot can change. And a lot will change…. But it still looks very promising!

Emulating the Exchange 2003 RUS for Out-of-Band Mailbox Provisioning in Exchange 2007

08 May 2009

by Ben Lye

Exchange's Recipient Update Service was important in Exchange 2000 or 2003 in order to complete the provisioning or updating of mailboxes created from an out-of-band process or from an Active Directory User. In Exchange 2007, full provisioning is automatic, but the functionality can replaced by Powershell scripting if it is required.

Exchange 2000 and Exchange 2003 included a component known as the Recipient Update Service or RUS. The RUS runs as a sub-process of the System Attendant, and is responsible for discovering partially provisioned mailboxes and fully provisioning them. The RUS is required because when user accounts are configured to have a mailbox using Active Directory Users and Computers (ADUC), they only have a few key attributes assigned. When the RUS discovers these partially configured accounts they are stamped with the remaining required attributes and a mailbox will subsequently be created in the appropriate Exchange database when the mailbox is either logged onto or receives a message.

Many Exchange administrators who have worked with Exchange 2000 or 2003 will have experienced times when the RUS hasn't worked at all or has worked slowly, and as a result have had problems with mailbox provisioning or updates.

In Exchange 2007 the RUS is no longer present because when a mailbox is provisioned in either of the Exchange management tools it is automatically fully provisioned straight away. This is great if mailboxes are only ever provisioned using the Exchange 2007 tools, but if you use an out-of-band process to provision mailboxes you will run into trouble as those mailboxes will never be picked up and fully provisioned by Exchange 2007.

For example, in Exchange 2003 it's possible to cause a mailbox to be provisioned simply by setting the LegacyExchangeDN, Mailnickname, and homeMDB attributes on a user object. The RUS will find the user account and will finish provisioning the mailbox. If you try this with Exchange 2007 you won't get a usable mailbox

Fortunately, there is a way to emulate the Exchange 2000/2003 RUS on Exchange 2007 using two Exchange 2007 PowerShell cmdlets, **Update-EmailAddressPolicy** which will apply an email address policy to mailboxes, and **Update-AddressList** which will update address list memberships.

The easiest way to run the "update" cmdlets is to pipe input to them from the "get" cmdlets. These commands will update all email address policies and address lists:

```
Get-EmailAddressPolicy | Update-EmailAddressPolicy
Get-AddressList | Update-AddressList
```

After running these commands the mailbox will be provisioned and will appear in the Exchange address lists, but it will appear in the Exchange Management Console as a "Legacy Mailbox." A legacy mailbox will still function for sending and receiving email, but it's not possible to enable or disable Exchange 2007 features such as Messaging Records Management or Unified Messaging. Legacy mailboxes can be converted to Exchange 2007 mailboxes using the **Set-Mailbox** cmdlet with the *-ApplyMandatoryProperties* parameter.

If your Exchange environment has only Exchange 2007 servers this command will update all legacy mailboxes to Exchange 2007 mailboxes:

```
Get-Mailbox -RecipientTypeDetails LegacyMailbox | Set-Mailbox –ApplyMandatoryProperties
```

If you have mailboxes hosted on legacy versions of Exchange and on Exchange 2007 you will need to filter the command to only touch those hosted on Exchange 2007. This command will do that by first getting only the Exchange 2007 servers then finding all the legacy mailboxes on those servers:

```
Get-ExchangeServer |
Where{$_.AdminDisplayVersion.ToString().SubString(0, 10) -eq"Version 8."} | ForEach{
Get-Mailbox -Server $_.Name -RecipientTypeDetails LegacyMailbox | Set-Mailbox
-ApplyMandatoryProperties}
```

The same problem exists for changes made to users outside of the Exchange Management tools. Any changes made within the Exchange 2007 tools are instantly reflected in the user's email addresses and address list membership, but changes which would impact email address polices and address list memberships that are made using out-of-band tools such as ADUC, ADSIEdit, or other tools for directly editing Active Directory will not show up in the Exchange 2007 mailbox properties until the same "update" commands are run.

For example, if there is an email address policy which applies email addresses in the format firstname.lastname, and the last name of a user is changed using an Exchange 2007 tool the user will instantly get a new email address with the new last name. If the same change is made using an out-of-band tool or process then the name change will not be reflected in the user's Exchange email addresses.

This problem can also be solved by running the **Update-EmailAddressPolicy** and **Update-AddressList** cmdlets.

So, if your environment requires out-of-band mailbox provisioning or user updates then the simplest solution is to put these three commands into a PowerShell script and schedule it to run as often as needed (don't forget that if you still have mailboxes hosted on legacy Exchange servers you need to change the last command to the filtered version shown above):

```
Get-EmailAddressPolicy | Update-EmailAddressPolicy
Get-AddressList | Update-AddressList
Get-Mailbox -RecipientTypeDetails LegacyMailbox | Set-Mailbox –ApplyMandatoryProperties
```

To run the **Update-EmailAddressPolicy** cmdlet you need to be delegated the Exchange Server Administrator role and have local Administrators rights for the Exchange server. To run **Update-AddressList** you need to be delegated the Exchange Organization Administrator role. To run the **Set-Mailbox** cmdlet you need to be delegated the Exchange Recipient Administrator role.

More Information on the cmdlets can be found on the Microsoft TechNet website:

Update-EmailAddressPolicy – http://technet.microsoft.com/en-us/library/aa996869.aspx.

Update-AddressList – http://technet.microsoft.com/en-us/library/aa997982.aspx.

Set-Mailbox – http://technet.microsoft.com/en-us/library/bb123981.aspx.

Using Exchange 2007 Transport Rules to Protect the First Entry in the Address Book

08 June 2009

by Ben Lye

Global Address Lists in MS Exchange can cause problems because the first person in the list often gets the reply. Ben Lye shows how one can eliminate any such problems with Global Address lists by creating a mail contact, a transport rule and a custom Delivery Status Notification.

I was recently asked to add an entry to the Outlook address book which would prevent mail being inadvertently sent to the person who normally appeared at the top of the Global Address List (GAL) – the person was getting frustrated by receiving e-mail which was clearly not intended for him.

There are several possible ways this could be achieved: a mailbox with an auto-reply or out-of-office rule, a public folder with an auto-response, a non-Exchange auto-responder, or simply a mail contact with an invalid external address, to name a few.

Ideally I wanted the e-mail to be stopped on Exchange without requiring an extra mailbox or public folder, and I wanted the sender should get a helpful error message. These requirements meant that a mailbox or public folder with an auto-reply rule could not be part of the ideal solutions, and neither could a non-Exchange auto-responder as that would mean that the email would have to leave the Exchange environment before being stopped. Additionally, using a contact record with an invalid address was not perfect either; because a user who e-mailed the contact would simply receive a rather unhelpful "address unknown" non-delivery report (NDR).

Fortunately Exchange 2007 provides mechanisms which can be used to provide a neat solution to this problem. The solution has three parts: a mail contact which will appear at the top of the GAL, a transport rule to prevent mail being sent to the contact, and a custom delivery status notification (DSN) to provide the user with information about why their message was not delivered.

Mail Contacts are Active Directory objects which are typically used to add e-mail addresses which are external to Exchange 2007 to the Global Address List. To create a new mail contact you must be delegated the Exchange Recipient Administrator role and the Account Operator role Active Directory container where you wish to create the contact.

Transport rules run on Exchange 2007 servers which have either the Hub Transport or Edge Transport role installed. They can be used to control the flow of e-mail messages within the Exchange 2007 organization and can be used for a variety of purposes including for restricting e-mail between certain individuals or groups, or for applying a footer to all e-mail destined for Internet recipients.

Custom DSN messages give Exchange 2007 administrators the facility to create new DSN messages for custom delivery notifications and the ability to customize existing DSN messages. They are a useful tool if you wish to provide users with links to further information such as links to self-help knowledge base articles, or contact information for help-desk staff.

To create transport rules and custom DSN messages you must be delegated the Exchange Organization Administrator role.

The first step in implementing this solution is to create a new mail contact which has a display name that will ensure it is shown as the first entry in the GAL. An easy way to do this is to prefix the display name with a period or underscore or any other valid character which does not normally appear in Exchange display names (spaces are prohibited as leading characters in display names). We also need to specify an e-mail address which is not in use by another e-mail enabled object.

The mail contact can be created in the Exchange Management Shell using the **New-MailContact** cmdlet:

```
New-MailContact -Name ."First Address Book Entry" -ExternalEmailAddress
"firstaddressbookentry@example.com" -Alias "firstaddressbookentry"
```

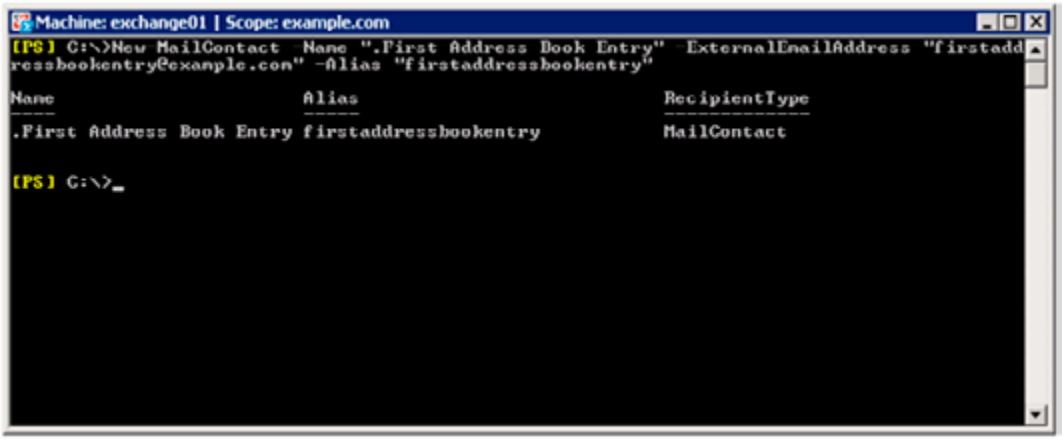

The new mail contact will appear in the Outlook address book:

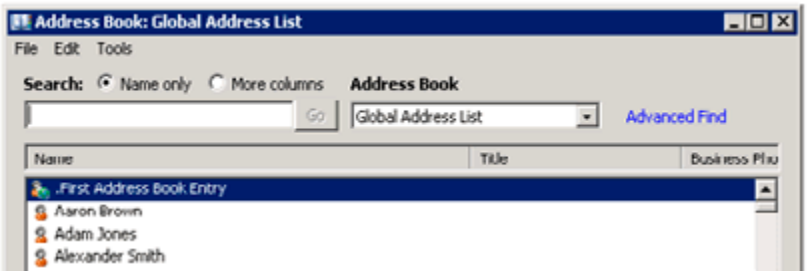

The second step is to create a new custom delivery status notification (DSN) message which will be sent to anybody who e-mails the new mail contact. Custom DSN messages can contain plain text or HTML, and in this case will provide useful information to the user pointing out that their message probably did not reach the intended recipient.

We'll create a DSN message for DSN code 5.7.10, which is the first available enhanced status code (the valid range is 5.7.10 through 5.7.999 inclusive).

The Exchange Management Shell cmdlet for creating DSN messages is **New-SystemMessage**:

```
New-SystemMessage -DsnCode 5.7.10 -Text 'You have accidentally sent an e-mail to the first
entry in the Outlook Global Address List.<br>You may wish to resend your message to the correct
recipient.' -Internal $True -Language en
```

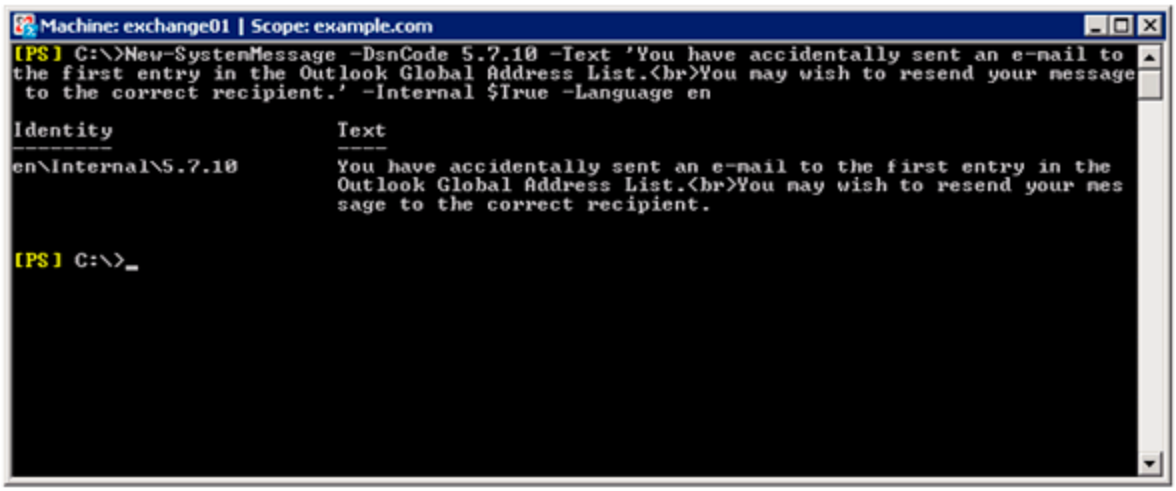

The final step is to create a new transport rule which will send the new DSN message to anybody who e-mails the new mail contact.

Transport rules consist of three components: conditions, actions, and exclusions. To create a new transport rule we must specify at minimum the action to be taken, but in this case we'll specify a condition and an action. The transport rule can be created in the Management Shell using these commands, incorporating the **New-TransportRule** cmdlet:

```
$Condition = Get-TransportRulePredicate SentTo
$Condition.Addresses = @(Get-MailContact ."First Address Book Entry")
$Action = Get-TransportRuleAction RejectMessage
$Action.RejectReason = "E-mail to the first entry in the GAL is not allowed."
$Action.EnhancedStatusCode = "5.7.10"
New-TransportRule -Name "Block e-mail to the first entry in the GAL" -Conditions @($Condition)
-Actions @($Action) -Priority 0
```

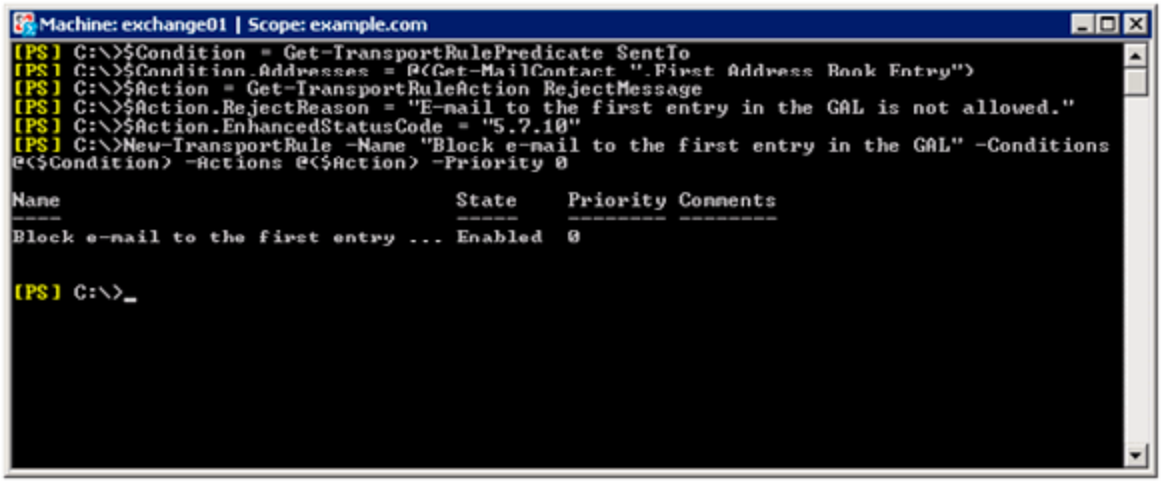

With the new mail contact in the Global Address List, the new DSN created, and the transport rule set up, if you sent a test e-mail message to the new contact you will receive this NDR message back:

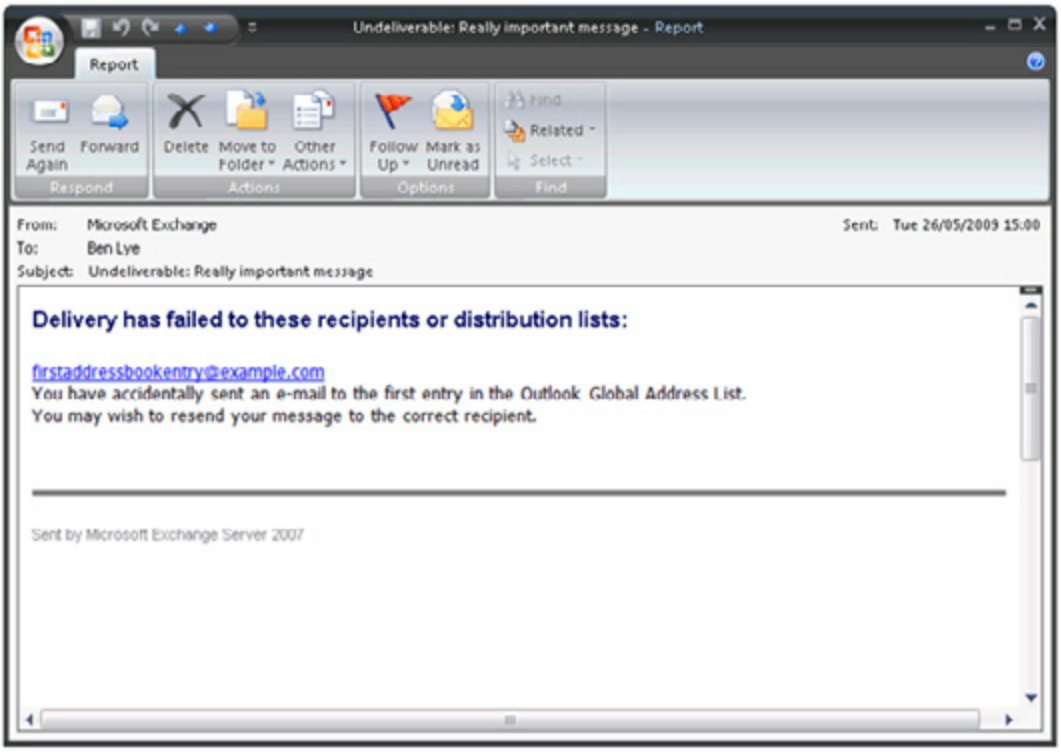

Using this solution my objectives have been met: I didn't have to create a mailbox or public folder, the email message doesn't leave the Exchange environment, and the sender receives a useful error message.

For more information about creating mail contacts:
http://technet.microsoft.com/en-us/library/aa997220.aspx.

For more information about custom DSN messages:
http://technet.microsoft.com/en-us/library/aa998878.aspx.

For more information about Transport rules:
http://technet.microsoft.com/en-us/library/aa995961.aspx.

Cluster Continuous Replication Network Design

18 June 2009

by Neil Hobson

Cluster continuous replication (CCR) is a means of providing a more resilient email system with faster recovery. It was introduced in Microsoft Exchange Server 2007 and uses log shipping and failover. The host-name feature of Exchange 2007 Service Pack 1 and the use of multiple redundant private network interfaces can improve the effectiveness of CCR.

Introduction

In a Cluster Continuous Replication (CCR) environment, each node of the cluster will usually have two network interfaces configured. One network interface will be used for client communications whilst the other will be used for intra-cluster communications. However, there are additional configuration options that you should consider when designing the network configuration of your CCR environment. In Exchange 2007 Service Pack 1, a new feature was introduced that allowed additional redundant networks to be added to a CCR environment for the purposes of log shipping and database seeding. With specific networks dedicated to log shipping and database seeding, the remaining network is dedicated to its task of servicing client communications. Such a configuration avoids situations where the network that is servicing client communications also has to process a large number of transaction logs. Additionally, since redundant networks are used, overall high availability of the CCR environment is enhanced.

In this article, I am going to describe a highly-available configuration that is required to support multiple redundant networks used for the purposes of log shipping and database seeding. The actual configuration process is slightly different depending on whether the CCR environment is being implemented on Windows 2003 or Windows 2008, so I will be highlighting the differences where appropriate. However, please note that the focus of this article is the Windows 2008 operating system. In the configuration that I am describing within this article, each node of the CCR environment is equipped with four network interfaces configured in the following fashion:

- Two network interfaces are configured as a network team and will be configured as the public network used for client connectivity.
- A single network interface will be configured as a private network for intra-cluster communications as well as transaction log shipping.
- Another single network interface will also be configured as a private network that is also used for intra-cluster communications as well as transaction log shipping.

With this configuration, Exchange will select one of the private networks for log shipping. If this private network were to fail, Exchange will select the remaining private network and will only revert to using the public network should both private networks fail or be otherwise unavailable.

Figure 1 shows this configuration, including the IP addresses in use. The cluster node server names are CCRA and CCRB respectively.

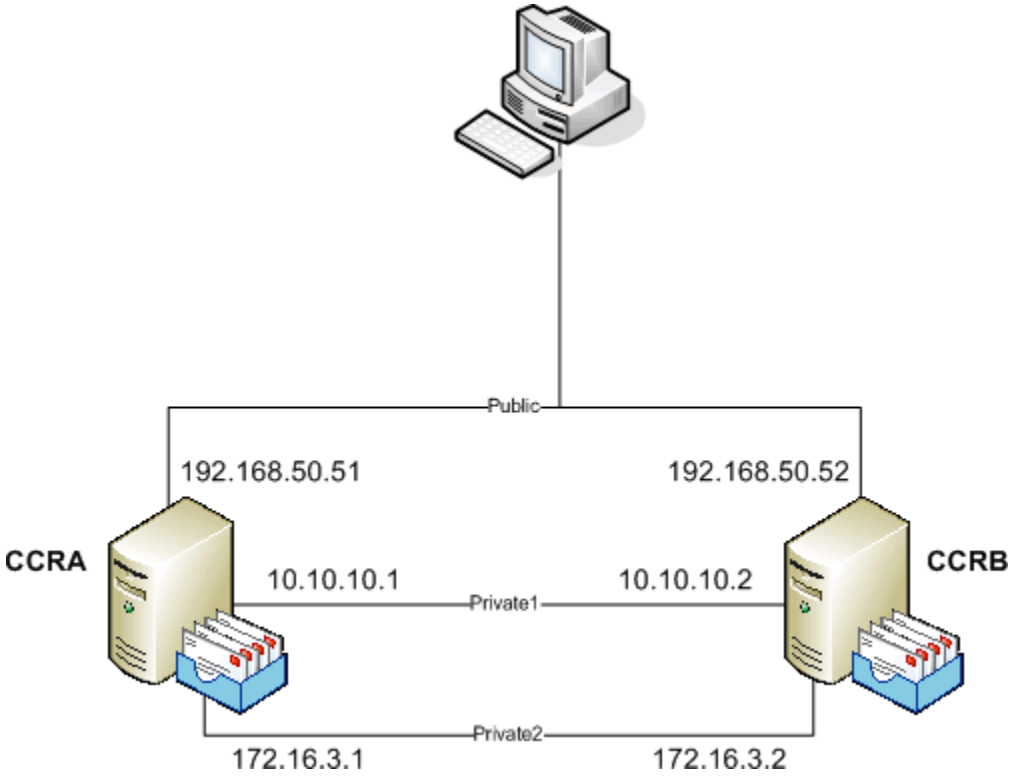

Figure 1: CCR Environment with Three Networks.

The advantage of having two separate network interfaces configured as private networks is that this configuration ensures that there is full redundancy for the private networks. However, this does not mean that it's not possible to implement redundant networks for log shipping and database seeding if there are only three network interfaces in each cluster node. It's perfectly possible but the configuration will be slightly different. With just three network interfaces available, the following configuration will be possible:

- Two network interfaces configured as a network team and configured as the public network used for client connectivity.
- A single network interface configured as a private network for intra-cluster communications as well as transaction log shipping.

With this configuration, Exchange will use the private network for log shipping and will only revert to using the public network should the private network fail or be otherwise unavailable. Therefore, utilizing four network interfaces in each cluster node gives more flexibility to the design and is the configuration that I am using in this article.

Network Teaming

Network teaming is simply the process whereby multiple network interfaces are grouped together to form a single logical network interface, with the objective being to offer fault tolerance and load balancing. In the network interface configurations mentioned previously within this article, it may have been noticed that only the public network interface is teamed. This configuration exists because Microsoft does not support teamed network card configurations for the private network interfaces in a clustered environment. More information on this can be found in Microsoft Knowledgebase Article 258750.

Rename Network Interfaces

It's recommended to rename the network connections so that they reflect the role that they are performing, as this can be very helpful when troubleshooting issues. Throughout this article, the network interfaces will be referred to using the following names:

- *Public*. This is the teamed network interface that services client connectivity.
- *Private1* and *Private2*. These are the two private network interfaces that service intra-cluster communications and will be configured for transaction log shipping. They can also be used for database seeding.

Configuration Procedure

Now that the background information on the network configuration has been described along with the setup that will be configured in this article, it's time to take a look at the actual configuration process involving continuous replication host names. These are additional host names that are assigned to each cluster node for the purposes of transaction log shipping and database seeding. The configuration process will be divided into the following broad stages:

- Choose the continuous replication host names and IP addresses.
- Configure the private network interfaces as mixed networks in the cluster management application.
- Disable strict name checking. This is only required if a CCR environment is being deployed on the Windows 2003 operating system.
- Configure name resolution.
- Enable NetBIOS on the private network interfaces. This is only required if a CCR environment is being deployed on the Windows 2008 operating system.
- Enable the continuous replication host names.
- Check that the correct continuous replication host names are being used.

Choose Replication Host Names and IP Addresses

Regardless of whether you are using Windows 2003 or Windows 2008 as your operating system of choice for your CCR environment, you must first plan up front the continuous replication host names and IP addresses that will be used. Each private network configured in the CCR environment requires a continuous replication host name and IP address assigned to it in addition to the existing IP address it will have. For example, it can be seen from Figure 1 that the cluster node CCRA has a private network interface called Private1 configured with the IP address of 10.10.10.1.

This private network interface now needs another IP address assigned to it and this IP address linked with a continuous replication host name. Extrapolating this information out, each cluster node within the CCR environment described within this article has two continuous replication host names and IP addresses assigned to it, since two independent private networks are being used for this purpose. It's these continuous replication host names that are used to identify which of the networks are used to ship the transaction logs from one cluster node to the other.

It's possible to choose any host names but I personally find it useful to ensure that they are linked in some way with each cluster node. For example, in the CCR environment that is being described within this article the two cluster nodes are called CCRA and CCRB. Therefore, for cluster node CCRA, I've chosen continuous replication host names of CCRAREP1 and CCRAREP2 whilst for cluster node CCRB

I've chosen continuous replication host names of CCRBREP1 and CCRBREP2. On cluster node CCRA, the continuous replication host name CCRAREP1 will be assigned to the network interface called Private1, whilst the continuous replication host name CCRAREP2 will be assigned to the network interface called Private2. A similar configuration will be applied to the cluster node CCRB. Each continuous replication host name is assigned a new IP address.

Putting all this information together, the example CCR environment configuration is shown in Figure 2. Note that the public network has been removed for reasons of clarity.

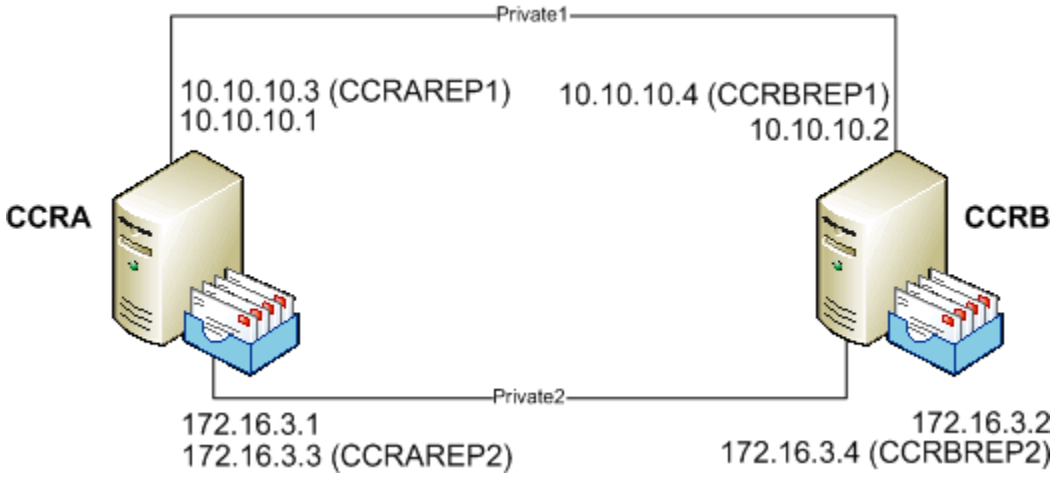

Figure 2: CCR Environment with Continuous Replication Host Names.

Mixed Cluster Networks

In a typical cluster configuration consisting of a single public network and a single private network, the public network will be configured within the cluster management application as a *mixed cluster network*. This means that this network will be used to process client communications as well as intra-cluster communications. Intra-cluster communication is often referred to as the *heartbeat*. Conversely, the private network will be configured within the cluster management application as a *private network*, meaning that it will only be used for intra-cluster communications.

One of the main configuration differences required when implementing redundant cluster networks for continuous replication host names is the fact that the private networks must be configured as mixed cluster networks. When using Windows 2003, this is achieved using the *Cluster Administrator* snap-in whilst in Windows 2008 it can be achieved using the *Failover Cluster Management* snap-in. Therefore, in the environment being described within this article, both the Private1 and Private2 networks have been configured as mixed cluster networks as shown in Figure 4, where the Failover Cluster Management snap-in is being used. Figure 3 shows the configuration when using the Cluster Administrator application on Windows 2003.

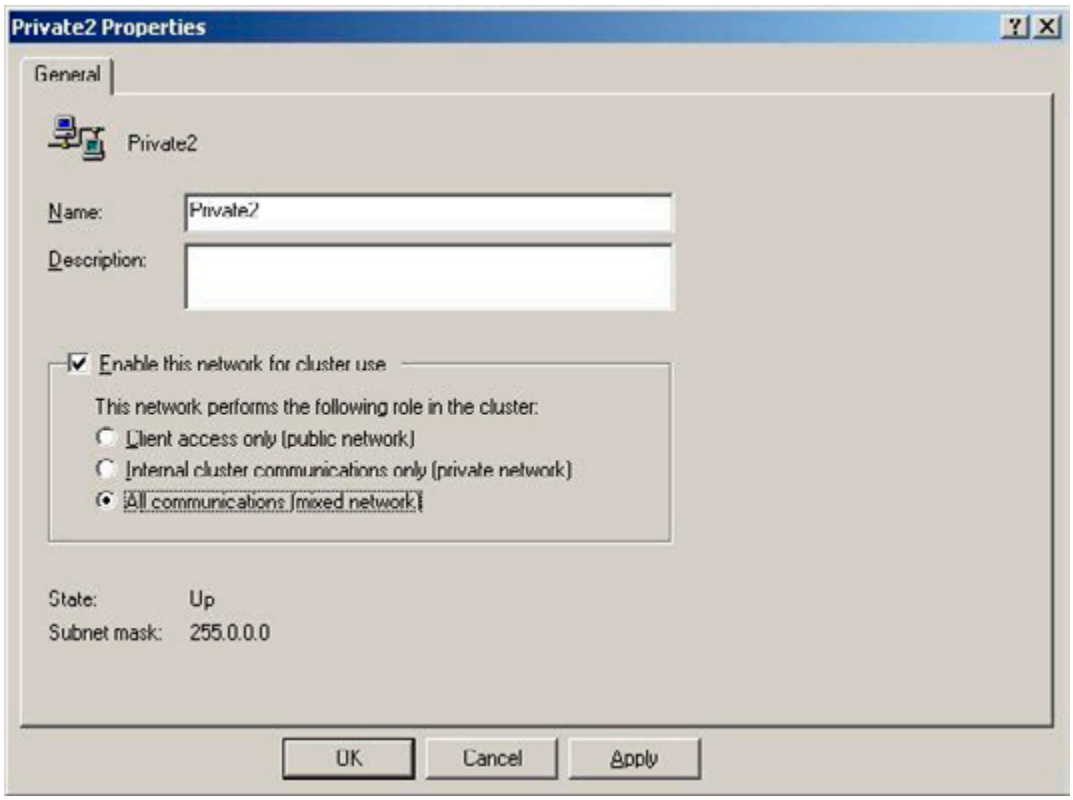

Figure 3: Mixed Cluster Network on Windows 2003.

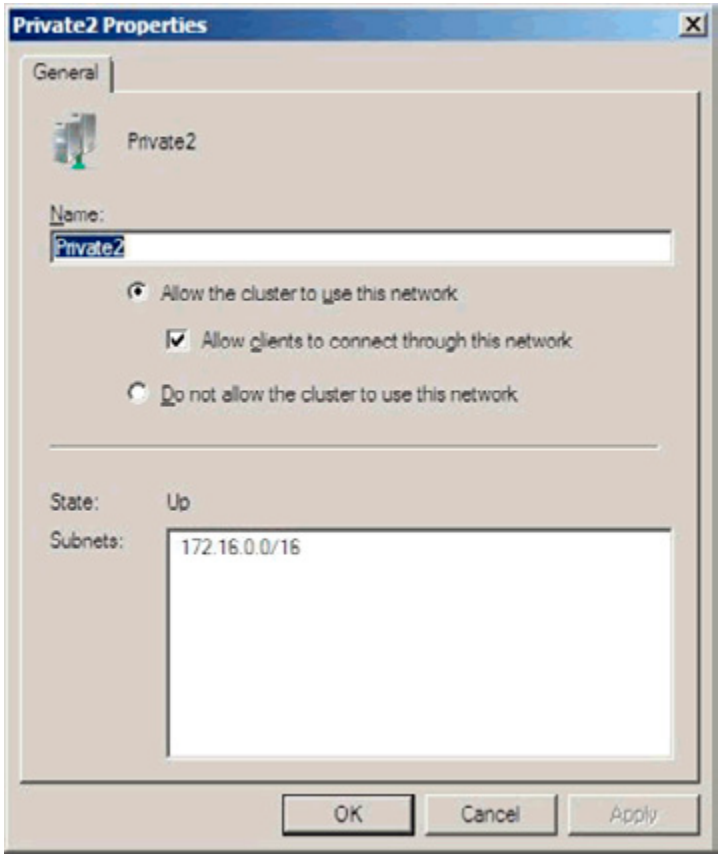

Figure 4: Mixed Cluster Network on Windows 2008.

Configuring the private networks as mixed networks within the cluster management application may seem counter intuitive based on the fact that mixed cluster networks generally service client communications. However, even though the private networks will not be servicing client communications in this case, they must still be configured as mixed cluster networks within the cluster management application so that cluster resources can be created for them.

Strict Name Checking (Windows 2003)

An additional configuration change must be made if the CCR environment has been deployed using Windows 2003. Strict name checking must be disabled on both cluster nodes to ensure that the Windows 2003 operating system correctly listens for connections being made using the continuous replication host names. This change does not need to be made if the CCR environment has been deployed using the Windows 2008 operating system.

To make this change, the registry must be modified on both cluster nodes. The required registry information is as follows:

```
Key: HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\lanmanserver\parameters
DWORD Value: DisableStrictNameChecking
Value: 1
```

After making this registry change, the server must be restarted for the change to take effect. Remember, this change must be made on both cluster nodes of the CCR environment.

Name Resolution

Regardless of whether the CCR environment is deployed using Windows 2003 or Windows 2008, some aspects of name resolution must be configured on each private network interface being used for continuous replication. There are two configuration requirements to complete.

The first requirement is to clear the check box named *Register this connection's addresses in DNS*. To do this in Windows 2008, follow these steps:

- Bring up the properties of the network interface.
- Select the *Internet Protocol Version 4 (TCP/IPv4)* object and then click the *Properties* button.
- In the resulting properties window, click the *Advanced* button.
- In the *Advanced TCP/IP settings* window, click the *DNS* tab.
- Clear the check box named *Register this connection's addresses in DNS* as shown in Figure 5.

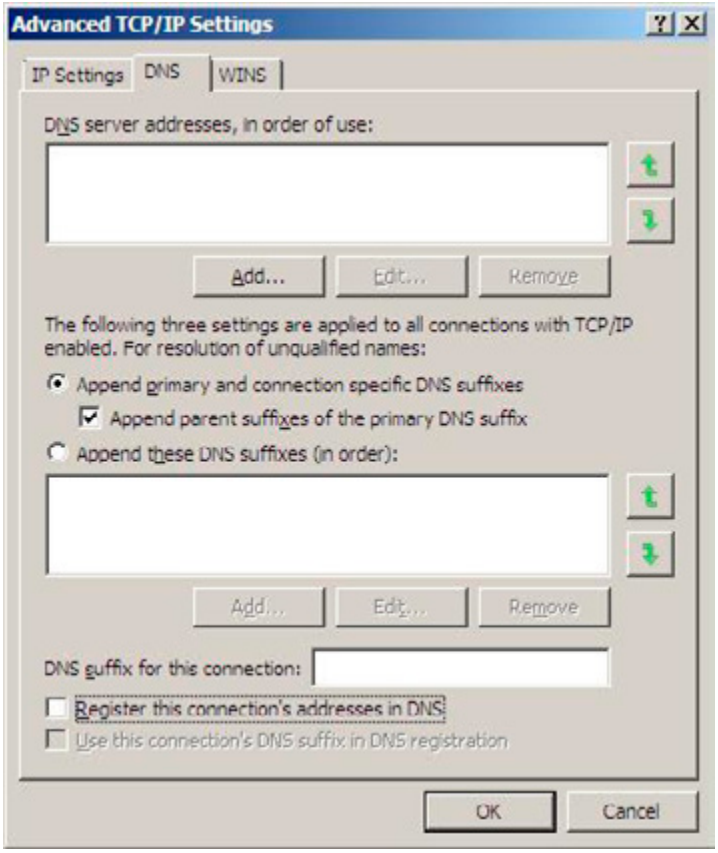

Figure 5: DNS Configuration Change.

Next the continuous replication host names must be entered in the Domain Name System (DNS), or host file entries must be created. Generally, DNS is recommended where possible as tracking manually created host file entries can sometimes be problematic. Therefore, in this article it is assumed that static DNS records have been created for the continuous replication host names. Don't underestimate the importance of this step as name resolution plays a vital role in allowing the log shipping or database seeding to take place via the redundant networks being configured. Therefore, in the example environment being described within this article four static DNS A records have been created, one for each continuous replication host name. These records are highlighted in Figure 6.

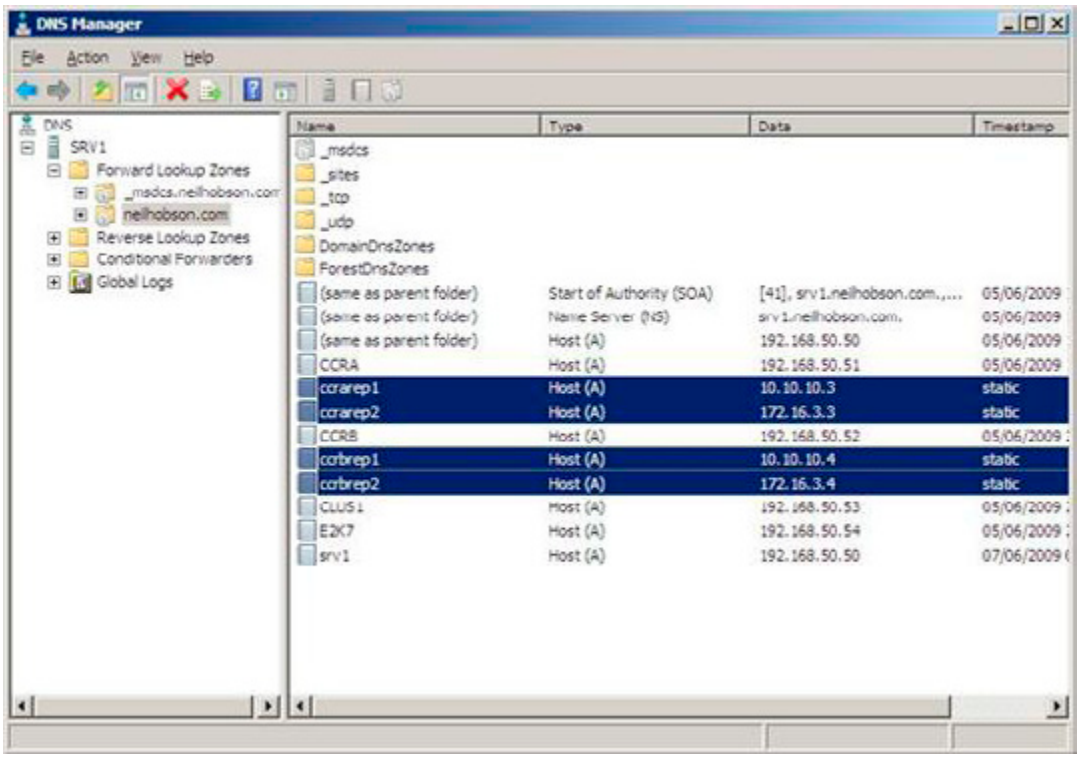

Figure 6: DNS Static Records.

Enable NetBIOS (Windows 2008)

If the CCR environment has been deployed using the Windows 2008 operating system, there is another important networking configuration that must be made on each private network interface being configured with continuous replication host names. Specifically, NetBIOS must be enabled on the private network interface to ensure that the new network name resource, that will be created later, will come online successfully. To enable NetBIOS on the private network interfaces, follow these steps:

- Bring up the properties of the network interface.
- Select the *Internet Protocol Version 4 (TCP/IPv4)* object and then click the *Properties* button.
- In the resulting properties window, click the *Advanced* button.
- In the *Advanced TCP/IP settings* window, click the *WINS* tab.
- Select the *Enable NetBIOS over TCP/IP* button as shown in Figure 7.

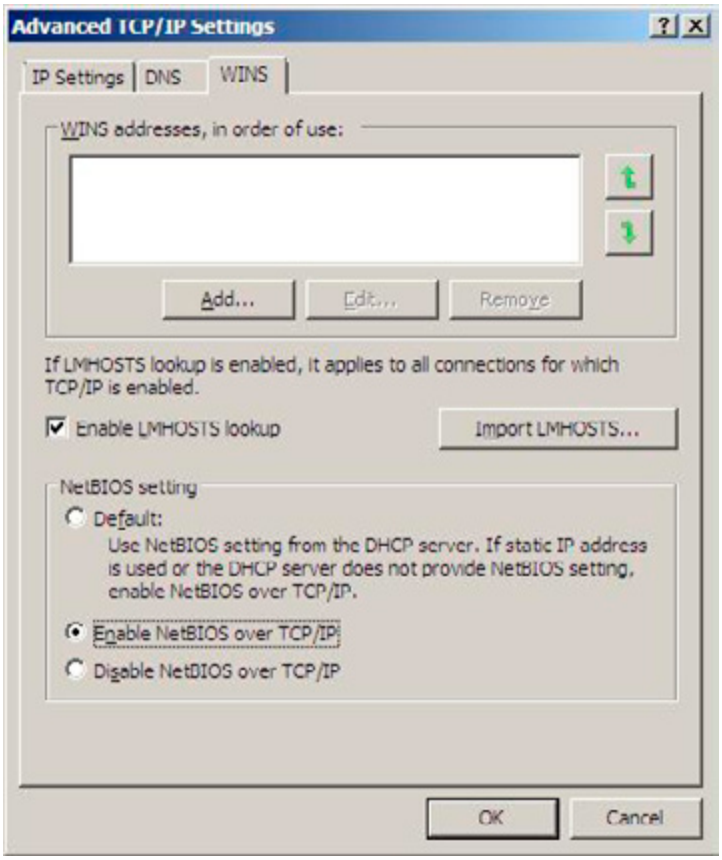

Figure 7: Enabling NetBIOS.

Enable Continuous Replication Host Names

The next step of the configuration process is to enable the continuous replication host names on the CCR environment. To do this, the *Enable-ContinuousReplicationHostName* cmdlet is run. Since the example environment in this article is using four continuous replication host names, it follows that the *Enable-ContinuousReplicationHostName* cmdlet will need to be run a total of four times to complete the required configuration. The parameters used in this cmdlet are as follows:

- *Identity*. This parameter is used to identify the name of the Clustered Mailbox Server (CMS) being configured. The CMS is the server name that the Outlook clients connect to in a CCR environment.
- *TargetMachine*. The TargetMachine parameter is used to identify the cluster node that is being configured.
- *HostName*. This parameter is the actual continuous replication host name being configured on the server referenced via the TargetMachine parameter.
- *IPV4Address*. The IPV4Address parameter is the IP address associated with the HostName parameter that has just been discussed. It's therefore useful to think of the HostName and IPV4Address parameters as a logical pairing.
- *Confirm*. When scripting PowerShell cmdlets, it can sometimes be useful to prevent the Exchange Management Shell from prompting you for confirmation. In this example, the Confirm parameter will be used to ensure that prompts for confirmation are not encountered during the running of the cmdlet.

Here are the four cmdlets that will be run to enable continuous replication host names within the example CCR environment being described within this article:

```
Enable-ContinuousReplicationHostName -Identity E2K7 -TargetMachine CCRA `
  -HostName CCRAREP1 -IPV4Address 10.10.10.3 -Confirm:$false
Enable-ContinuousReplicationHostName -Identity E2K7 -TargetMachine CCRA `
  -HostName CCRAREP2 -IPV4Address 172.16.3.3 -Confirm:$false
Enable-ContinuousReplicationHostName -Identity E2K7 -TargetMachine CCRB `
  -HostName CCRBREP1 -IPV4Address 10.10.10.4 -Confirm:$false
Enable-ContinuousReplicationHostName -Identity E2K7 -TargetMachine CCRB `
  -HostName CCRBREP2 -IPV4Address 172.16.3.4 -Confirm:$false
```

It can be useful to have the cluster management program open before running these cmdlets, as it's possible to see that a new cluster group is created each time that one of the cmdlets is run. For example, Figure 8 shows the Windows 2008 Failover Cluster Management application after the running of the first cmdlet listed above.

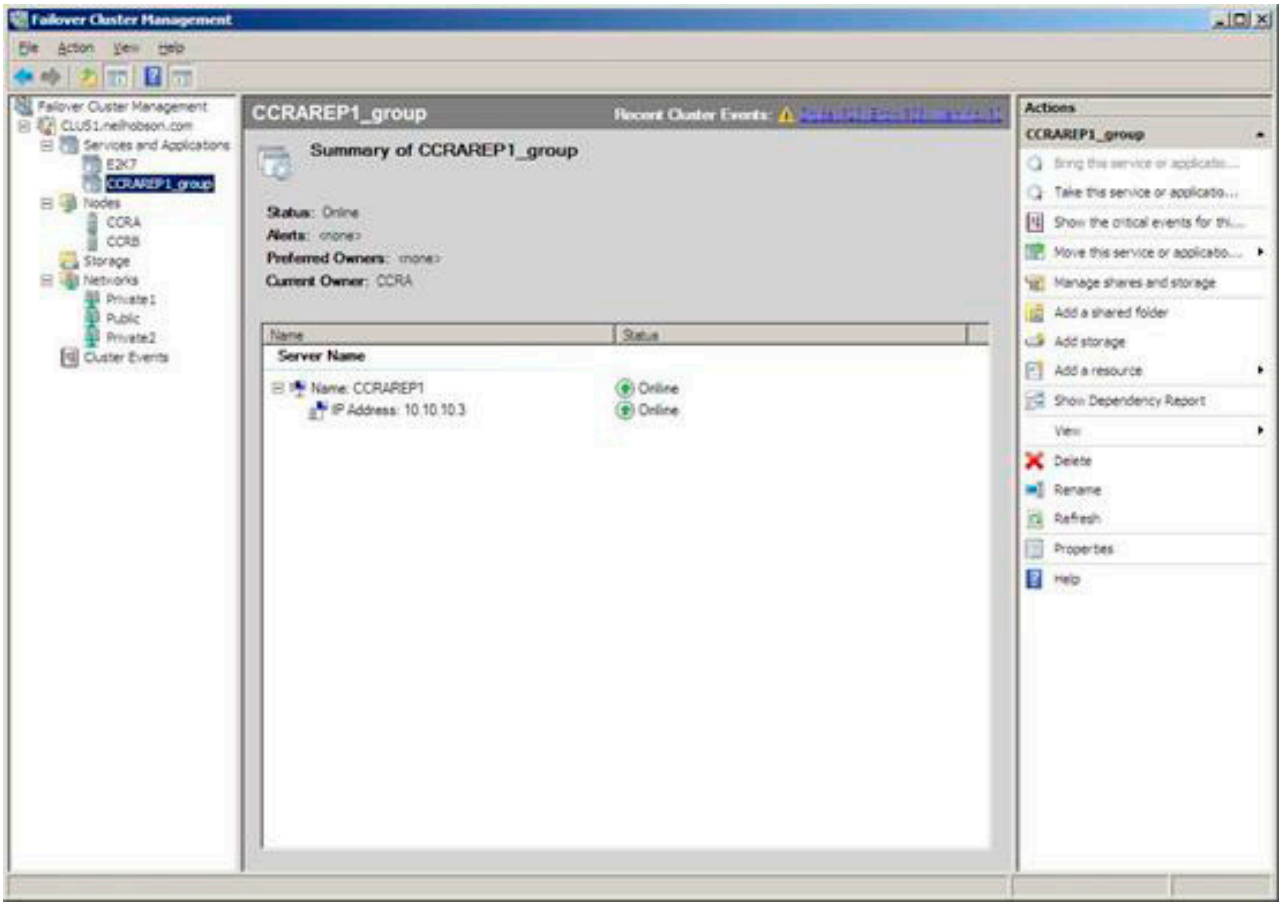

Figure 8: New Cluster Resource Group 5.

It was discussed earlier in this article that NetBIOS must be enabled on the private network interfaces in order that the cluster resources can be brought online successfully. If this is not performed the following error will be seen when running the *Enable-ContinuousReplicationHostName* cmdlet:

```
Cluster Common Failure Exception: Failed to bring cluster resource Network Name (name) in cluster group (cluster group name) online. The event log may contain more details. Cluster Common Failure Exception: The cluster resource could not be brought online by the resource monitor. (Exception from HRESULT: 0x8007139A)
```

Figure 9 shows this error with the Exchange Management Shell.

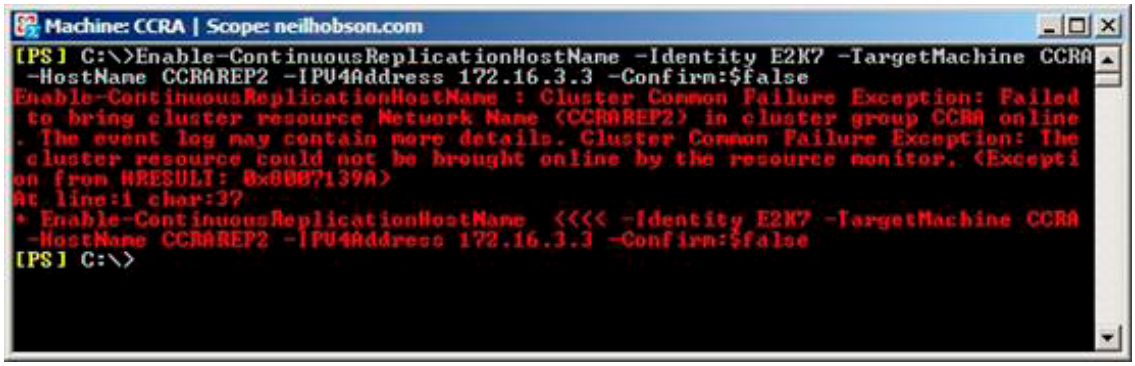

Figure 9: Resource Error in Exchange Management Shell.

Within the Failover Cluster Management application, it can be seen that the corresponding Network Name resource has failed to come online, as can be seen in Figure 10.

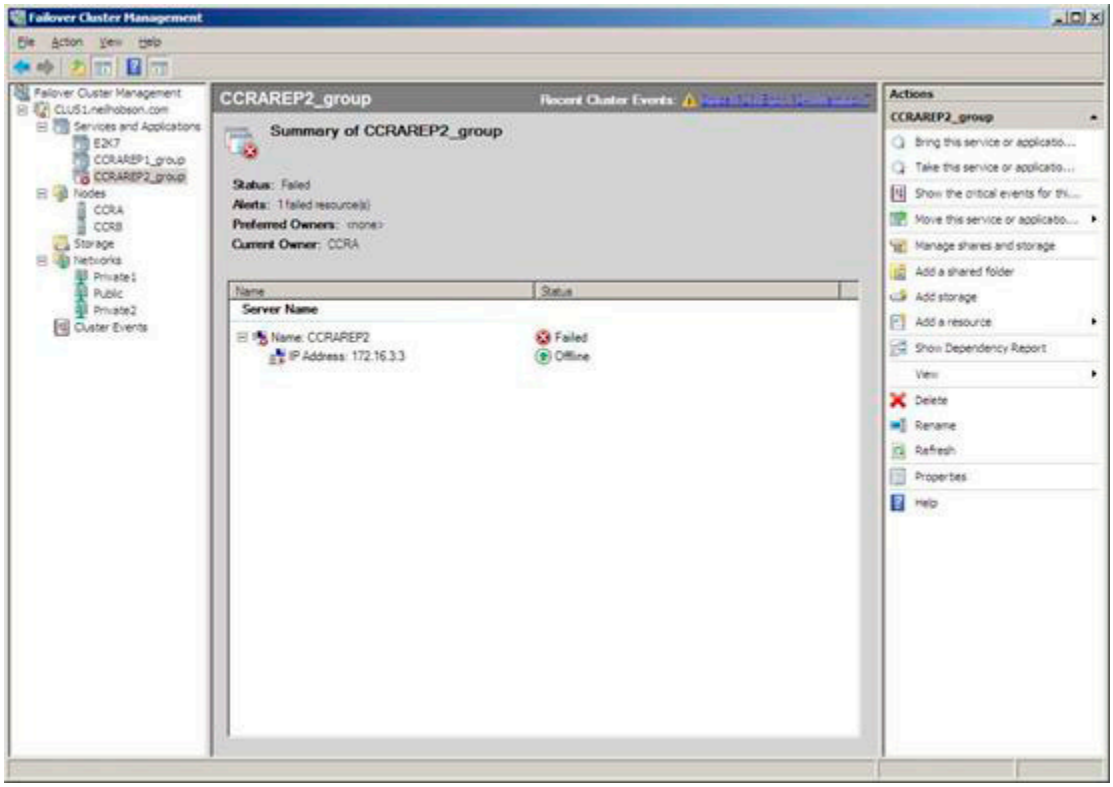

Figure 10: Resource Error in Failover Cluster Management.

Once NetBIOS has been enabled on the private network interface, the configuration process can continue successfully as shown in Figure 11 and Figure 12. Note that a new cluster group has been created for each continuous replication host name.

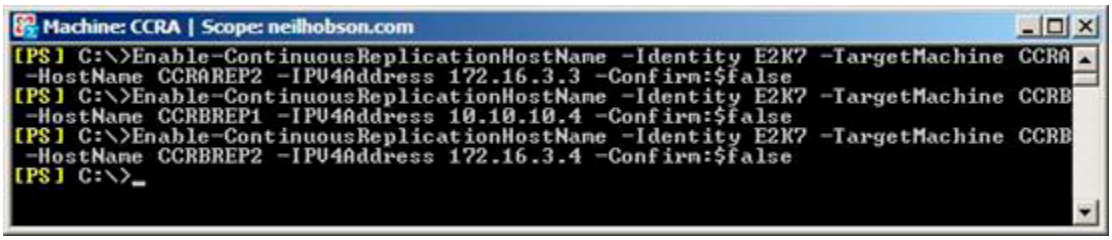

Figure 11: Enabling Continuous Replication Host Names.

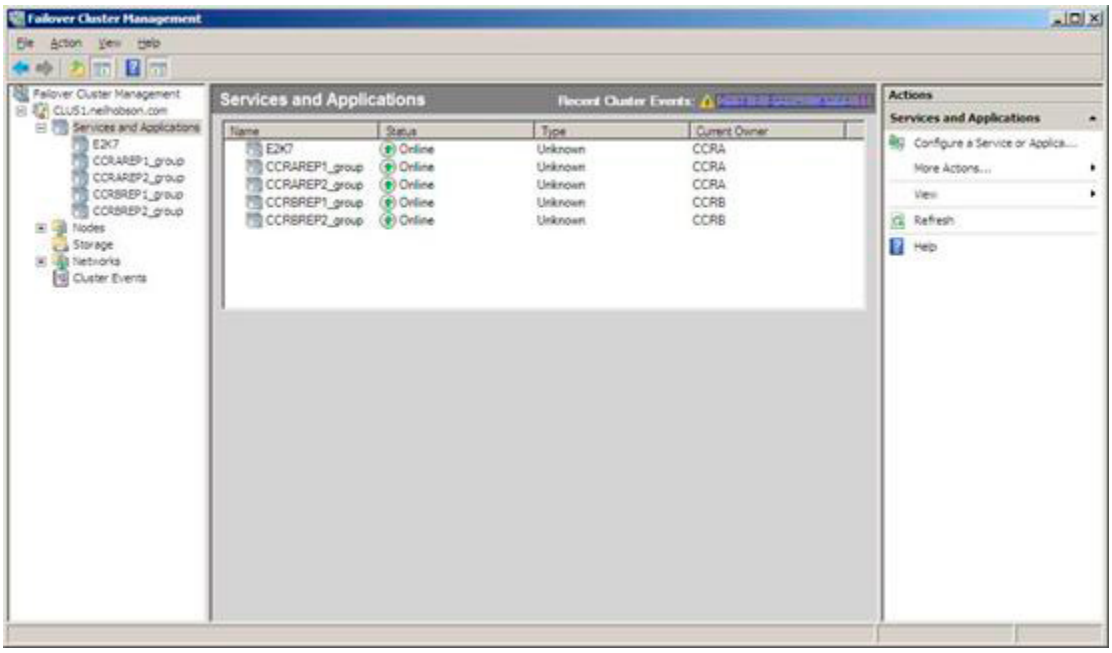

Figure 12: Continuous Replication Host Name Cluster Groups.

Log Shipping Using Continuous Replication Host Names

Now that continuous replication host names have been enabled, it will be satisfying to confirm that Exchange is actually making use of these networks for log shipping purposes, rather than using the public network. To do this, it's possible to examine the output of the *Get-ClusteredMailboxServerStatus* cmdlet as shown in Figure 13.

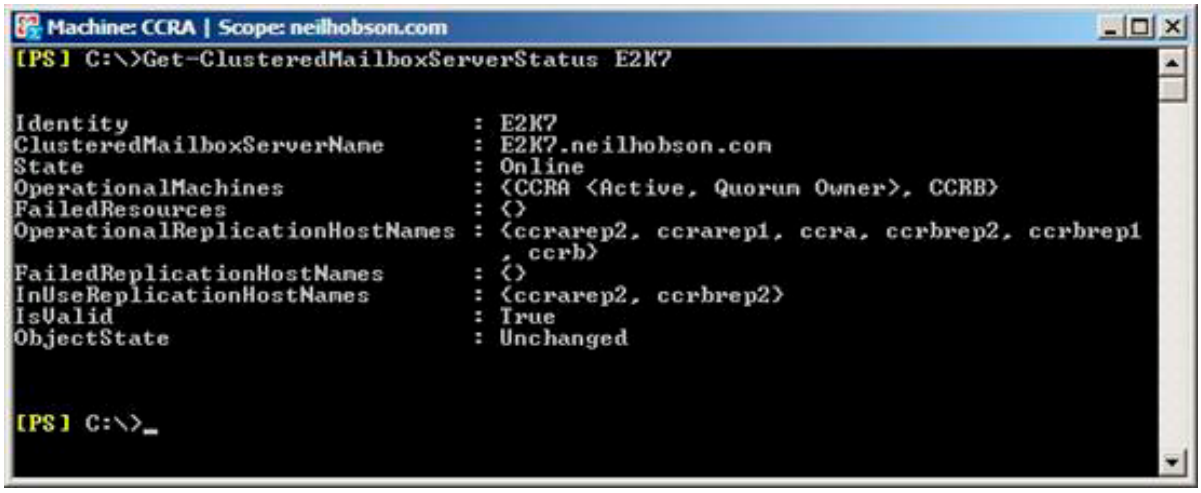

Figure 13: Get-ClusteredMailboxServerStatus Output.

In Figure 13 it can be seen that the OperationalReplicationHostNames parameter contains the full list of continuous replication host names available for the CCR environment to use for log shipping purposes. In the example configuration presented within this article, there are two redundant networks available for use and the Microsoft Exchange Replication Service is responsible for choosing between them. It can be seen from Figure 13 that the InUseReplicationHostNames parameter shows which of the continuous replication host names has been selected by the Microsoft Exchange Replication Service for use.

The Microsoft Exchange Replication Service checks the status of the networks every five minutes and adjusts the configuration as required. For example, if the private network associated with the continuous replication host names CCRAREP2 and CCRBREP2 is disabled, the Microsoft Exchange Replication Service then chooses the remaining redundant network as can be seen in Figure 14. Notice how the CCRAREP2 and CCRBREP2 host names are now classified as FailedReplicationHostNames.

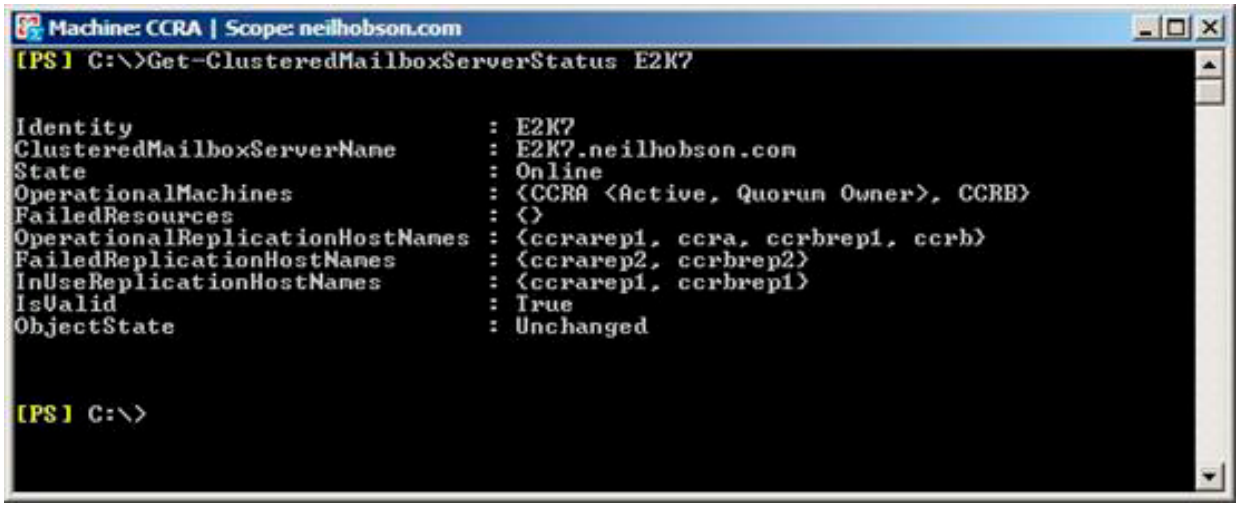

Figure 14: New Network in Use.

Database Seeding via Redundant Networks

There may be occasions where it's a requirement to re-seed one or more databases in a CCR environment. Naturally, re-seeding one or more large databases can cause a significant amount of network traffic to be generated and consequently it is desirable to ensure that this re-seeding process takes place using one of the redundant networks that have been configured.

To do this, it's possible to use the *DataHostNames* parameter of the *Update-StorageGroupCopy* cmdlet. With this parameter, a maximum of two redundant networks can be specified although note that only one network can be used at a time. In other words, it's not possible to perform a parallel seeding operation across two networks at the same time. In the scenario where it's a requirement to re-seed the default mailbox database in the default first storage group from cluster node CCRA to cluster node CCRB, the cmdlet that can be used to take advantage of one of the redundant networks is:

```
Update-StorageGroupCopy "E2K7\First Storage Group" -DataHostNames CCRAREP1
```

Summary

CCR is an excellent high availability solution that can be deployed with just two network interfaces in each cluster node. However, there's no doubt that the ability to deploy redundant networks using the continuous replication host name feature of Exchange 2007 Service Pack 1 greatly enhances the overall performance and high availability of a CCR environment. If you are about to implement a CCR environment, be sure to give due thought to implementing redundant networks for the purposes of log shipping and database seeding. In particular, take the time to decide whether you should be implementing multiple redundant private network interfaces as described within this article.

Building an Exchange Server 2007 environment

29 June 2009

by Jaap Wesselius

Of course, changing a 32,000 mailbox system, based in 40 Exchange Servers, to a centralised 25,000 mailbox solution is not a difficult task as long as you follow the Microsoft guidelines; It just takes patience, and time to understand Exchange's storage design principles. Jaap describes the design process and shares his experiences with the chosen server hardware.

I have been working on a project recently where we had to migrate 25,000 mailboxes from a decentralized Exchange Server 2003 with 40 servers to a centralized Exchange Server 2007 environment. Designing and building an environment like this is fun and in this article I will explain the approach we took for designing this environment. I will not cover the actual migration and decommissioning the Exchange Server 2003 platform, just the design phase, especially for the mailbox servers.

Exchange Server 2003

The original platform that the customer was using was built on Windows Server 2003 and Exchange Server 2003 and was in use for only 3 years. There were 40 Exchange Server 2003 mailbox servers spread across the country hosting around 38,000 mailboxes of which 12.000 inactives. There was one access point to and from the Internet and at this location two Exchange Server 2003 front-end servers were located. These servers were providing POP3 (just a few users), Outlook Web Access and ActiveSync for Windows Mobile users.

Before designing anything for Exchange Server 2007 the existing environment needs to be investigated. The following figures were found, Type I, II and III are the type of users (or usage scenario):

	Type I	Type II	Type III
msg sent/day/mbx	10	20	30
msg rec/day/mbx	31	67	100
Total msg per day	41	87	130
msg /mbx per day (MB)	6	17	44
msg /mbx per week (MB)	42	119	308
average msg size (KB)	150	200	350
Total MBX	23384	2632	863

Table 1. Usage scenario on Exchange Server 2003.

All users were using Outlook 2003. Desktop were running in online mode due to the roaming profile, laptops were running in cached mode. Outlook Web Access is used by only a small number of users and so is POP3. There were approximately 3,000 Blackberry device and approximately 1,000 Windows Mobile devices.

Exchange 2003 Storage Usage

Besides the usage profile of the current messaging environment it is also important to know how the current mailboxes are sized. And how many mailboxes are actually active. Using the Exchange Server 2003 current sizing, an appropriate storage design for Exchange Server 2007 can be made.

At the beginning of the project approximately 27,000 mailboxes could be separated into three categories:

Mailbox size	Number of users
Less then 500 MB	~23,000
Between 500 MB and 2 GB	~3,000
Between 2 GB and 8 GB	~1,000

Table 2. Mailbox limits and the number of mailboxes.

The customer wanted to maintain the mailbox limits in the Exchange Server 2007 environment.

As it turned out the mailbox types did follow the mailbox profiles we found earlier, in general a large mailbox had a higher IO profile compared to a mailbox which was of a smaller type.

Exchange Server 2007

One of the most important design aspects in Exchange Server 2007 is a proper storage solution. Without a proper storage solution there's a serious risk to end up with an Exchange environment that does not perform well.

To help customers design a proper storage solution Microsoft has created a tool to assist, the Microsoft Mailbox Storage Requirement calculator which can be found here: Exchange 2007 Mailbox Server Role Storage Requirements Calculator.

When you open the storage calculator you can enter variables like the number of users, the concurrency, the usage profile, etc., and the storage calculator will calculate the amount of storage, the number of I/O (Input/output) operations per second, etc. Please note that the storage calculator is updated regularly and that the output of the current version can differ from for example last year's version. The last version per June 2009 is V.16.9.

Using the storage calculator and the mailbox profiles and sizes the following numbers were found for sizing the Exchange Server 2007 mailbox servers:

Mailbox limit	< 500 MB	< 2 GB	< 8 GB
Number of mailboxes	23,000	3,000	1,000
Number of mailboxes per server	4,000	2,000	800
Number of servers needed	5	2	2
Max mailbox database size	200 GB	200 GB	200 GB
Number of mailboxes per database	400	100	25
Number of databases per server	11	15	20
Total Database IOPS per server	2,100	3,000	2,000
Total log file IOPS per server	580	820	550

Table 3. Exchange Server 2007 mailbox server sizing.

> ***Note***
>
> *The actual design was made in the summer of 2007 with an older version of the storage calculator, based on Exchange Server 2007 RTM. With SP1 and with a newer version of the storage calculator different values will be found.*

So now we know the total number of Exchange Mailbox servers: 9. The number of Hub Transport Servers and Client Access Servers can be derived from this number, or better, from the number of processors in the servers. For this customer no hard numbers were used though. The customer wanted to deploy Exchange Server 2007 into two separate rooms (in one building), so in each room 2 Hub Transport Servers and 3 Client Access Servers were used.

A Unix based anti-virus and anti-spam solution was already available so an Edge Transport Server is not used.

Hardware used for Exchange Server 2007

The customer was using a Standard Building Block (SBB) for their entire IT infrastructure. This SBB is based on a Dell PowerEdge 2950 with a dual Quad Core Intel Xenon processor, 2 local disks in a RAID-1 configuration, 16 GB internal memory and two Intel dual port network adapters. These servers are used for all Exchange Server 2007 servers. All servers are running Windows Server 2003 X64, Standard Edition or Enterprise Edition.

The storage being used for the Mailbox Servers is an iSCSI solution from Equallogic. The Equallogic PS-series should scale well and are very flexible, even in an enterprise environment like this. Well, so did the sales men tell us ;-)

Before this deployment I've seen Equallogic PS-series storage in several scenarios, but never beyond three or four units. This is quite a different environment where one would expect a fiber channel storage solution from EMC or HP, but not an Equallogic iSCSI solution.

It took us three months of calculating, testing (JetStress), debugging, implementing a new firmware (which had to be written by Equallogic), calculating again, testing again and implementing before we got the Equallogic environment up-and-running the right way, sufficient for servicing the environment I outlined earlier.

Because EqualLogic virtualizes storage and handles volumes differently that mainstream storage the exchange storage calculator has to be used differently. The important values are the IOPS which are needed and because the storage calculator still uses disk spindels to calculate requirements the values differ. Diskspindels do not matter (or at least less) in an EqualLogic design as all data is spread over 14 disks per member in the pool anyway. Also the controller logic within the Equallogic and the cache memory used in an Equallogic mean you have to interpret the numbers differently.

This is how we designed the storage solution:

- Two Equallogic portals were designed. A portal is the target where the iSCSI initiators connect to.
- Each portal has three so called pools. A pool is a logical separation of units (that use the same RAIDSet) where data can be written.
- Two pools are configured each with three EqualLogic PS-5000X arrays; each array is equipped with 300 GB SAS disks (25TB usable storage). These will be used for storing the mailbox databases.
- One pool is configured with one EqualLogic PS-5000E array; each array is equipped with 1TB SATA disks (12TB usable storage). This will be used for storing the log files from the various Storage Groups.

Because EqualLogic virtualizes storage a volume is spread over multiple arrays in a pool to a maximum of three possibly utilizing 42 disks giving maximum performance and IOPS.

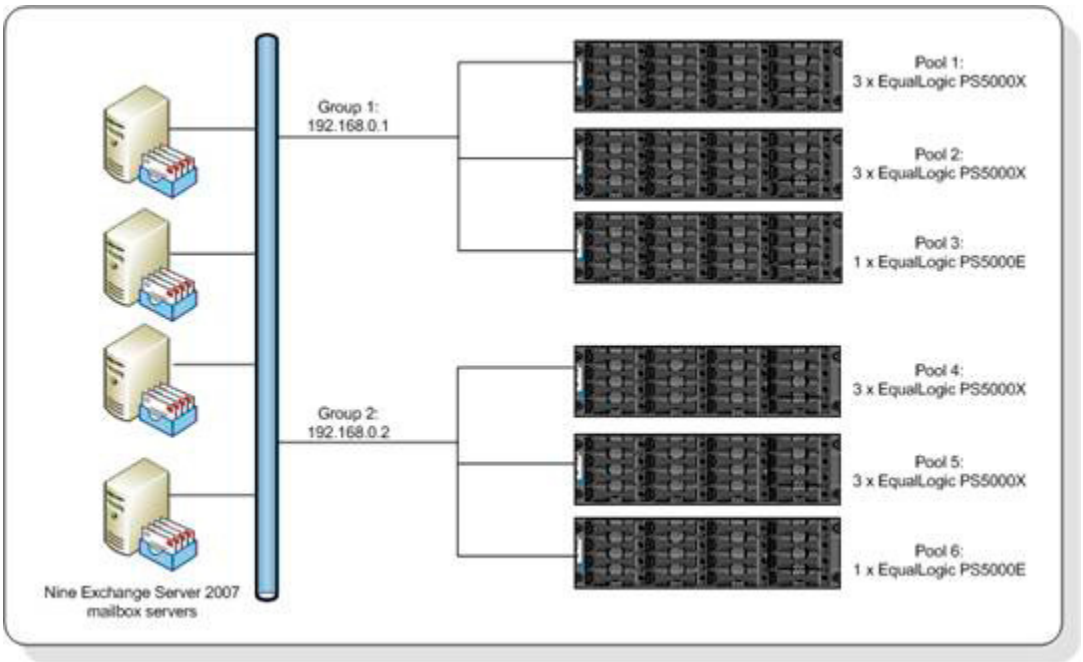

Figure 1. Schematical overview of the mailbox servers and the storage solution.

Note

In total 14 Equallogic units were used, but please not that since a CCR cluster is used each side of the cluster has to be equipped with identical storage. The passive side of the CCR cluster is equipped with 14 Equallogic units as well, so a total of 28 iSCSI units were used!

All units are configured in a RAID-50 scenario, this will give the best performance on these particular devices. In total a net capacity of approximately 50 TB of storage is created this way and this should be sufficient to host the 25,000 mailboxes and provide enough for future growth.

Mailbox Server configuration

So how is the mailbox server configured you might ask. All fourteen arrays are connected to a Cisco switch. The controllers are redundant; each controller has 3 active network interfaces and 3 passive interfaces. All servers are configured with two (Intel) GBit network adapters, each connected to the Cisco switch as well.

One problem we initially ran into was the amount of ISCSI connections we needed. Using two network cards in the server and three network cards in the Equallogic units resulted in six ISCSI connections per volume on the ISCSI portal. ISCSI connections are setup according the following rule; one connection per network interface on the server per array on which the data resides. When using 140 databases this means that 280 volumes are needed which resulted in 1680 ISCSI connections. The initial firmware (version 3.x) was not capable of handling this amount of connections (only 512 per portal), so we had to wait for the version 4.x firmware which was capable of handling this amount of connections (512 per pool). But still we had to use a separate pool for LOG volumes since one array only has 2 ISCSI connections per volume. This resulted in a total number of 560 ISCSI connections per portal instead of 860 ISCSI connections leaving enough space to grow.

A classic Windows limitation when using a large number of LUNs on a server is the drive letter limitation. There are only 26 characters in the alphabet, and 5 are already in use on the server, so only 21 are left. When configuring 20 Storage Groups 40 LUNs are used, so Volume Mount Points are configured on each mailbox server, resulting in only two drive letters being used. One for the databases and one for the LOG volumes.

When configuring an environment like this you really have to know how to script. Scripts were used to configure all the LUNs on the Equallogic environment, scripts were used to create all the partitions (be aware of the disk alignment!) and format the drivers. Best practice is a 64Kb disk offset which is not default in Windows.

Finally a PowerShell (Exchange Management Shell) script was used to create and mount all the databases on the Exchange Server 2007 mailbox servers.

For backing up the Exchange environment a Microsoft System Center Data Protection Manager (DPM) 2007 solution was built. DPM is using a disk-to-disk-to-tape backup solution. A full backup is created every day and an incremental backup is created every two hours. Every week an entire full backup is written to tape for long term, off-site storage.

Performance

The performance of the Equallogic PS-series really surprised me, they were performing better than I initially expected. During testing we manged to get 35,000 IOPS from the storage which was a lot more than we needed. After migration 25,000 mailboxes to the new platform all nine mailbox servers are still running fine with a pretty low number of disk I/O's and a fairly low processor utilization.

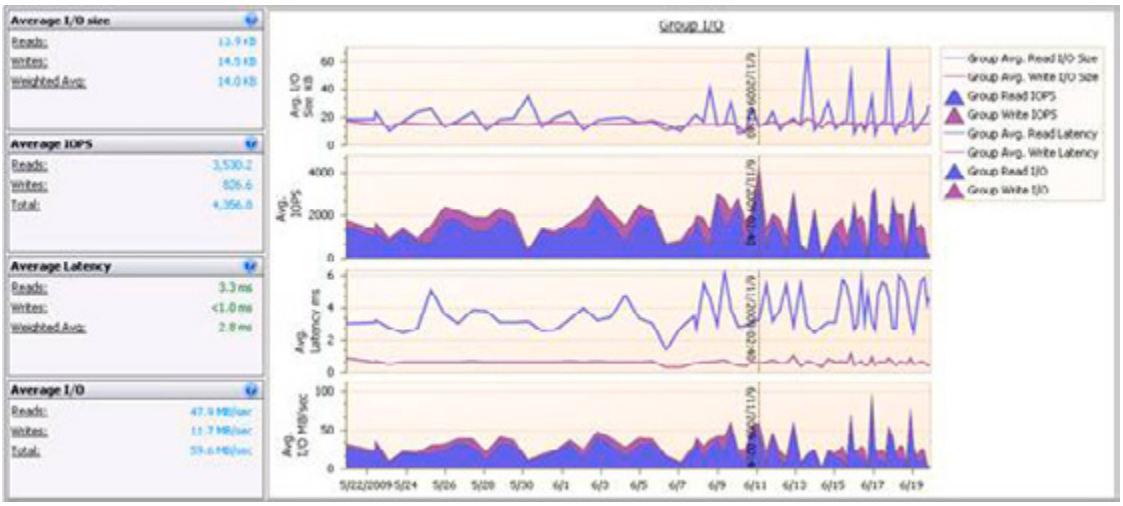

Are there challenges? Oh yes, think about the Online Maintenance. When using such an amount of databases and this sizing OLM really has a hard time to finish in time, so only two databases are configured each night to run the OLM. Each database is running OLM once every week; otherwise it cannot finish its run.

Reseeding a database is nothing else than a file copy over the network from the active node in the CCR cluster to the passive node. Copying a 150 GB or 200 GB mailbox database can take some time; imagine you have to reseed five databases.

You also need a good performing backup solution to have the backup finish within the available timeframes. Also monitoring becomes more important when dealing with these amounts of databases.

Conclusion

Designing a 25,000 mailbox solution is not a difficult task as long as you follow the Microsoft guidelines; use the Exchange Storage Calculator for designing a proper storage solution and understand exchange storage design principles. Failing to do so will always result in a poor performing Exchange server environment.

Using a storage solution from Dell/Equallogic is also possible, but special care has to be taken to design and implement the members the correct way, using the right number of portals, pools and members. If you do so, it will perform well. Please note that in this scenario the Equallogic solution was only used for storing the Exchange Server mailbox databases and logs. No other data, such as like SQL databases, is stored on the EqualLogic arrays.

Disclaimer

This article is a true scenario, but the entire project from the Exchange Server 2003 inventory up to decommissioning the Exchange Server 2003 servers took more than a year with three people, working full time on all aspects. The documentation that was created during the project covers more than 300 pages, so it is not possible to cover everything in great detail. My main objective was to explain a bit more about using the Equallogic hardware in combination with a 25,000 mailbox environment.

An Introduction to Messaging Records Management

07 July 2009

by Ben Lye

There are a number of features in Exchange that can be used in creative ways to solve problems in Exchange as they crop up. Ben Lye recently found a great use for Messaging Records Management when he hit the problem of a meeting-room mailbox that stopped replying to booking requests. Here he shows how to apply policies that apply to all folders of a particular type, and how to schedule a regular task to do it.

Messaging Records Management (MRM) is a feature of Exchange 2007 which lets administrators define policies to control the storing of e-mail data. Essentially MRM allows you to create folder-level rules which determine how long messages stored in that folder should be retained. MRM policies can be applied to default mailbox folders or custom mailbox folders, however applying policies to custom folders requires the use of Enterprise CALs.

MRM can be used to ensure compliance with company policy or legal requirements, or other regulations. MRM can also be used as an administrative tool.

Recently I had to fix an issue with a meeting room mailbox which had stopped replying to booking requests. It turned out that the mailbox was over the database's default Prohibit Send quota, so was unable to send any messages. The problem was that all the old meeting requests, which had been deleted after they had been processed by the Calendar Attendant, were still in the Deleted Items folder, and that was causing the mailbox to exceed its quota.

The quick-fix was to empty the deleted items folder, but this didn't really fix the root cause of the problem - the mailbox would continue to grow over time. Another option was to apply a higher quota, but this would just make the problem come back again later when the new quota was reached. Ideally I wanted to find a way to make sure that deleted messages in resource mailboxes were permanently removed.

This is where MRM comes in - as well as applying policies for storing e-mail messages, MRM can be used to apply policies to *delete* e-mail messages. To solve the problem with my resource mailbox I created an MRM policy which would automatically remove messages from the deleted items folder of all resource mailboxes.

There are five steps to setting up an MRM policy:

- Choose the folder you want the policy to apply to.
- Create the managed content settings for the folder.
- Create a managed folder mailbox policy.
- Apply the managed folder mailbox policy to the mailboxes.
- Schedule the managed folder assistant to run.

To set up the MRM settings and policies you must be delegated the Exchange Organization Administrator's role.

I already know that I want to apply the policy to the Deleted Items folder; however I want to create a new managed folder so that if there are future needs for different managed content settings on the Deleted Items folder in other mailboxes, they can be accommodated. Because the Deleted Items folder is a default folder I can use managed default folders.

To create a new managed default folder we use the **New-ManagedFolder** cmdlet

```
New-ManagedFolder -Name "Deleted Items (30-day limit)" `
-DefaultFolderType "DeletedItems"
```

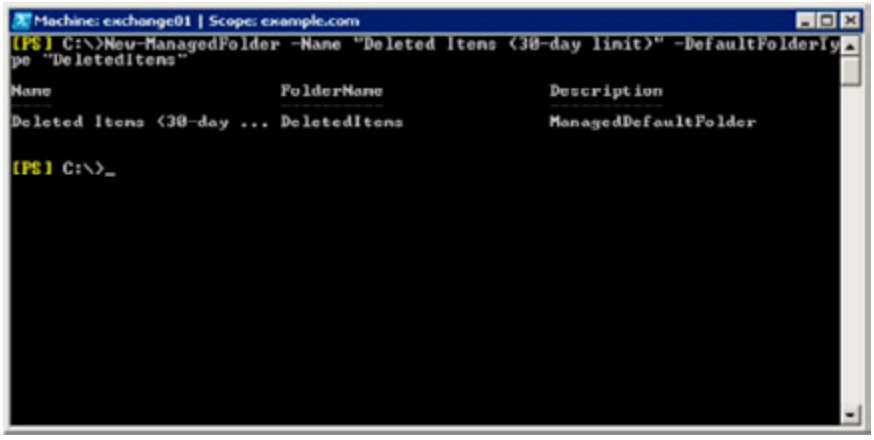

With the new managed folder created the next thing to do is create the managed content settings. The managed content settings define the actual retention criteria.

Creation of the managed content settings can be done through either the Exchange Management Console or the Exchange Management Shell.

I want settings which will remove items from the deleted items folder 30 days after they were received by the resource mailbox. The deleted messages should be available for recovery according to the database's standard deleted item recovery policy.

To create the managed content settings we will use the **New-ManagedContentSettings** cmdlet:

```
New-ManagedContentSettings -Name "Remove Deleted Items After 30 Days" `
-FolderName "Deleted Items (30-day limit)" `
-MessageClass * `
-RetentionEnabled $true `
-RetentionAction DeleteAndAllowRecovery `
-AgeLimitForRetention 30 `
-TriggerForRetention WhenDelivered
```

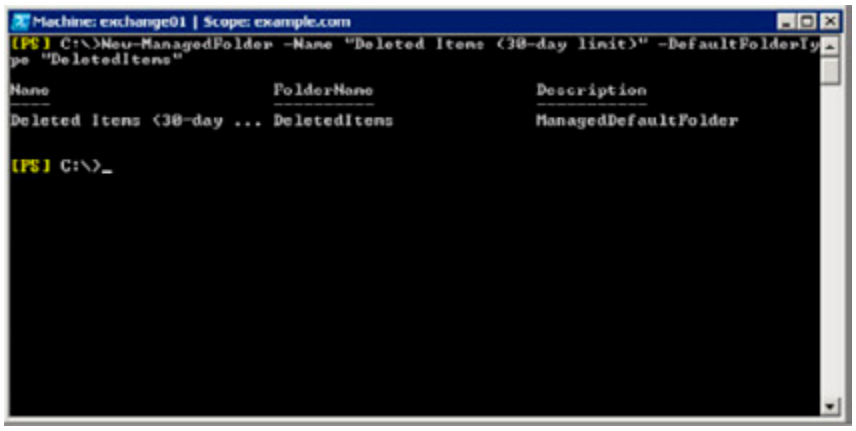

Because a mailbox can only have one managed folder mailbox policy applied managed folder mailbox policies are used to group managed folder content settings so that multiple settings can be applied to a mailbox. The next step is to create the managed folder mailbox policy.

To create a new policy containing the new deleted item managed content settings we use the **New-ManagedFolderMailboxPolicy** cmdlet:

```
New-ManagedFolderMailboxPolicy -Name "Resource Mailbox Policy" `
```

```
-ManagedFolderLinks "Deleted Items (30-day limit)"
```

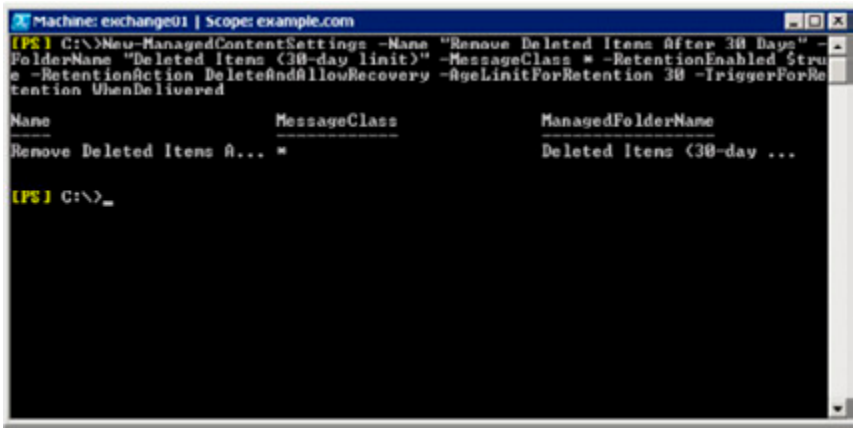

The next step is to apply the new managed folder mailbox policy to the resource mailboxes using the **Set-Mailbox** cmdlet (once for room resources and once for equipment resources). You will be warned that Outlook clients prior to Outlook 2007 do not support all MRM features, and clients older than Outlook 2003 SP2 are not supported):

```
Get-Mailbox -RecipientTypeDetails RoomMailbox |
Set-Mailbox -ManagedFolderMailboxPolicy "Resource Mailbox Policy"
Get-Mailbox -RecipientTypeDetails EquipmentMailbox |
Set-Mailbox -ManagedFolderMailboxPolicy "Resource Mailbox Policy"
```

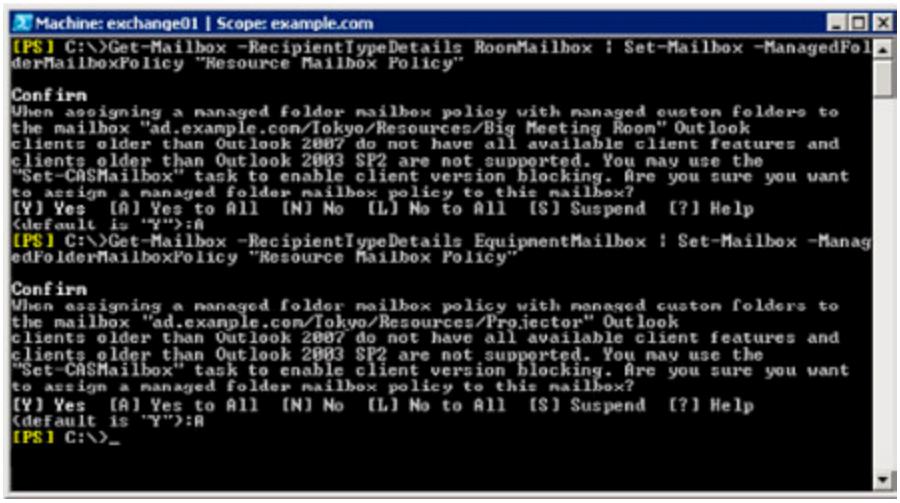

The final step is to schedule the managed folder assistant (the process which actually enforces the policy) to run. This is done with the **Set-MailboxServer** cmdlet.

Each Exchange server only has one schedule for the managed folder assistant and you should schedule it to run as often as needed to meet you requirements. To run the managed folder assistant on all Exchange 2007 mailbox servers daily between 1am and 3am the command would be:

```
Get-ExchangeServer |
Where { $_.AdminDisplayVersion.ToString().SubString(0, 10) -eq "Version 8." `
-and $_.ServerRole -eq "Mailbox" } |
ForEach { Set-MailboxServer -Identity $_.Identity `
-ManagedFolderAssistantSchedule "Sun.1:00 AM-Sun.3:00 AM," `
"Mon.1:00 AM-Mon.3:00 AM," "Tue.1:00 AM-Tue.3:00 AM," `
"Wed.1:00 AM-Wed.3:00 AM," "Thu.1:00 AM-Thu.3:00 AM," `
"Fri.1:00 AM-Fri.3:00 AM," "Sat.1:00 AM-Sat.3:00 AM" }
```

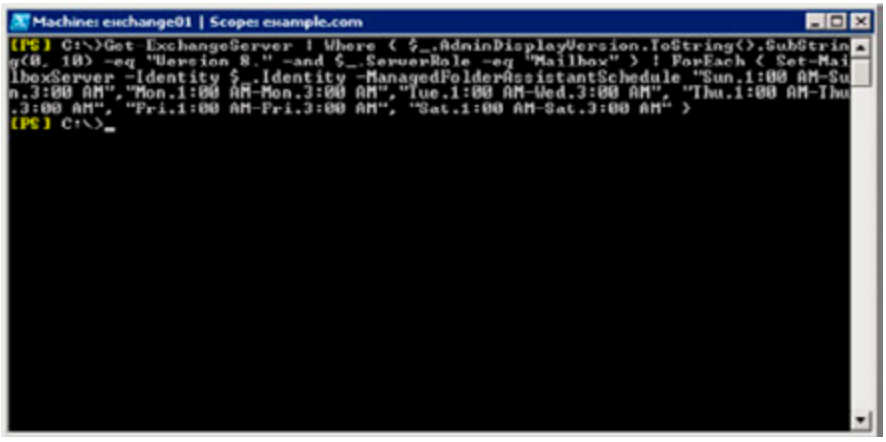

The managed folder assistant can also be started on demand using the **Start-ManagedFolderAssistant** cmdlet. Running the managed folder assistant can be a resource-intensive process, and it should be scheduled to run during off-peak hours.

The application of this MRM policy will ensure that my resource mailboxes will remain small, as the content of their deleted items folders will be automatically removed.

More information on Messaging Records Management and messaging policies can be found in Microsoft TechNet:

http://technet.microsoft.com/en-us/library/aa998599.aspx

Installing Hyper-V and Beyond

28 July 2009

by Jaap Wesselius

> ***In his previous article "Hyper-V, an introduction" Jaap Wesselius explained about the Hypervisor, the parent partition, the child partition, and Integration Components. In this article Jaap discusses installing Hyper-V, all kinds of Virtual Hard Disks, Virtual Networks, and some best practices.***

Installing Hyper-V

Hyper-V is a Server Role within Windows Server 2008. This means that you have to install Windows Server 2008 before installing Hyper-V. The hardware requirements for installing Hyper-V are:

- The processor needs to support hardware virtualization (AMD-V or Intel VT).
- Hardware enforced Data Execution Prevention (DEP). The Intel XD bit needs to be enabled or the AMD NX bit needs to be enabled.
- The processor needs to be an X64 processor.

Furthermore you need plenty of physical memory in the server. For testing purposes the absolute minimum is 4 GB of memory, the amount of memory needed in a production environment is dependant of the services you want to run.

Windows Server 2008 Hyper-V is a Server Role in the Standard Edition, Enterprise Edition and the Datacenter Edition. From a technical perspective there's no difference, it's a licensing issue:

- Install Standard Edition and you are allowed to run one Virtual Machine using the same license.
- Install Enterprise Edition and you are allowed to run four Virtual Machines using the same license.
- Install Datacenter Edition and you are allow to run as many Virtual Machines as the server can handle.

> ***Note***
>
> *There's also a Hyper-V Server available. This is a very small Operating System (command-line based) that's only capable of running the Hypervisor, nothing else. There are a few graphical UI's, like the date and time applet, but it's mainly command-line. Hyper-V Server is a free download, but when configuring Virtual Machines you need one license for each Virtual Machine!*

Install Windows Server 2008 using your corporate standards and bring it up to date with the latest hotfixes or service packs. Windows Server 2008 is now running on "bare metal" and there's no difference with any other Windows Server.

Logon to the server and open the Server Manager. In the left pane select "Roles" and in the right pane you'll see "Roles: 0 of 18 installed." Click on "Add Roles" and select "Hyper-V" in the Select Server Roles Window.

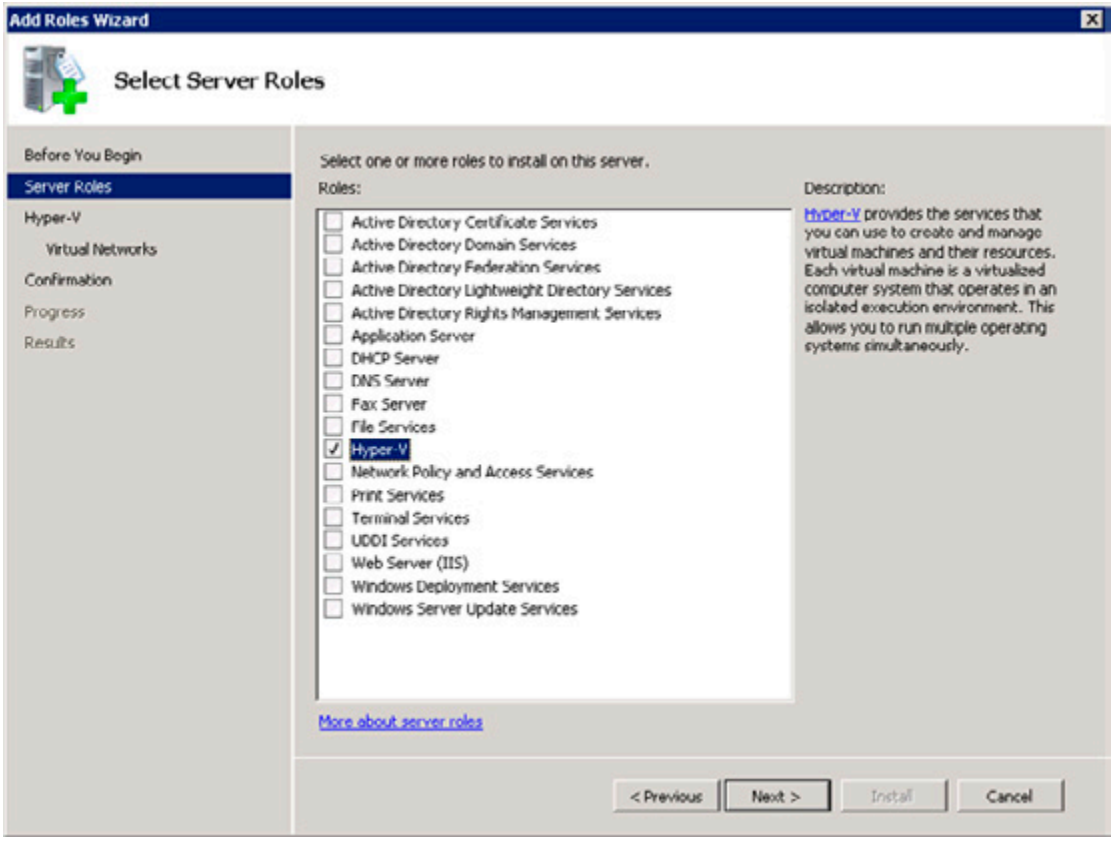

Figure 1. Need to create a new image.

Click Next to install the Hyper-V Server role. The wizard also includes the possibility to create Virtual Networks. We'll cover that later in this article so leave all options blank and click Next to continue.

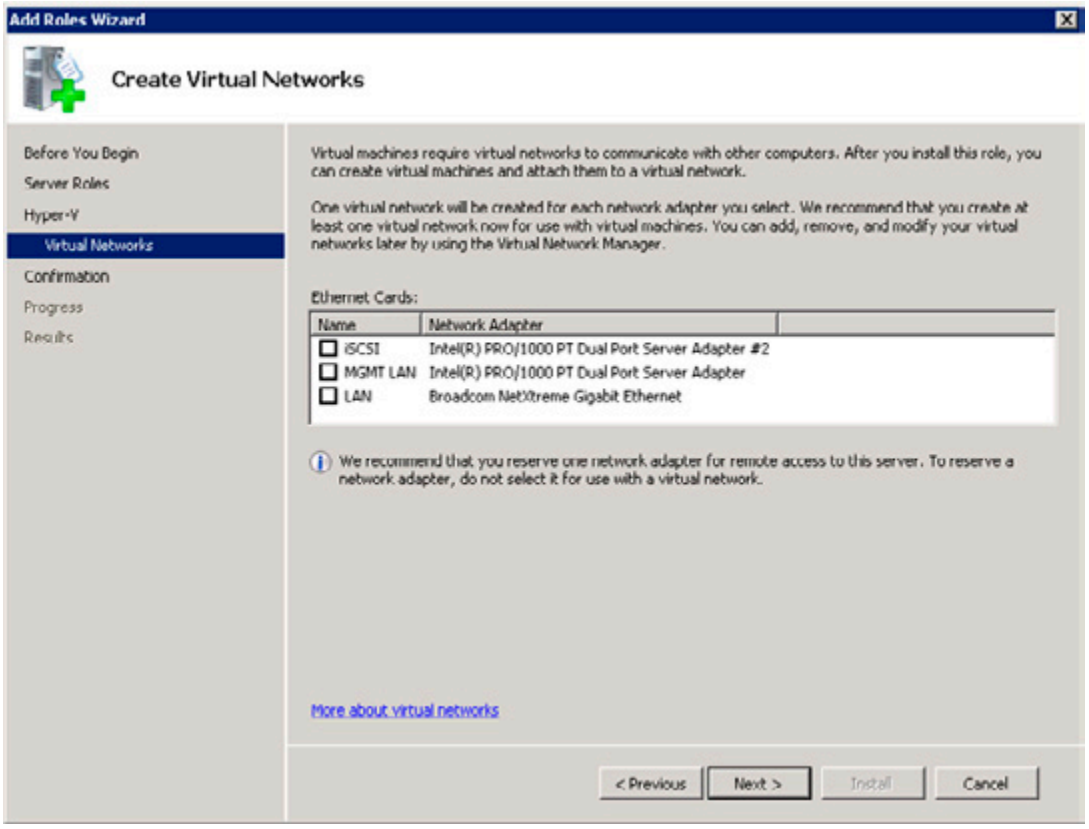

Figure 2. In the Add Roles wizard there's the possibility to create Virtual Networks. You can create these later as well.

Right now the actual Hypervisor "slides" under the Operating System, making the original Windows Server 2008 instance a Virtual Machine. Also the VM Worker processes (responsible for emulating devices), some WMI interfaces, the VM Bus and the Virtual Storage Provider (VSP) are installed.

When finished the server needs to be rebooted. During the reboot process the Hypervisor will be started initially which will hand-over the boot process to Windows Server 2008, which is now the Parent Partition. The Parent Partition is also referred as the Root Partition.

When you logon after rebooting you'll see nothing special, it just looks like Windows Server 2008, and it is. The only new part is the Hyper-V Manager that can be found in the "Administrative Tools" menu.

The Hyper-V Manager is an MMC snap-in, with a server pane (left), a results pane (middle) and an actions pane (right). In the server pane you'll see only the server that you're logged on to, additional Hyper-V servers can be added later on.

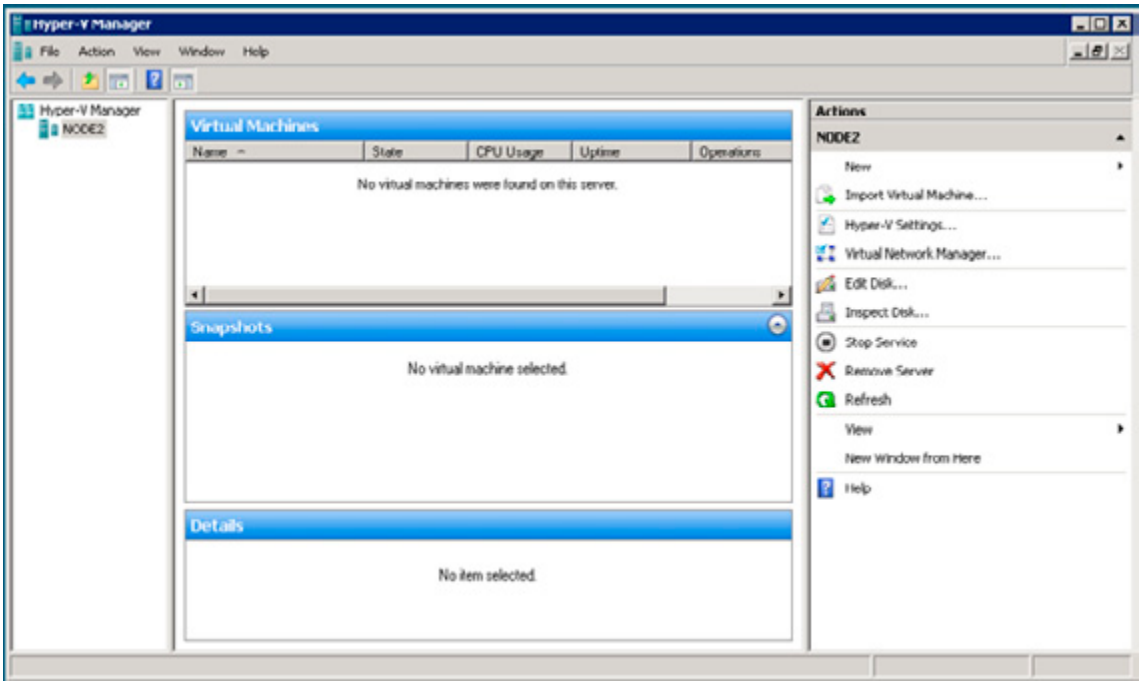

Figure 3. The Hyper-V Manager just after installing Hyper-V.

One of the first things that has to be done is changing the settings. Click on "Hyper-V Settings..." in the actions pane and enter new paths for the Virtual Hard Disks and the Virtual Machine. By default these will be placed on the C:\ drive, but from a performance perspective these should be configured on a separate drive. This can be a physical disk in the server, a storage cabinet attached to the server or a SAN (iSCSI or Fiber).

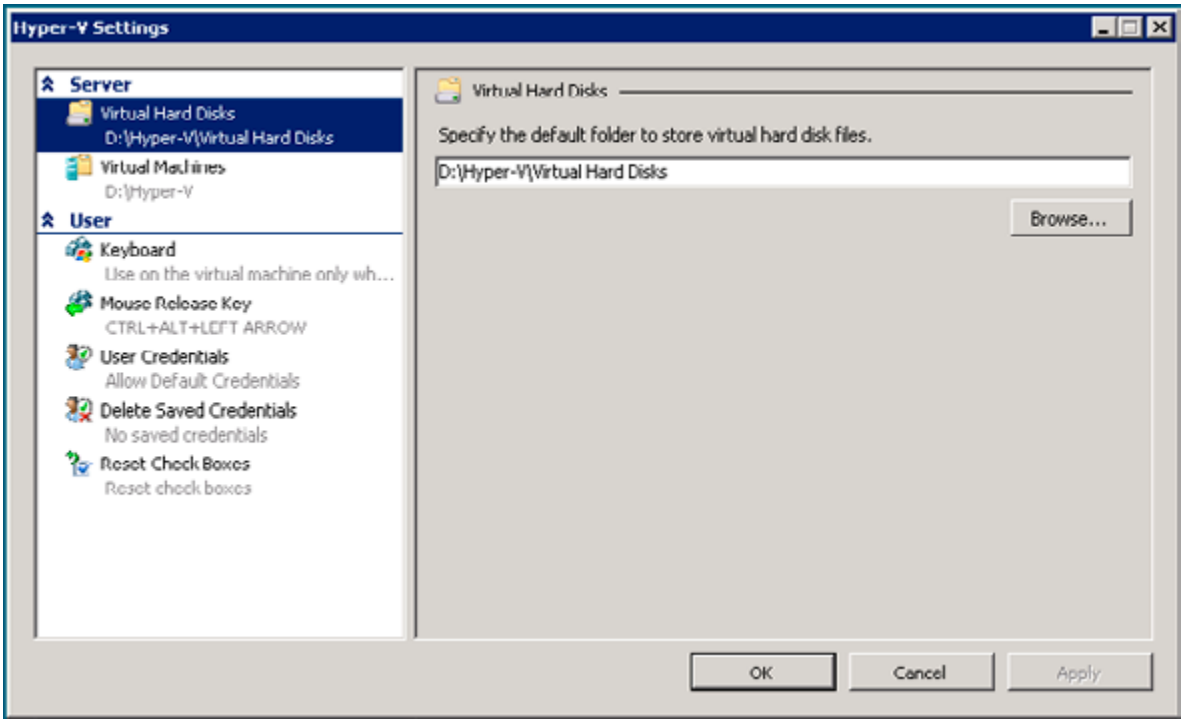

The Virtual Machines are using the D:\Hyper-V directory in this example. This means that the configuration files will be placed in this directory. But also snapshots will be placed in this directory. Therefore it is a best practice not to place the Virtual Machines on the C:\ drive. When using snapshots extensively this disk will fill up rapidly with snapshot information.

Virtual Hard Disks

Hyper-V uses Virtual Hard Disks for storing data. A Virtual Hard Disk is stored as a .VHD file on the disk of the Hyper-V server. There are three types of Virtual Hard Disks:

- **Dynamically Expanding Disks** – The dynamically expanding disk is a VHD file that's initially very small (2 MB) and that grows on demand. When more space is needed Hyper-V will automatically assign more space to this VHD. A 50 GB disk will start as a 2 MB file, but the Virtual Machine will see it as a 50 GB disk. This is perfectly suited for test- and development environments.

- **Fixed Size Disks** – The fixed size disk is a VHD file that has the complete size allocated before it can be used. A 50GB disk also means a 50 GB VHD file. Pre-allocating 50 GB for a Virtual Disk can take a considerable amount of time. Microsoft recommends using fixed size disks for using Virtual Machine in a production environment since it does not have the overhead of growing the VHD file.

- **Differencing Disks** – Differencing disks consists of two virtual hard disks. One virtual hard disk is designated as a read-only disk, changes made by the Virtual Machine are written to the second disk. This will allow us to create a "master image" and use several Virtual Machines based on the master image. This is a perfect solution for creating a test environment within minutes.

There's also a fourth type of disk, the pass-through or dedicated disk. This is a dedicated disk on the Parent Partition that's connected to a disk controller of a Virtual Machine. The complete disk is used in this configuration without the VHD file. Therefore it has no overhead. Besides the fixed size disk Microsoft also recommends using this kind of disk in a production environment. Check the Technet site for more information.

The pass-through disk can be a physical hard disk in the Hyper-V host, but it can also be a LUN on an iSCSI or Fiber Channel SAN.

To create a new Virtual Hard Disk open the Hyper-V Manager, in the Actions Pane click New and select "Hard Disk…." The New Virtual Hard Disk Wizards shows up with a welcome screen. Click Next to continue. In the Choose Disk Type window you can select what type of VHD needs to be created. Please be aware that creating a large fixed VHD file will take a considerable amount of time!

I will cover a dedicated or pass-through disk later in this article, after the creation of a Virtual Machine.

Virtual Networks

By default Virtual Machines are not connected to any network, so they are pretty useless actually. To connect Virtual Machine to each other, to the Parent Partition or to the network outside the Hyper-V host Virtual Networks need to be used.

Three types of Virtual Networks are available:

- **Private Virtual Network** – this is a virtual network that's only available for Virtual Machines running on the Hyper-V server. There's absolutely no possibility for the VM's to connect to the outside world, or to the Hyper-V server.

- **Internal Virtual Network** – This is a virtual network that's available for the Virtual Machines, but also for the Parent Partition. VM's can connect to the Parent Partition using the Internal Virtual Network. When Internet Connection Sharing (ICS) is enabled on the Parent Partition the VM's can connect to the outside world using this connection.

- **External Virtual Network** – this is a virtual network that's connected to the network card on the Hyper-V server. Each Virtual Machine is capable using this network card to connect to the outside world. This also means that other computer on the network can "see" the Virtual Machines just as regular computers.

It is a Microsoft recommendation to use multiple network interfaces. One interface should be used for management purposes and be only available for the host. One network interface should be used for an External Virtual Network so that VM's can access the physical network as well. If you're using an iSCSI solution an additional network interface for this should be used as well. Check the following Technet article for more information: 7 Best Practices for Physical Servers Hosting Hyper-V Roles.

So, for having a Virtual Machine communicating with the outside world we have to create an External Virtual Network. Logon to the server and open the Hyper-V Manager. In the tasks pane select the Virtual Network Manager.

Click Add, enter a name like "Public Virtual Network" and select the network interface that needs to be used by this Virtual Network. In this example this is the first Broadcom that's available.

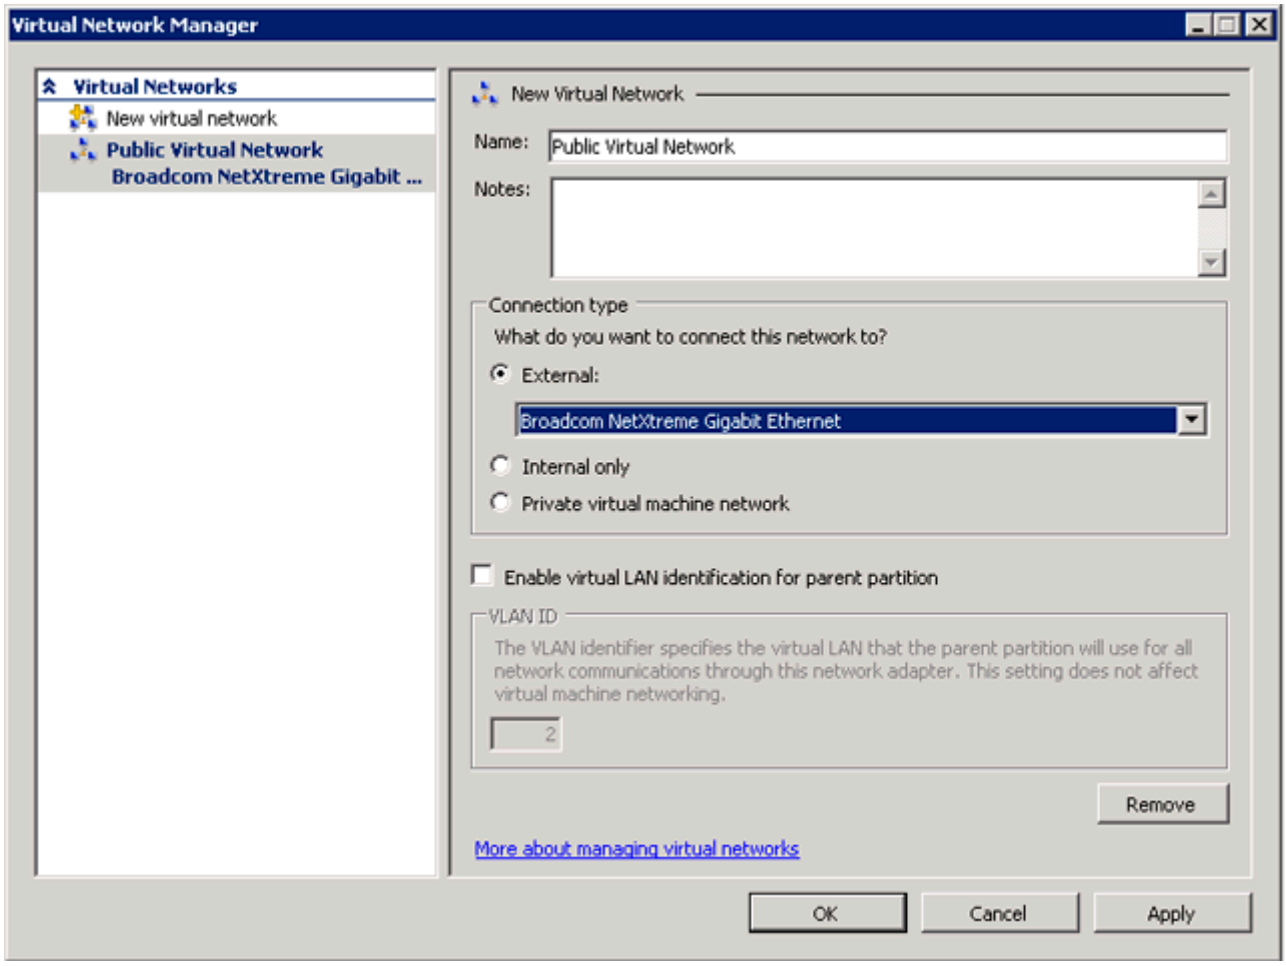

Figure 4. During creation of an External Private Network you have to select the physical network interface.

When you click OK the new Virtual Network will be created. A warning message will pop-up indicating you may temporarily lose network connectivity with the Hyper-V host. If you're connected via a Remote Desktop sessions you may have to re-establish the connection.

When the new Virtual Network is created and you check the Network Connections you'll notice a new entry, this is the new Virtual Network.

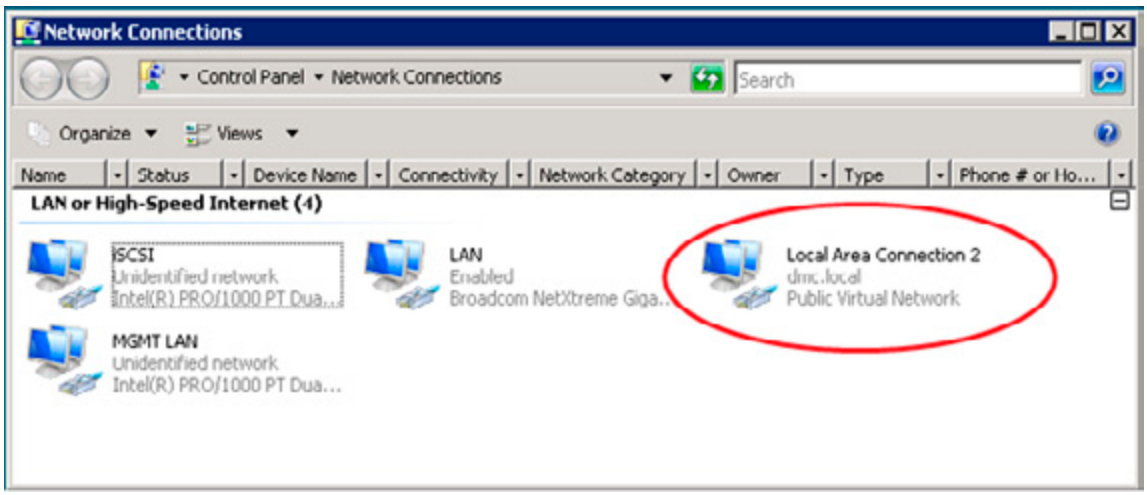

Figure 5. The new Network Connection after the creation of a External Virtual Network.

When you create a Virtual Machine and connect it to this new Virtual Network the Virtual Machine will be able to access resource on the public network.

Virtual Machines

On Hyper-V there are two types of Virtual machines, Supported and Unsupported Guest Operating Systems. This has nothing to do with the ability to run in a Hyper-V Virtual Machine, but whether it is directly supported by Microsoft:

- **Supported Guest Operating Systems** – these are Guest Operating Systems that are fully supported running on a Hyper-V server where the Guest Operating System can fully utilize the Hyper-V Infrastructure like the VMBus. Examples of these Operating Systems are Windows Server 2008, Windows Server 2003 SP2, Windows 2000 Server with SP4, Windows Vista SP1, Windows XP Professional SP2 and SUSE Linux Enterprise Server 10 with Service Pack 2. For a complete list of all supported Operating Systems check the Microsoft site: http://support.microsoft.com/kb/954958. Guest operating systems that are supported on a Hyper-V virtual machine

- **Unsupported Guest Operating Systems** – these are Guest Operating Systems that can run on a Hyper-V server but do not utilize the Hyper-V infrastructure. These are running in an emulated environment and cannot use the native Hyper-V architecture. This is the VM Worker process. An emulated environment is running in User Mode in the Parent Partition. When a VM makes a call to for example a disk or a network it goes through the VM, to the emulator and then to the actual hardware. This results in a serious I/O penalty, resulting in a slower performance.

Suse Linux Enterprise Server is a Supported Guest Operating System. This will run under Hyper-V and if you encounter an issue you can call Microsoft Product Support. SCO Unix is an Unsupported Guest Operating System. This will run perfectly under Hyper-V but it is not supported by Microsoft, but fully supported by SCO. For more information check the site of SCO.

Now we are going to create a Virtual Machine running Windows Server 2003 SP2. You can use the Windows installation CD or you can use an ISO image. The latter is usually faster.

Logon to the Hyper-V server and start the Hyper-V Manager. In the Actions pane click New and select Virtual Machine. The New Virtual Machine Wizard will show up, click Next to continue. Enter a name for the new Virtual Machine, for example "Windows Server 2003 X64."

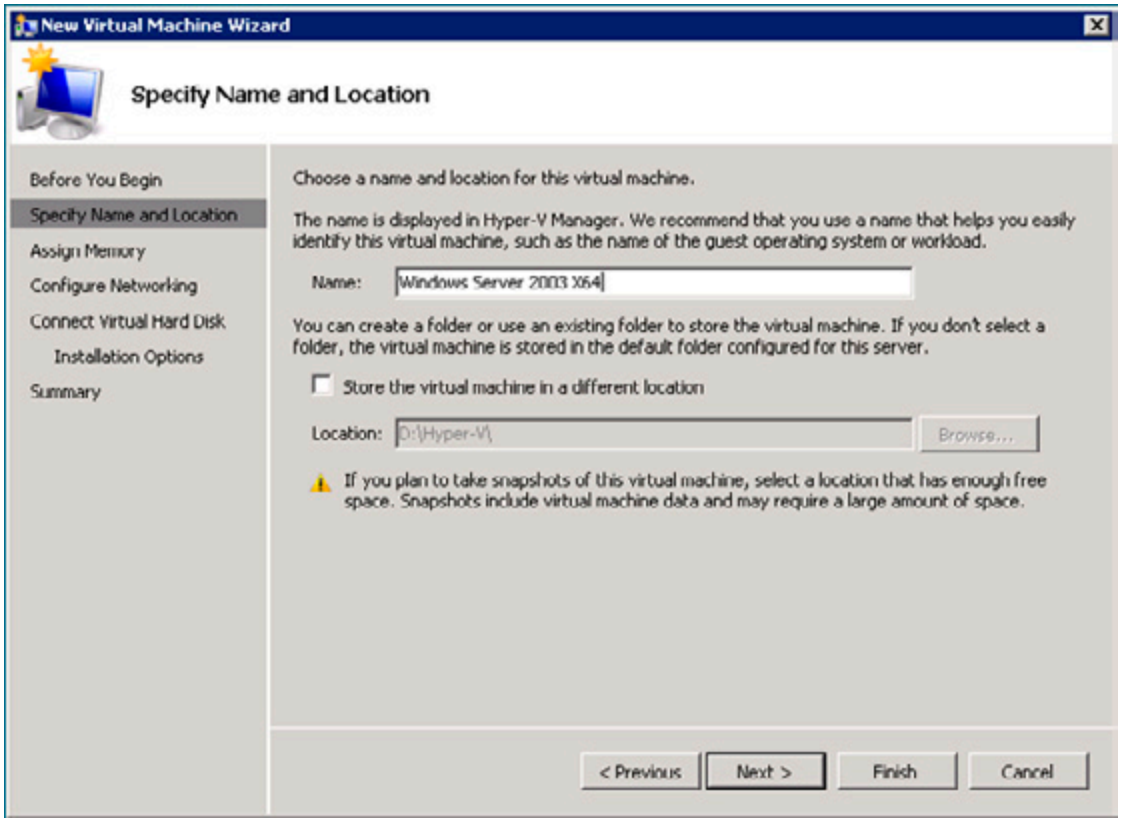

Figure 6. The New Virtual Machine Wizard when creating a new Virtual Machine.

The location will be same as we configured in the Hyper-V settings earlier in this article. You can check the "Store the virtual machine in a different location" checkbox when you want to have the new Virtual Machine on another location (i.e. another disk). Click Next to continue.

Enter the amount of memory that you want to assign to the VM. Please not that there's no real difference with a physical machine. If you fail to assign enough memory the new Virtual Machine will be slow and sluggish. Another thing to remember is that unlinke VMWare, there's no way to overcommit memory. If you have 16 GB of RAM available, you can only assign 14 Virtual Machines 1 GB of RAM (assuming approximately 2 GB is used by the Parent Partition).The next stop is to connect the new Virtual Machine to a Virtual Network. Connect the Virtual Machine to the "Public Virtual Network" that we created earlier.

The new Virtual Machine needs a Virtual Hard Disk (VHD) to use. Since we did not create a VHD in advance use the default setting "Create a Virtual hard disk."

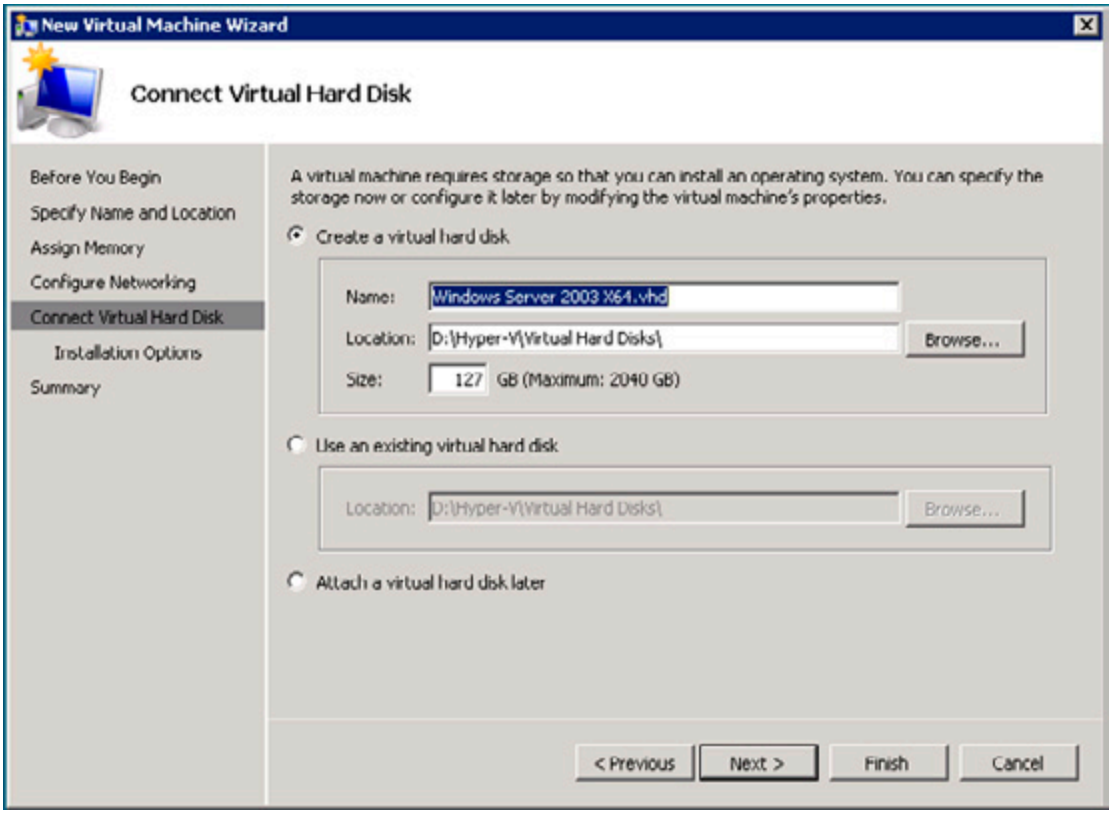

Figure 7. A new Virtual Hard Disk will be created.

If we created a Virtual hard Disk in advance, for example a fixed size Virtual Hard disk you have to select "use an existing virtual hard disk" and select the Virtual Hard Disk in the Location field. If you're using disk intensive applications like SQL Server, MySQL or Exchange Server fixed disks or pass-through disks should be used for best performance. Click Next to Continue.

Since we will be installing a new Operating System we have to select the installation media. This can either be a CD/DVD or an ISO image, whichever your prefer.

Click Finish to end the New Virtual Machine Wizard. The new Virtual Machine is now ready to be booted and we can continue to install the Operating System in the Virtual Machine.

Double click on the new Virtual Machine in the Hyper-V Manager and press the "Virtual Power button" to boot the new Virtual Machine. Install Windows Server 2003 as you are used to...

Note

If you are accessing the Hyper-V server using Remote Desktop then it's not possible to use your mouse in the new Virtual Machine until you've installed the Integration Components. And Integration Components can only be installed after finalizing the setup of Windows, so the setup itself can be challenging. When you have access to the console itself then it's no problem and the mouse works as expected.

Finishing the installation

When the setup of Windows Server 2003 in our example is finished you have a new Virtual Machine that cannot do anything. When you check the Windows installation for example you'll notice that it doesn't have a network interface. The Virtual Machine is running in an "isolated box" and cannot do anything. The first thing you have to do is install the Integration Components. Remember my previous article with the Virtual Storage Provider (VSP), the Virtual Storage Client (VSC) and the VM Bus? The Integration Components contain the software that's needed for working natively with Hyper-V. Open the Virtual Machine, select Action in the Virtual Machine menu and select "Insert Integration Services Setup Disk." The Integration Components wizard will be started and the software will be installed. When finished reboot the Virtual Machine and it will work just like a normal server!

Using a dedicated or pass-through disk

Creating a dedicated or pass-through disk requires a different approach. This is not really a Virtual Disk, but it is a physical disk connected to a virtual disk controller in a Virtual Machine. On the Hyper-V server, open the Server Manager and under storage select Disk Management. Make sure that the disk you want to use as a dedicated disk is Offline. This will ensure that the Parent Partition cannot access the disk. If it can access the disk, and write data on it, it will give unpredictable results in the Virtual Machine!

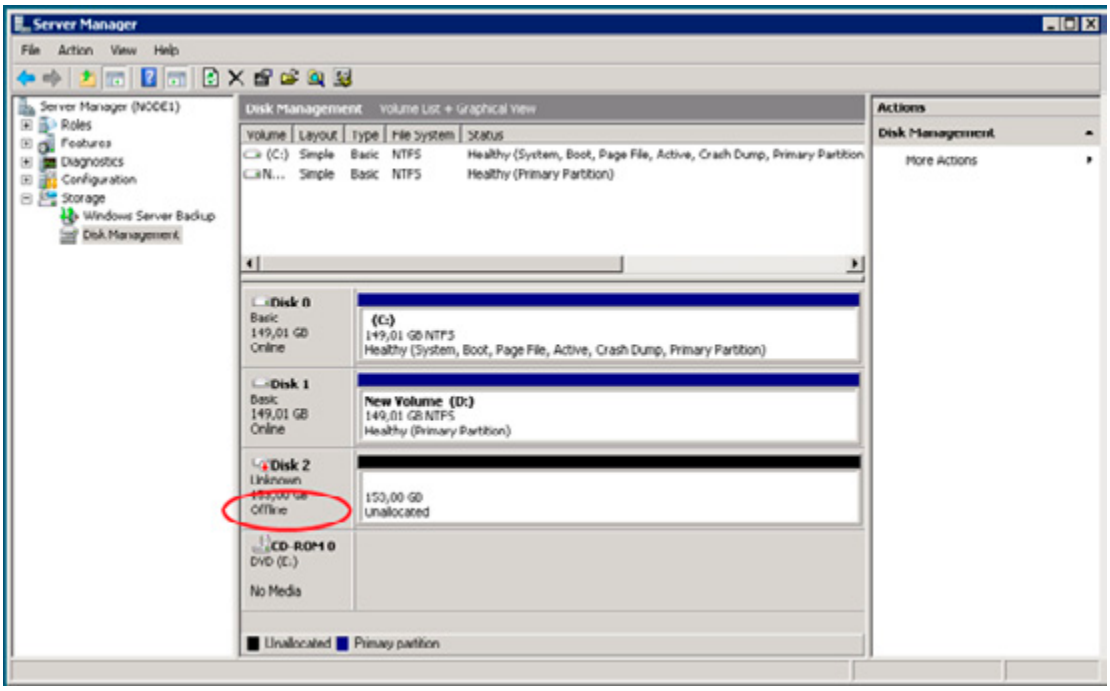

Figure 8. The dedicated disk needs to be offline in the Parent Partition!

In the Hyper-V Manager right-click the Virtual Machine and select Settings. Select the IDE Controller, select Hard Drive and click Add. Under Media select Physical Hard Disk and select the disk you want to assign to this controller.

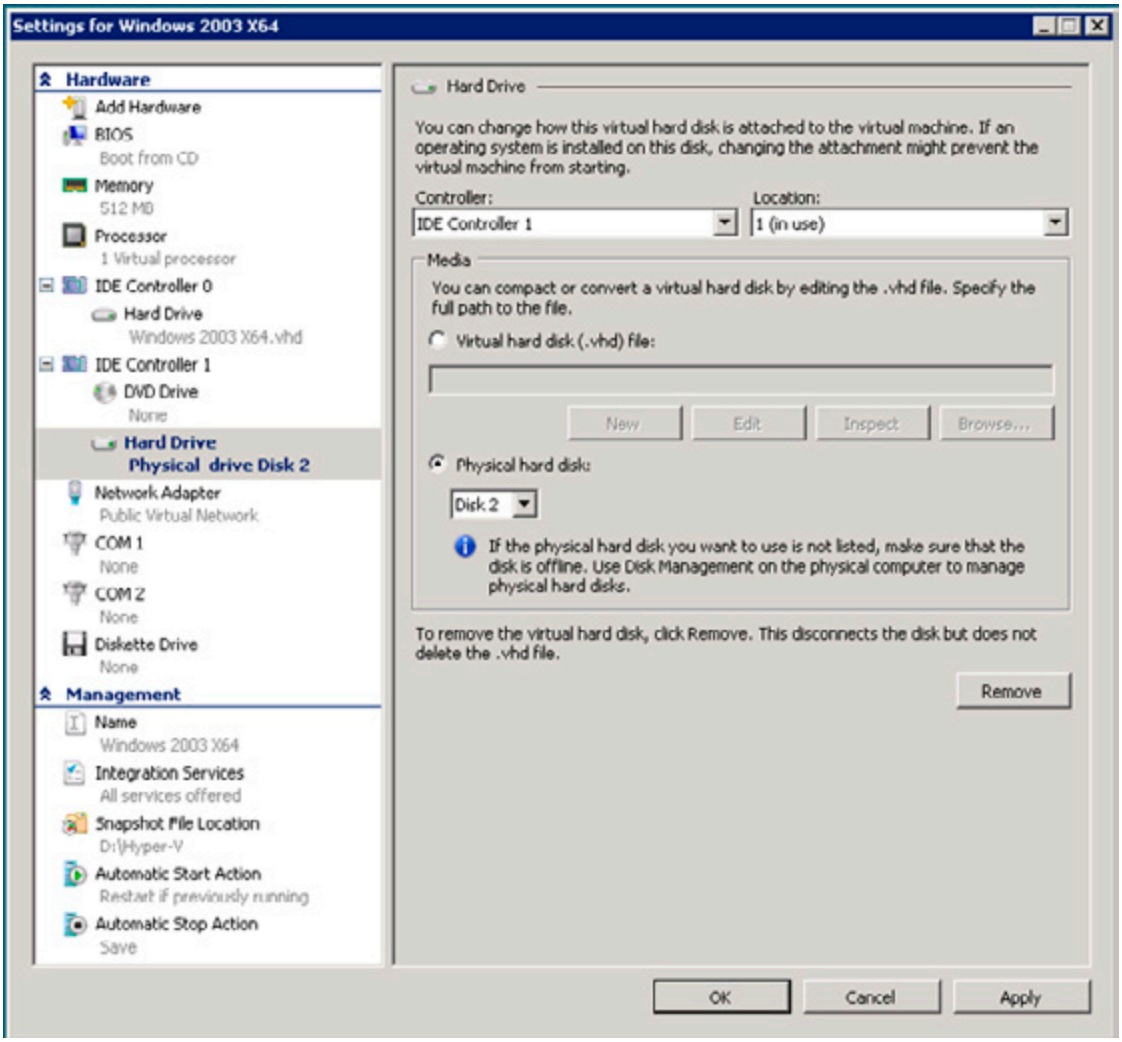

Figure 9. Add the physical disk to the disk controller. Note that this is the same disk as in Figure 8!

When you boot your Virtual Machine and open Disk Management the new disk will show up, ready to be partitioned and formatted.

Is is possible to boot from a pass-through disk? Yes, but only if assigned to the IDE controller. SCSI Controllers can be used for data disks. When creating a Virtual Machine that needs to boot from a pass-through disk you have to create the VM Configuration first, then add the pass-through disk to the IDE controller and then start up the VM.

Conclusion

In this second article I tried to give more information regarding the installation of Hyper-V, Hyper-V's Virtual Networks, Virtual Hard disks and Virtual Machines running under Hyper-V. In my next article I'll discuss the deployment of Virtual Machines and the use of System Center Virtual Machine Manager (VMM) 2008 in a Hyper-V environment.

Restricting Outlook Client Versions in Exchange 2007

29 July 2009

by Ben Lye

There are good reasons for preventing old versions of Outlook from connecting to Exchange Server. You'll probably, at least, want to do it for security. Before you do so, you'll also need know what versions are out there being used so you can make sure that blocking of legitimate users is prevented. Ben Lye explains how it is done.

Outlook has been around for a long time, and there are many versions with different features and varying levels of security fixes. It's not always desirable to let any version of Outlook connect to Exchange Server as you may require your clients to utilize a specific feature set, such as Messaging Records Management, or to have certain security fixes. You may also have corporate policy which dictates that a particular version or versions of Outlook are used.

Fortunately Exchange has the ability to restrict which versions of Outlook can connect by blocking MAPI client versions, a feature which was introduced in Exchange 2000 Service Pack 1. In Exchange 2007 MAPI client blocking can be implemented on a per-server basis using a registry change, or on a per-mailbox basis using the Exchange Command Shell.

Additionally, Microsoft recommends implementing Outlook client blocking as a best practice, and if you run the Exchange Best Practices Analyzer against a server which does not have client blocking enabled it will suggest that you configure it.

Determining Which Client Versions are in use

Before implementing client version blocking it's a good idea to know versions are in use. With this information you can tell which clients need to be upgraded to a newer version before blocking is implemented, or simply which clients will no longer be able to connect after it is implemented. In Exchange 2007 client version information is retrieved using the **Get-LogonStatistics** cmdlet.

Get-LogonStatistics accepts a mailbox, a mailbox database, or a server name as input and returns statistics including user name, logon time, last access time, client name, and client version.

For example, to list the client versions used to access a single mailbox, the command is:

```
Get-LogonStatistics JSmith | ft UserName,ClientVersion,LogonTime
```

To list the client versions for all clients connecting to a specific server:

```
Get-LogonStatistics –Server SERVER01 | ft UserName,ClientVersion,LogonTime
```

To list the client versions for all clients on all mailbox servers, and export the results to a CSV file:

```
Get-MailboxServer | Get-LogonStatistics | `
Select UserName,ClientName,ClientVersion,LogonTime | `
Export-Csv -Path ExchangeClientVersions.csv
```

Once you have identified the clients in use in your organisation and taken any remedial action necessary you can move on to blocking any further access by unwanted clients.

Determining Which Client Versions to Block

The client version is determined by the version of Emsmdb32.dll on the client. This is not necessarily the same as the Outlook version, or the version of any other DLL or executable files. This table shows the version of Emsmdb32.dll for major releases or updates of Outlook since the release of Office XP.

Release	Emsmdb32.dll Version
Office XP RTM	10.0.2627.1
Office XP SP1	10.0.3416.0
Office XP SP2	10.0.4115.0
Office XP SP3	10.0.6515.0
Office 2003 RTM	11.0.5604.0
Office 2003 SP1	11.0.6352.0
Office 2003 SP2	11.0.6555.0
Office 2003 SP3	11.0.8161.0
Office 2007 RTM	12.0.4518.1014
Office 2007 SP1	12.0.6211.1000
Office 2007 SP2	12.0.6423.1000

Table 1: Emsmdb32.dll version by Office release.

There are some important points to note:

- The MAPI client version numbers listed in Table 1, and those in the results of the Get-LogonStatistics cmdlet, are in the format x.0.y.z. When specifying MAPI versions to be blocked you must use the format x.y.z. For example, the version number for Outlook 2003 RTM becomes 11.5604.0.

- When setting per-server restrictions it is very important to avoid restricting clients with version numbers 6.y.z as Exchange Server makes use of MAPI for server-side component connections and uses MAPI versions within the 6.y.z range (with the version number potentially varying by Exchange component and patch level). This does not apply to per-mailbox restrictions.

- Microsoft recommends that at a minimum you block all MAPI clients with version numbers equal to or earlier than 5.3164.0.0.

Single versions are blocked by specifying the version in the format <*version*>, an open-ended range is blocked by using the format -<*version1*> or <*version2*>-, and a specific inclusive range is blocked by using the format <*version3*>-<*version4*>. Multiple sets of client versions can be disabled using a comma or semi-colon separated list.

Range Type	Example	Effect
<version>	11.5604.0	Block the specified MAPI version.
-<version>	-11.0.0	Block the specified version number, and all previous versions.
<version>-	11.0.0-	Block the specified version number, and all newer versions.
<version>-<version>	11.0.0-11.9.9	Block the specified version numbers, and all clients between the specified versions.

Table 2: Example client version blocking syntax.

Some example blocking settings to use:

Blocking Setting	Effect
11.5604.0	Block Outlook 2003 RTM
-5.9.9;7.0.0-11.9.9	Block all clients older than Outlook 2007
12.0.0-	Block all versions of Outlook starting with Outlook 2007 and including all future versions
-5.3164.0	Block Microsoft recommended Outlook versions
-5.99;7.0.0-	Block all MAPI clients, except for Exchange Server components

Table 3: Example client version blocking settings.

Implementing the Blocking Settings

As mentioned earlier, restrictions can be implemented per-server or per-mailbox. Per-server restrictions are implemented via a registry change, and per-mailbox restrictions are implemented via the Exchange Management Shell – it is not possible to use the Exchange Management Console.

If both server and mailbox restrictions are used the most restrictive combination of both settings applies, additionally a server restriction cannot be overridden by a mailbox setting.

For example:

Server Restriction	Mailbox Restriction	Net Effect
-5.3164.0	-11.9.9	Mailbox can only be accessed using Outlook 2007
-5.9.9;7.0.0-	11.0.0-11.9.9	Mailbox cannot be accessed by any MAPI client (which is not executing on the Exchange server)

Table 4: Cumulative effect of restrictions.

To implement a per-server restriction

Note

Incorrectly editing the registry can cause serious problems that may require you to reinstall your operating system. Problems resulting from editing the registry incorrectly may not be able to be resolved. Before editing the registry, back up any valuable data.

You need to be a local administrator on the Exchange server in order to edit the registry.

- Start the registry editor on your Exchange 2007 Mailbox server.
- Locate the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\MSExchangeIS\ParametersSystem registry key.
- Right-click ParametersSystem, select New, and then select String value.

- Name the new string value "Disable MAPI Clients."
- Right-click Disable MAPI Clients, and then click Modify.
- Enter the restriction setting, for example -5.9.9;7.0.0-11.9.9.
- Close the registry editor.

The change will be effective within 15 minutes, or to make it effective immediately you can restart the Microsoft Exchange Information Store service. Once the change takes effect any existing client connections which do not meet the version requirements will be terminated.

To implement a per-mailbox restriction

The **Set-CASMailbox** cmdlet is used to implement per-mailbox restrictions. To use the `Set-CASMailbox` cmdlet you must be delegated the Exchange Recipient Administrator role.

To prevent a mailbox from using Outlook clients prior to Outlook 2007 the command is:

```
Set-CASMailbox JSmith -MAPIBlockOutlookVersions "-11.9.9"
```

To remove a restriction for a mailbox:

```
Set-CASMailbox JSmith -MAPIBlockOutlookVersions $null
```

To prevent all mailboxes in a particular database from using clients other than Outlook 2007 RTM:

```
Get-Mailbox -Database "SERVER01\SG1\Database 1" | `
Set-CASMailbox -MAPIBlockOutlookVersions "-12.4518.1013;12.4518.1015-"
```

When an Outlook 2003 or Outlook 2007 user tries to connect with a restricted client version they will receive the message "Your Exchange Server administrator has blocked the version of Outlook that you are using. Contact your administrator for assistance."

Users of older clients will receive the message "Cannot start Microsoft Outlook. The attempt to log on to the Microsoft Exchange Server computer has failed."

More information on Get-LogonStatistics and Set-CASMailbox can be found in TechNet:

Get-LogonStatistics – http://technet.microsoft.com/en-us/library/bb124415.aspx.

Set-CASMailbox – http://technet.microsoft.com/en-us/library/bb124415.aspx.

Using Twitter and PowerShell to Find Technical Information and Join a Community

30 July 2009

by Jonathan Medd

Using PowerShell and a little bit of .NET Framework and XML knowledge, it is possible to delve a little deeper into the information which is potentially available to you from Twitter. Jonathan explains about Twitter and shows how to use Powershell to access twitter automatically.

For somebody so interested in new technology, particularly within the Enterprise workplace, I remain healthily sceptical of the latest new fads and how useful they might actually be. A colleague of mine had for some time been pestering me to use social networking sites saying I was missing out on some great tools for information gathering and also being part of various online communities. As far as my initial look took me I couldn't really see much benefit past a bit of fun time wasting at idle parts of the day and potentially having to deal with some privacy issues.

After a while though, seeing how enthused about it he was and how much he seemed to be getting out of it I gave in and signed up. There are plenty of social networking sites out there these days, each of which has their possible different uses in different situations – in this article I'm going to look at the micro-blogging service Twitter and how you can use it to get technical information online and also participate in a community.

As with many ideas the best ones can typically be the simplest ones. Using messages of up to 140 characters, tweets, can be published to your Twitter page and consequently anybody can be subscribed to your updates with a variety of different applications. Based on SMS messaging technology they typically describe what you are doing right at that moment in time, but instead of being read by just one person you have sent a text to, anybody following you on Twitter is able to read them. These can range from the informative "I just fixed this technical problem by doing x....." to the banal "Just missed the bus, I'm going to be late for work." It's these trite comments which are often highlighted as Twitter being a great time waste, but if you ignore these there is some great technical content just waiting for you to grab hold of it.

Getting Started

First of all you will need to sign up for an account at Twitter. Once complete you can use the web interface as a basic starting point.

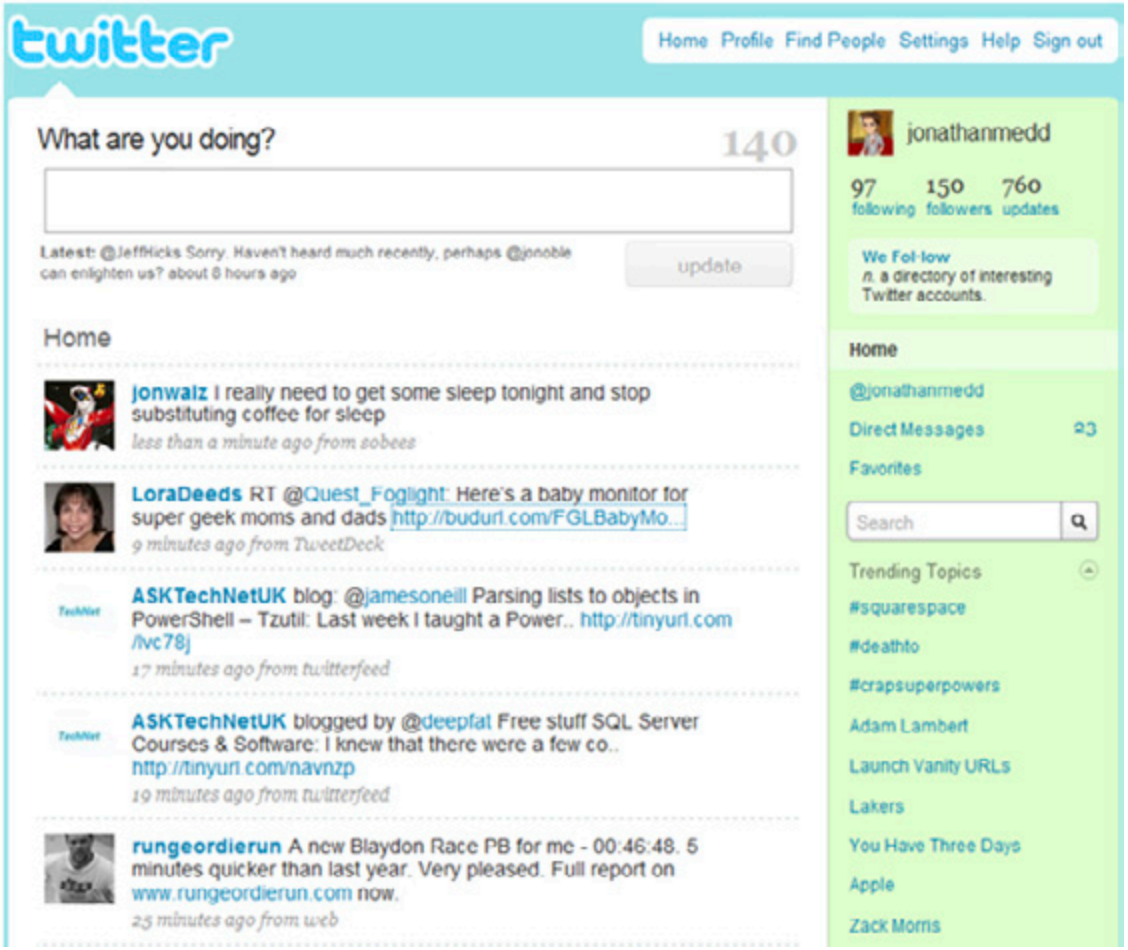

This is a typical view of a Twitter home page. From here you can post updates, add Twitter users you wish to follow, see recent updates from those you are following, view a list of those following you, carry out some basic searching and view trending topics (some of the current most popular words being talked about on Twitter, which usually gives a good barometer for some of the most commonly talked about news stories of that point in time).

Tips for getting started

Twitter is not very useful if you are not following anybody! Use the search bar above or http://search.twitter.com/ to perform a search of Twitter on a topic. You will see a list of updates including the topic; anything that looks interesting simply start following the person who published the tweet.

- Be careful though, don't follow too many people all at once initially otherwise you'll be subjected to a tidal wave of tweets which you can't keep up with and you'll be wondering why you started using Twitter.
- You don't have to follow everybody who follows you. Although it may seem the correct etiquette to do so, again you may end up with updates you are not really interested in – check out the person's recent updates first to see if they appeal to you.
- Don't worry if you don't look at Twitter for a few hours or even days – it's not like email where you need to go back and catch up, rather something you can dip in and out of as and when you have the time or inclination.

- Be careful what you post. Once it's there, although you can delete it, it doesn't take long for a message written in haste to be sent round the globe can cached on a server somewhere.

Also it's probably not a good idea to tell the world that you are going on holiday and leaving an empty house for two weeks. (See:Note to Self: Don't Tweet Vacation Plans) If you want to be extra cautious you can approve each person who wishes to follow your Twitter updates before you let them see the updates.

If you are somebody who already likes to engage with a technical community and share content you can use Twitter to help promote a new post on your blog, a new episode of a podcast or an update for an open-source project you work on.

For instance as well as manually posting a link to a new blog post you have just written, why not use a service like http://twitterfeed.com to automatically publish a tweet for the new post. Or over at the Get-Scripting Podcast we have created a Twitter account for the podcast and use it to keep listeners up-to-date with show releases and what's coming up on future shows.

It never ceases to amaze me of new uses for Twitter that people find. At a technical event I attended recently the presenter encouraged those in the audience to send him questions via Twitter during his presentation. I have also attended another event in the past which was a Q&A session; during the session I put out a tweet saying where I was and ended up asking questions for people elsewhere sending me them via Twitter. Another example comes from a situation where I had made a forum post which had gone unanswered for a few days; I was particularly keen to get an answer so I tweeted about it and within a couple of hours I had some answers which had previously started to look like they were not going to appear.

A number of corporate and marketing types have begun to cotton on to this real-time and potential personal interaction with customers and will provide Twitter feeds about their products and company. Whilst IT pros may not always be particularly interested in marketing spiel it can be useful to say find out when new updates for a product are released or maybe follow the product manager who may ask what kind of features / updates people would like to see in a product. It's great that these opportunities now exist in these more open times and even massive companies like Microsoft who may in the past have hidden away the technical people behind the marketeers are now much more open and easier to interact with.

Twitter Clients

Fairly soon you will find that the basic web page for Twitter doesn't have enough features for your now hopefully growing Twitter use. Since Twitter publishes an API for accessing their service there are a plethora of clients out there you can use instead. A couple of the most commonly used clients on desktop and laptop machines are Twhirl, and TweetDeck, both of which are free to use. Using a client gives you a feature rich experience for using Twitter.

Twhirl

Twhirl gives you features like notifications for new messages, spell checking, coloured replies (green) and direct messages (yellow), posting images to twitpic and automatically finding tweets mentioning your username – in the web client you won't see replies from people unless you are following them.

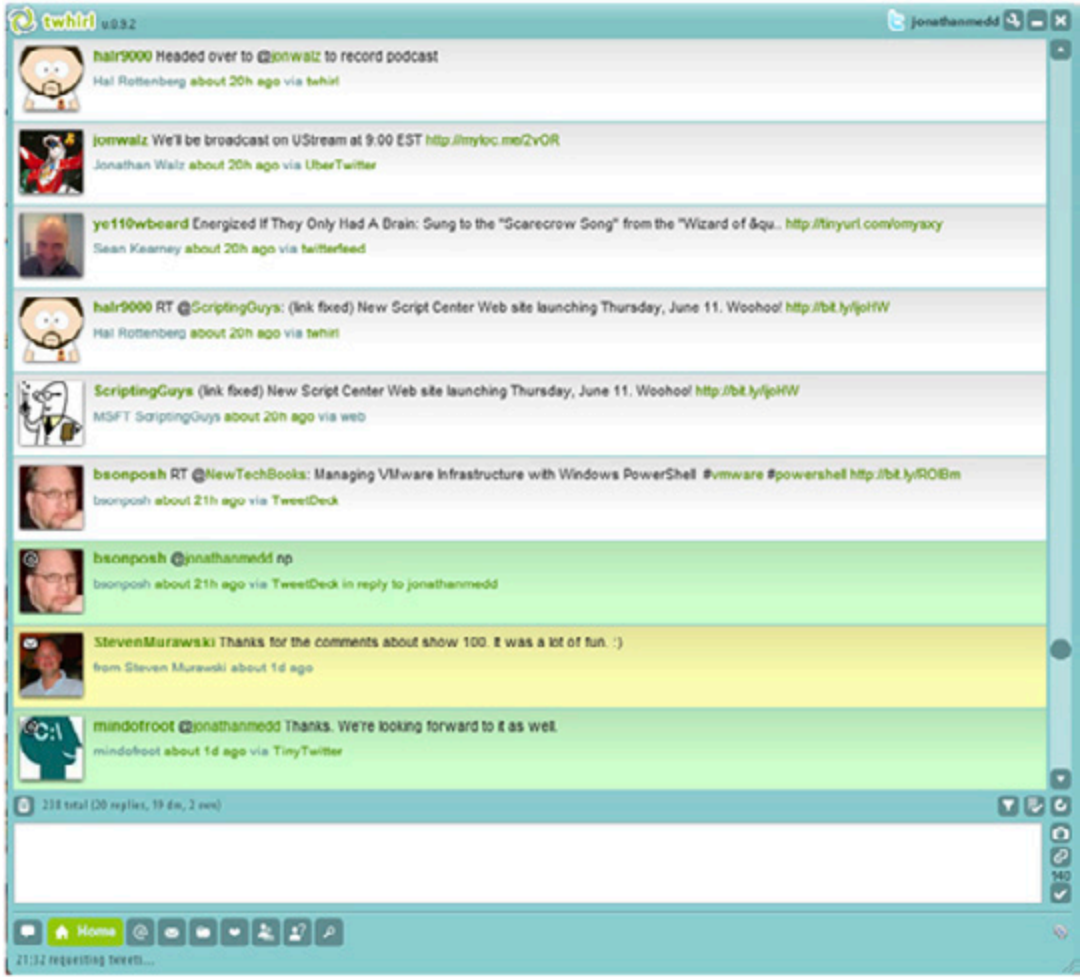

TweetDeck

TweetDeck offers fairly similar features to Twhirl, with one particularly great extra feature – the ability to maintain extra columns based on searches. So say for instance you are interested in what people are talking about Exchange 2010, you can create a search based on the phrase "Exchange 2010" and view the results in a column next to your main Twitter feed. The difference being that the results are dynamic, so the client will poll for new results allowing you to keep track of conversations.

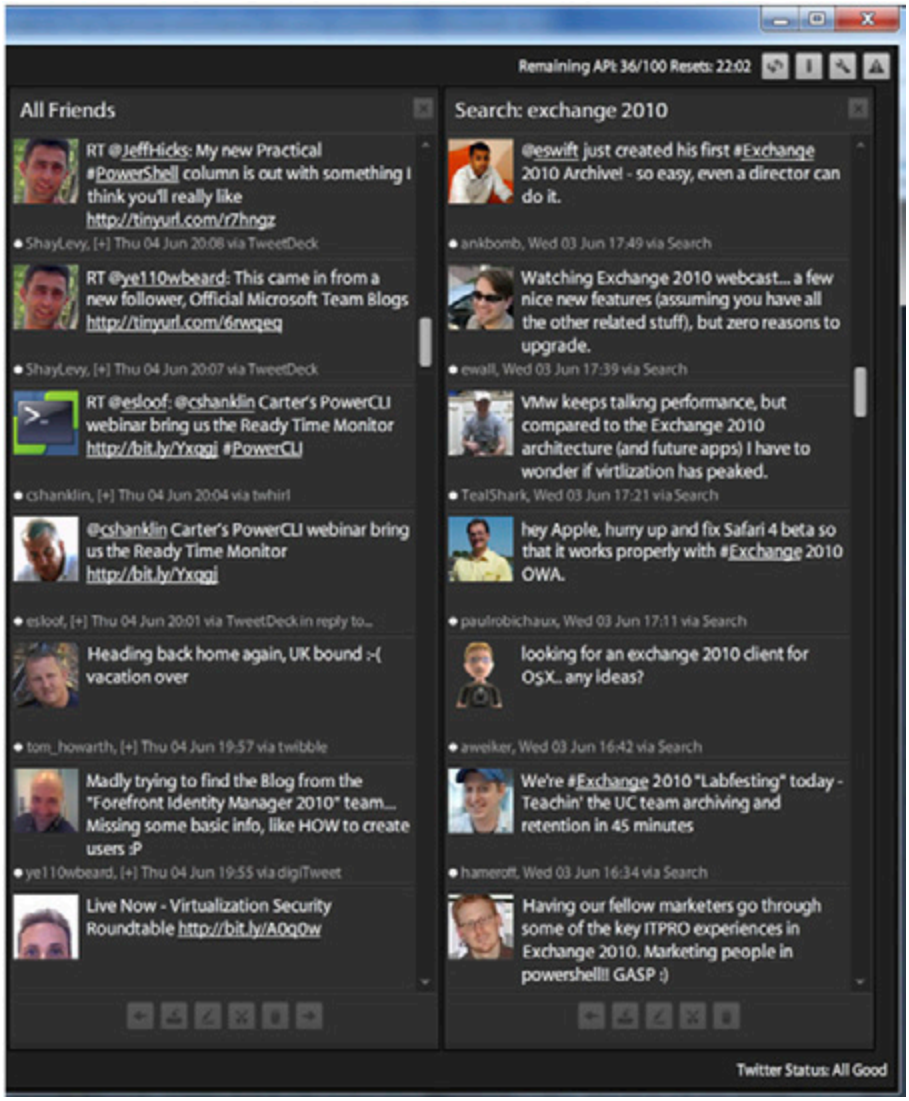

Once you start to see information coming in like this, not only will you be really up-to-date with all the latest news of the topics you are interested in, you'll find yourself wanting to engage with these communities, join in by posting your own information and starting conversations with people with similar interests.

Mobile devices

Most IT people these days will likely be carrying around some kind of Internet connected mobile device for everything from keeping up with emails to remote administration, naturally there are Twitter clients for these devices too. I won't cover too much on these since you are far more likely to use Twitter on a mobile device for things like up-to-date travel information, but you may wish to view some topics whilst out and about.

Windows Mobile

Twikini seems to be one of the most recommended clients for Windows Mobile based devices. It has most of the features you would expect from a Twitter client.

iPhone

TwitterFon gets a lot of good reviews and appears generally popular.

As with most of these kinds of things it's probably a decision for personal preference and worth trying out a few options to see what works best for you.

PowerShell and Twitter

Since the Twitter API enables you to access their offerings via a web service if you had the inclination you could of course write your own client. However, if you don't want to go quite that far you can use PowerShell and some very straightforward .NET code to get some very interesting information out of Twitter. James O'Neill has published on his blog a set of PowerShell functions which he has put together so that you can interact with Twitter direct from the command line.

Tip

The downloadable functions are stored in a .ps1 file, i.e. a PowerShell script file. To make them accessible in PowerShell you can either include them in your profile so they are available every time you run PowerShell, or for one off use add them to your current session like so:

```
C:\Temp>. ./twitter.ps1
```

(Note the dot, then a space, then a second dot)

Most of the information you wish to get from Twitter is available via an XML file; PowerShell supports manipulating XML files really well and is particularly straightforward for a beginner. The code is stored in the Function Get-TwitterSearch for ease of use to run the searches.

> ***Note***
>
> *The following code is slightly amended from James's original function since at the time of writing Twitter appear to be making their search available via Atom rather than RSS; the RSS search gives intermittent errors)*

```
Function Get-TwitterSearch {
 Param($searchTerm, [switch]$Deep)
 if ($WebClient -eq $null) {$Global:WebClient=new-object System.Net.WebClient }
 $results=[xml]($webClient.DownloadString("http://search.twitter.com/search.atom?rpp=100&page=1
&q=$SearchTerm"))
$Searchitems=$results.feed.entry
 if ($Deep) { $refresh = $results.feed.link | Where-Object {$_.rel -eq 'refresh'}
   $MaxID= $refresh.href.split("=")[-1]
   2..16 | foreach { $Searchitems += ([xml]($webClient.DownloadString("http://search.twitter.
com/search.atom?rpp=100&max_id=$maxID;&page=$_&q=$SearchTerm"))).feed.entry} }
 $SearchItems
}
```

At that point we could stop the search there, but that would only have given us 100 results. It may well be better to get more information, in which case the search function has a "deep" switch which essentially runs the search another 15 times further back each time in the Twitter timeline. A good example for why you might wish to do this is searching for people who Twitter about topics you are interested in and might be interesting to follow.

If we carry out a "deep" search and store the results into a variable we can then use standard PowerShell techniques to manipulate the data into something potentially interesting. For instance earlier we used TweetDeck to search on "Exchange 2010" and it would show recent and new posts about that topic. Let's run the same search using the Get-Twitter search function.

```
$twittersearch = Get-TwitterSearch "Exchange 2010" –deep
```

Now let's group the results by the author of the tweets, sort them by the authors with the most tweets, pick out the top 20 and display them in a table with their count and the author's name.

```
$twittersearch |
  Group-Object {$_.author.name} |
  Sort-Object count -Descending |
  Select-Object -first 20 |
  Format-Table count,name –auto
```

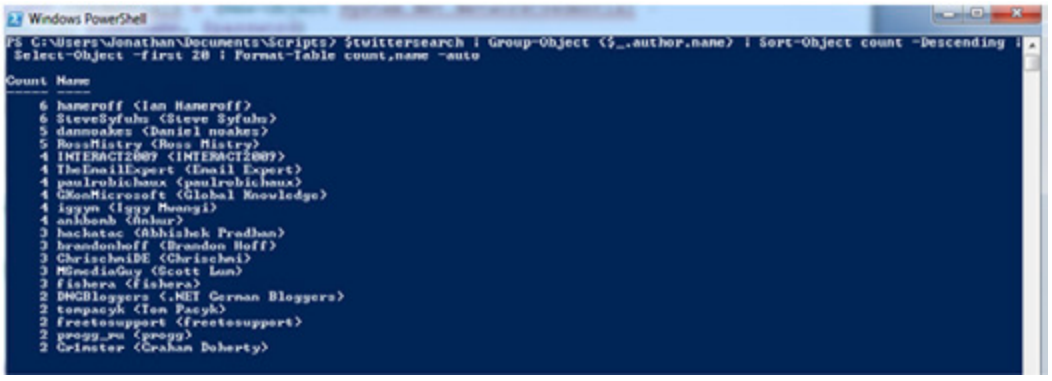

You can then check out the authors in question by looking at http://twitter.com/*authorname* for their recent updates and see if they might be interesting to follow. Exchange 2010 is still a relatively new topic so the numbers above aren't massive. The below screen shows a search on PowerShell and gives a pretty good indication of the people who Twitter about it a lot.

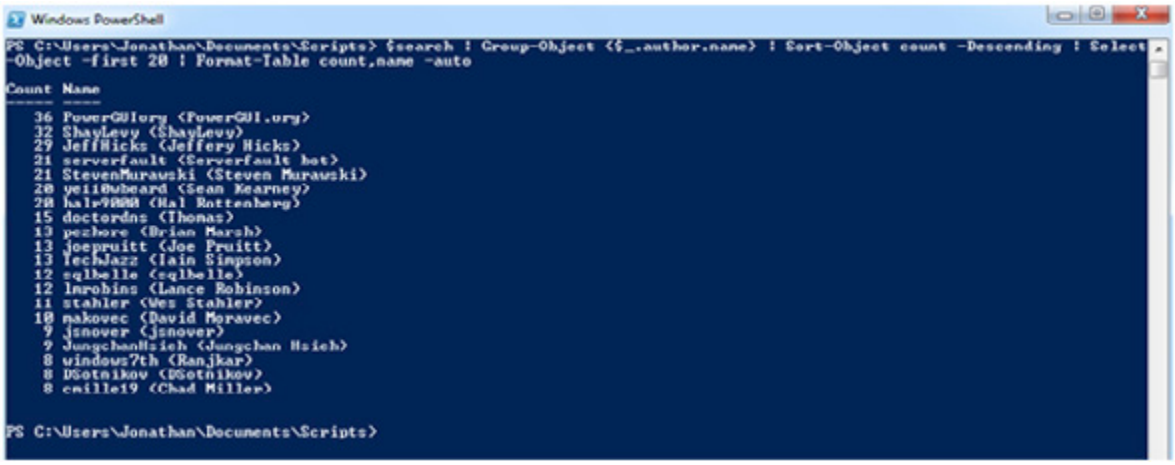

You can obviously now run these searches based on the topics of your choice and find people worth following.

> ***Note***
>
> *Be aware that you may get HTTP errors like "502 Bad Gateway" or "503 Service Unavailable" when searching, particularly if you run multiple "deep" searches or are using it at peak Twitter usage times. This is because the Twitter service uses rate limits to throttle back accounts running too many searches; they also currently sometimes have capacity issues at peak times which can mean running these kinds of processes are not quite as reliable as you would hope.*

Something else you can check out is your list of Twitter friends. On the webpage it's quite difficult to see a complete list of people you are following; below James has put the ability to pull this data down into a PowerShell function. (Note that I have amended the function slightly from the original so that if you are following more than 100 Twitter users you will see them all; the original function only pulled down the first page of 100)

Essentially we again create a new WebClient object, give it the credentials to use for Twitter, pull down the data in XML format and return the data about the users.

```
Function Get-TwitterFriend {
param ($username, $password, $ID)
if ($WebClient -eq $null) {$Global:WebClient=new-object System.Net.WebClient }
$WebClient.Credentials = (New-Object System.Net.NetworkCredential -argumentList $username,
$password)
$nbrofpeople = 0
$Friends = @()
if ($ID) {$URL="http://twitter.com/statuses/friends/$ID.xml?page="}
else {$URL="http://twitter.com/statuses/friends.xml?page="}
do {$Friends += (([xml]($WebClient.DownloadString($url+($nbrofpeople/100 +1)))).users.user)
$nbrofpeople += 100
} while ($Friends.count -eq $nbrofpeople)
$Friends
}
```

So to run the function and see your friends listed in alphabetical order with the Name, Location and Description properties, you would type:

```
Get-TwitterFriend username password |
  Sort-Object Name |
  Format-Table Name,Location,Description
```

This would give you results like those shown below.

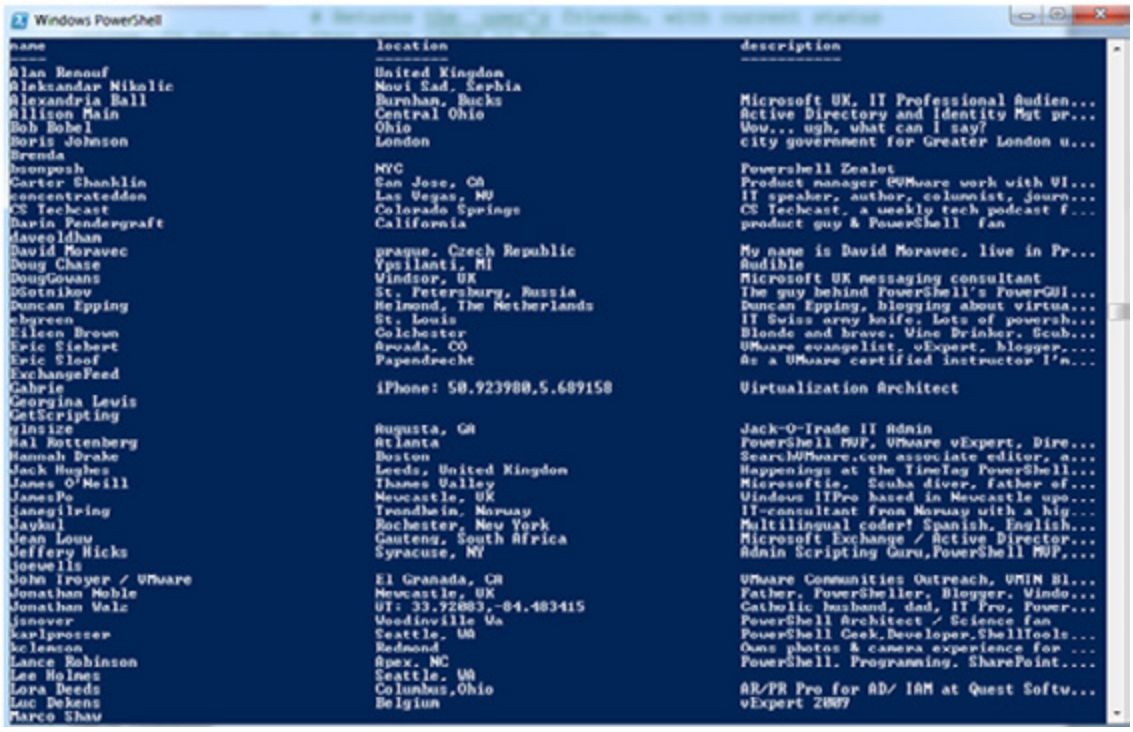

There are nine other functions in James's Twitter function library for you to play with including things like Get-TwitterReply to see replies you have sent, Publish-Tweet to send a message to Twitter or Add-TwitterFriend to add people to your friends list. For instance say you have run one of the previously mentioned searches and found the top 20 people talking about the topic you are interested in; you could then pipe the results onwards and using the Add-TwitterFriend function add all of them to your friends list in one easy go.

Twitter PowerPack for PowerGUI

If you're not 100% comfortable with working totally from the shell yet, but are interested in some of the functionality that it can provide, then I have plugged James's functions, plus some other extras, into a PowerPack for Quest Software's free product PowerGUI. This provides a graphical front-end to PowerShell scripts and produces results into a grid format. Simply download the PowerGUI tool from the http://powergui.org website and the Twitter PowerPack from the library and you will see the below choices.

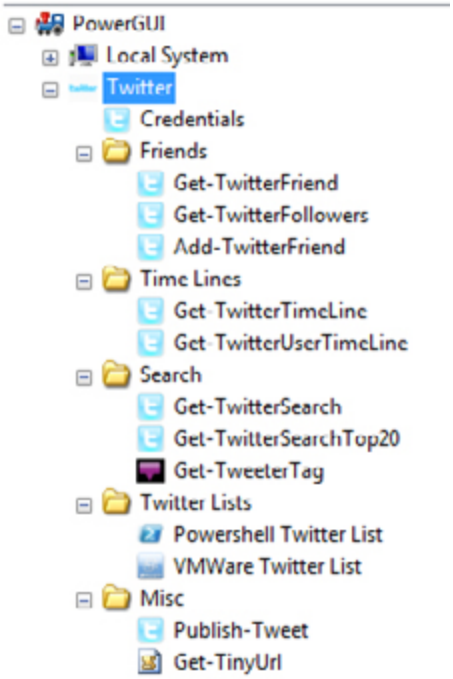

First of all use the **Credentials** node to store your Twitter username and password into a global variable so that you don't need to supply them each time you run one of the scripts. You can then run any of the script nodes and see the results in the grid pane. In the below example I'm using the **Get-TwitterFollowers** script to see the list of people who follow me and I've clicked on the name column to sort them alphabetically.

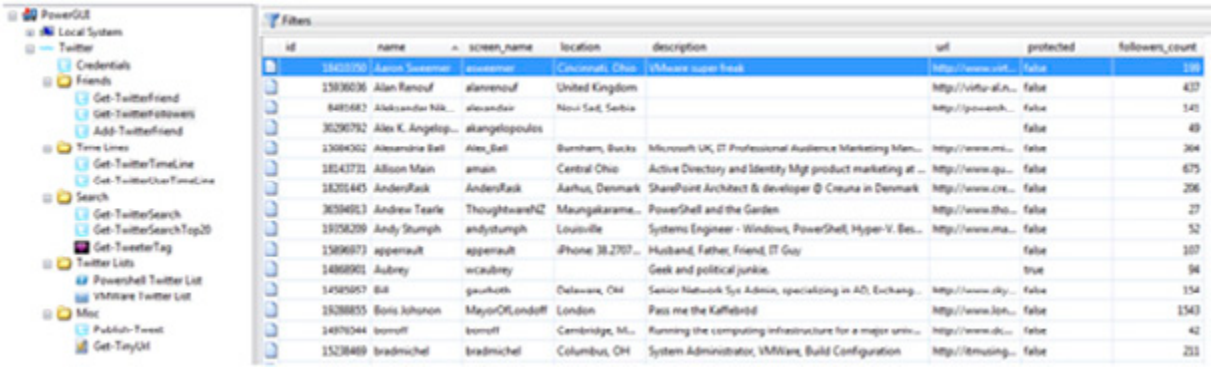

Another way to find people on Twitter to follow is to use one of the maintained lists of people who talk about particular subjects. A couple of these lists feature in the Twitter PowerPack, below is the VMWare list maintained over at the http://Virtu.al.net website.

Clicking the script node will produce a list of the VMWare Twitterers; you could select multiple or all people and then use the **Add-TwitterFriendFromList** action based on the previously mentioned **Add-TwitterFriend** function to add these all to your follow list in one go.

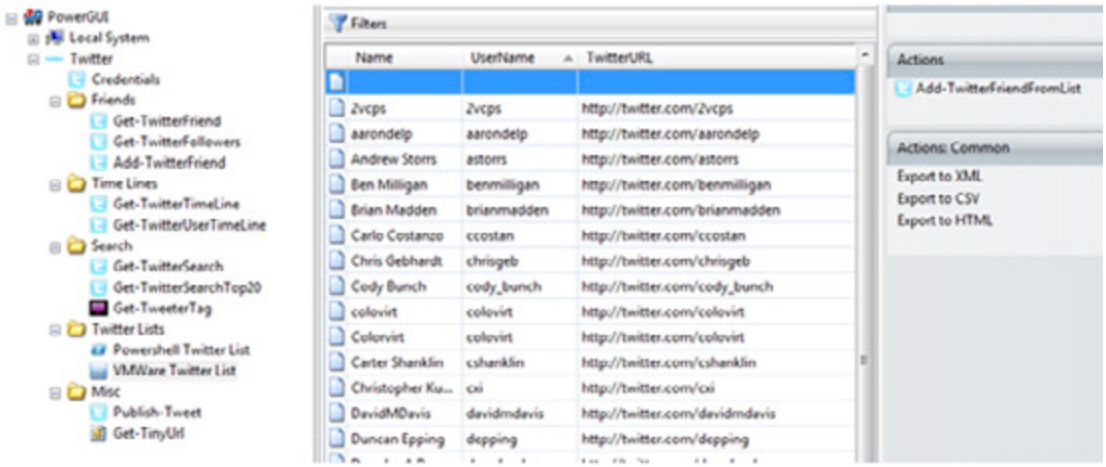

A potential drawback of these lists is that they are statically maintained. To use a more dynamic approach you could use a website like http://www.tweetertags.com where Twitter users tag themselves with subjects they are interested in and you can search these tags for people to follow.

Not surprisingly you can search these tags using PowerShell and I've included that in the Twitter PowerPack. For instance using the **Get-TweeterTag** node and searching for "Exchange" tags on the website returns the following list of users; you could then use the **Add-TweeterFriend** action to add a multiple selection of people to your follow list.

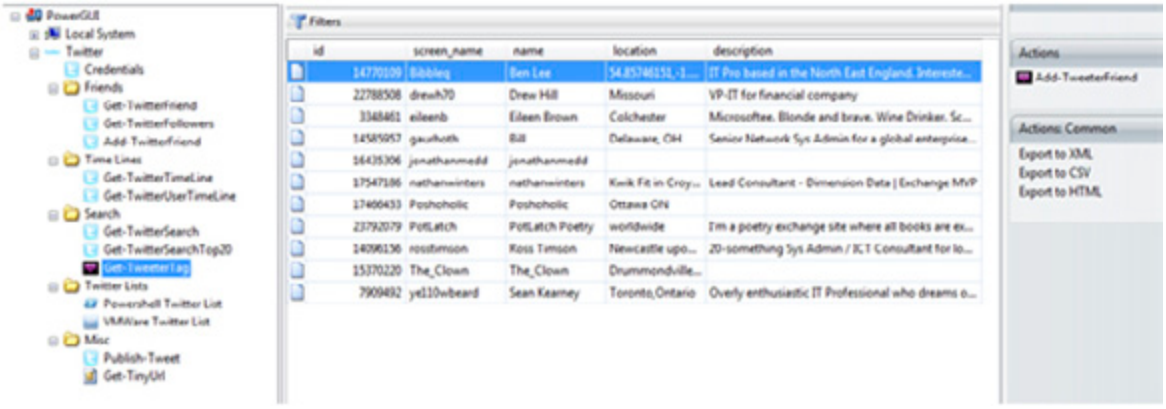

Conclusion

Twitter is one of those things where you can get as much out of it as you want to; it can be anything from a fun time-wasting tool to something really useful to find information and people around the globe with similar interests to yourself. The various Twitter clients for differing devices can help you keep in touch with these communities wherever you are.

Using PowerShell and a little bit of .NET Framework and XML knowledge you can delve a little deeper into the information which is potentially available to you from this social network.

Update: Exchange Server 2010 Release Candidate

18 August 2009

by Jaap Wesselius

Exchange Server 2010 Release Candidate is now available for download from Technet, so we asked Jaap to review it for us.

In an earlier article (Exchange Server 2010 – The first public beta version) that I wrote in April of 2009, I discussed the first public beta of Exchange Server 2010. Microsoft has now released an update of this beta, Exchange Server 2010 Release Candidate which can be found here: .

The Release Candidate is now feature complete, so no more new features will be added. The only thing Microsoft now has to do until the release of Exchange Server 2010 later this year is bug fixing. So what is new and what has changed in this Release Candidate version of Exchange Server 2010?

Changes and new features in Exchange Server 2010

The following are the most visible additions in the Release Candidate, compared to the first public beta:

- **Archiving** – Microsoft added archiving capabilities to Exchange Server 2010. When a new mailbox is created there's the possibility to add an archive to the mailbox (of course you can add the archive later on as well). The archive is a secondary mailbox in the same database where the primary mailbox resides. The secondary mailbox is not visible to other users and therefore do not show up in the Global Address List. The archive is visible in Outlook 2010 and Outlook Web Access 2010, unfortunately Outlook 2007 does not support this feature. On the client side, no extra configuration is needed, when the Autodiscover functionality is configured correctly the archive will show up in your Outlook profile.

- Users can use the archive to store information manually or Exchange administrators can create Messaging Records on the Exchange server to automatically process messages and store them in the archive. There's a big difference between the mailbox and the archive: the mailbox is available both online and offline, the archive is only available online. This means that the offline copy of the user's mailbox (the .OST file) contains only the items in the mailbox and *not* in the archive.

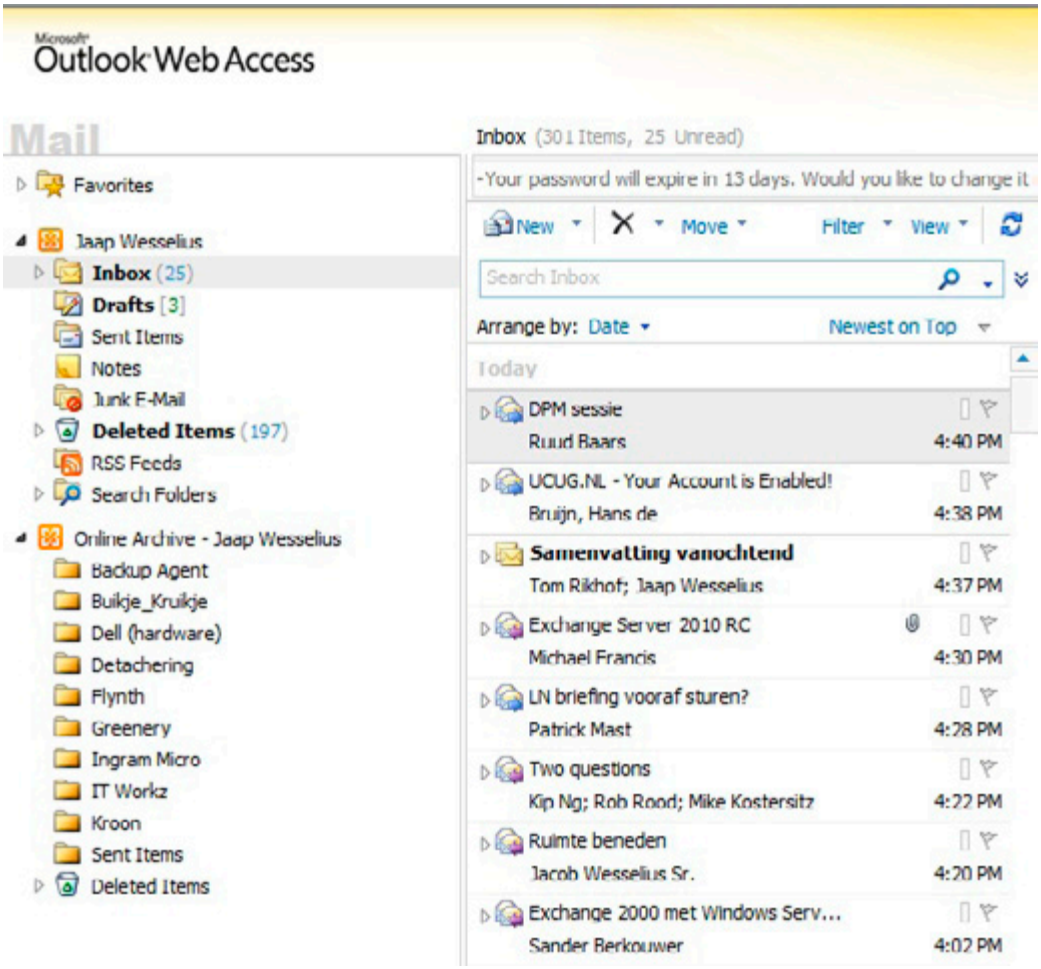

Figure 1. The online archive is visible in Outlook Web Access 2010.

- **Backup** – Exchange Server 2010 Release Candidate contains a plug-in for Windows Server Backup (WSB). Exchange Server 2010 does not support streaming backups anymore, only VSS (snapshot) backups are supported. Unlike NTBackup in Windows Server 2003 R2 and earlier there's nothing visible about WSB supporting Exchange Server. When creating a backup of a volume containing Exchange Server 2010 databases the status indicator reveals that something is done with Exchange Server databases, but that's all. In REF _ Ref238257218 \h Figure 2, drives F:\, G:\ and H:\ contain Exchange Server 2010 databases (and log files) and these disks are selected for backup. Windows Server Backup is responsible for checking the databases for consistency (VSS itself does nothing about consistency, it just makes the shadow copy) and this is the only thing that's visible about the Exchange support in Windows Server Backup. Windows Server Backup itself is pretty limited when it comes to backing up the Exchange Server environment. It can only create full backups (incremental or differential backups are not supported) and only the Active Copy of an Exchange Server 2010 database can be backed up. Creating a backup of a Passive Copy of a database is not supported. The next version of Microsoft System Center Data Protection Manager (DPM) will fully support Exchange Server 2010 and offer a lot more features than Windows Server Backup.

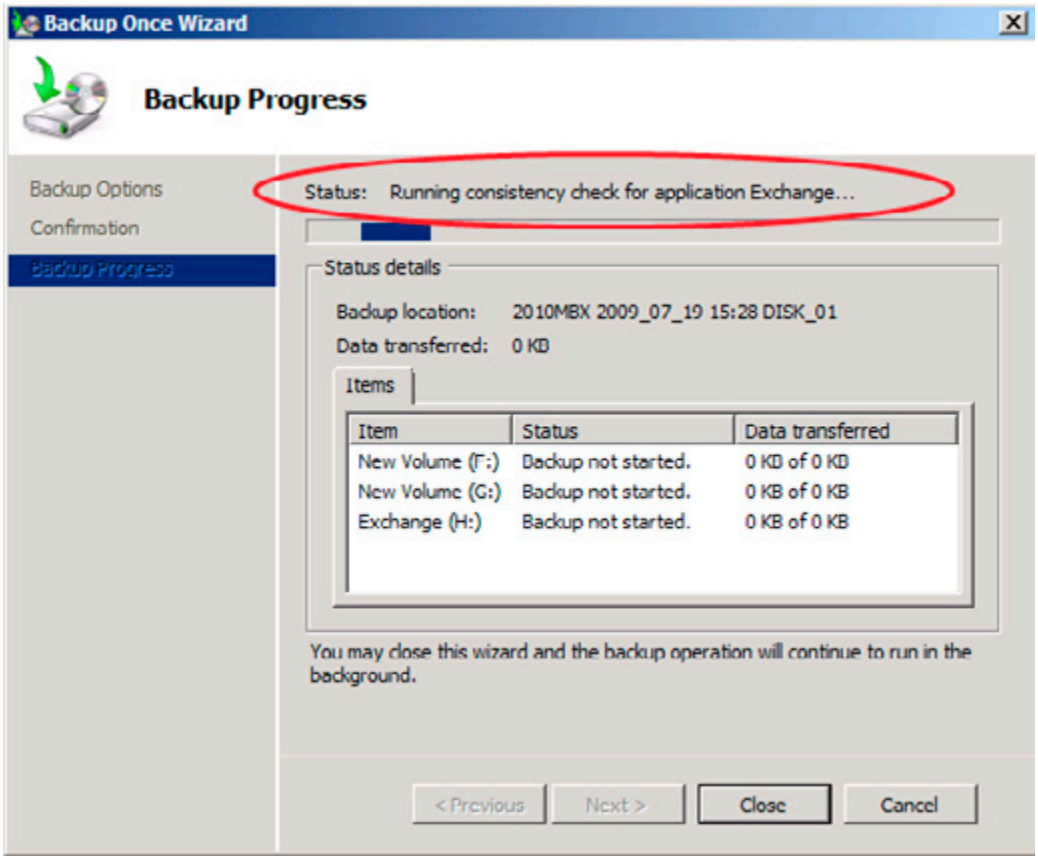

Figure 2. Windows Server Backup checking Exchange databases for consistency.

- **Sizing** – Microsoft is now releasing some information regarding sizing. In my earlier article I already explained that one of the design goals of Exchange Server 2010 was to host a database and the accompanying log files on one (large) SATA disk. The maximum database size that's recommended in Exchange Server 2010 will be 2 TB (10 times the recommended maximum size of a database in an Exchange Server 2007 server in a Cluster Continuous Replication (CCR) configuration!). It is up to the Exchange administrator to determine if a 2 TB database is useful or not, think about the implications when reseeding the database (creating a new copy) is necessary. In this case the complete database will be copied over the network. Or when a utility like ESEUTIL or ISINTEG needs to work on a 2TB database…. But it is still to early to say anything about the performance of Exchange Server 2010 with such a configuration.

- **Powershell** – The beta of Exchange Server 2010 came with two flavors of the Exchange Management Shell: the local version and the remote version. This turned out to be very challenging and numerous issues occurred. The Release Candidate of Exchange Server 2010 now comes with only one Exchange Management Shell. This Exchange Management Shell still supports managing both the local Exchange server as well the a remote Exchange server.

- **Windows Server 2008 R2 support** – The beta of Exchange Server 2010 was only supported on Windows Server 2008. Windows Server 2008 R2, even build 7000, was not supported on Exchange Server 2010 beta. Now Windows Server 2008 R2 is released also Exchange Server 2010 is supported on this platform. All of the prerequisite software needed for installing Exchange Server 2010 (except for the filter pack needed for the Mailbox Server role) is available by default on Windows Server 2008 R2 so this can speed up the installation of Exchange Server 2010. Also the improvements in Windows Server 2008 R2 (like performance) make this a better for running Exchange Server 2010 (but that's my personal opinion ;-)

- **Improvements in the setup** – when setting up an Exchange Server 2010 Client Access Server there's the possibility to enter the external domain name when the Client Access Server is Internet facing. If this option is used all external setting settings on the Client Access Server (for the Offline Address Book, ActiveSync and Exchange Web Services) are automatically configured.

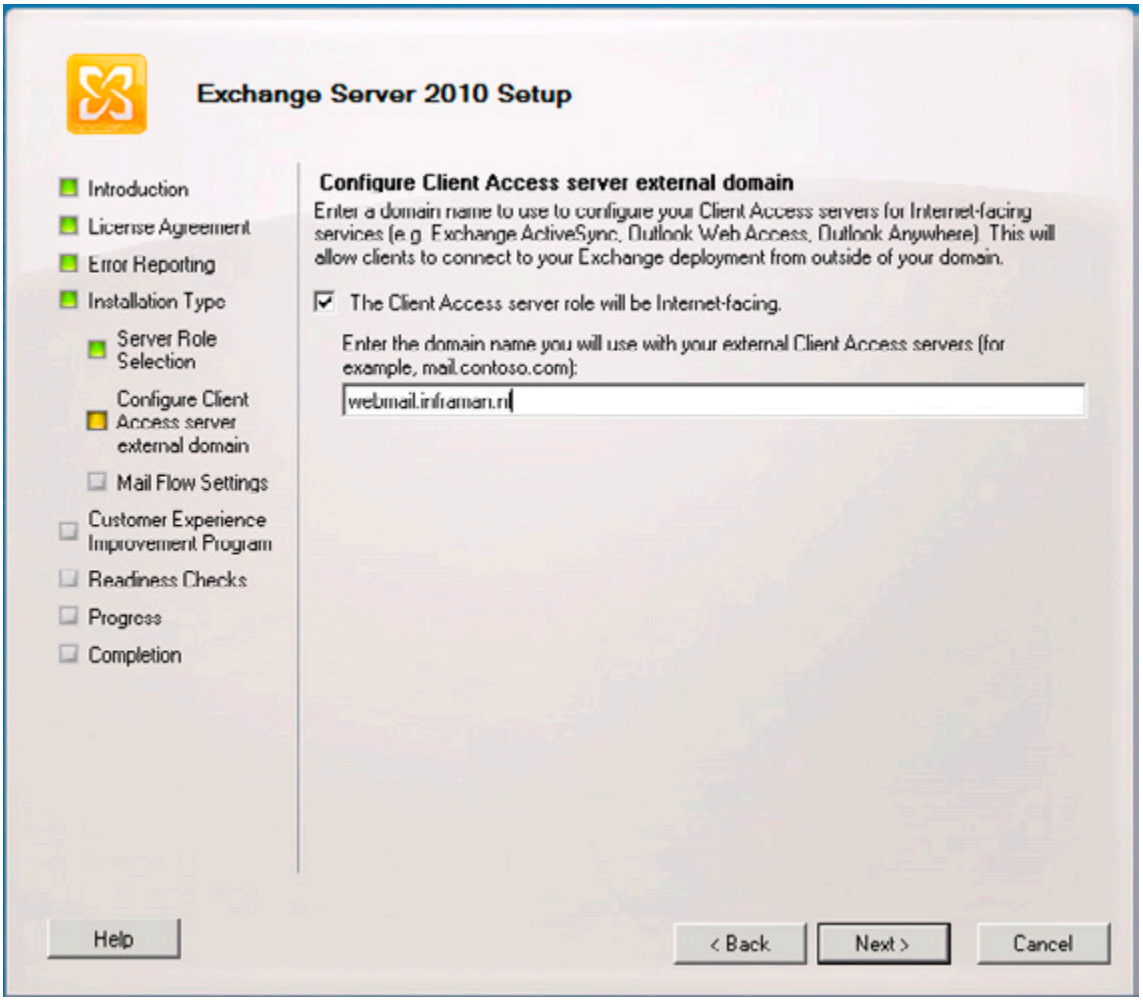

Figure 3. Enter the external domain name when the Client Access Server is internet facing.

If you prefer not to use this option than the mentioned settings can be changed later using the **Set-OABVirtualDirectory**, **Set-WebServicesVirtualDirectory** and the **Set-ActiveSyncVirtualDirectory** cmdlets in the Exchange Management Shell.

And how does it work?

And how does the Release Candidate of Exchange Server 2010 feels in real live? Compared to the public beta this Release Candidate works really great. We've tested an interesting configuration with two combined Hub Transport and Client Access Servers, two Edge Transport Servers and two dedicated Mailbox Servers. These Mailbox Servers are configured in a Database Availability Group where two Mailbox Databases have a copy on the other server. There's also a Public Folder database available for Outlook 2003 clients. Please note that Public Folders cannot be used in combination with the Database Replication feature. When creating multiple copies of Public Folders the Public Folder replication has to be used.

For datacenter resilience all Exchange Server 2010 server roles can be installed in multiple datacenters, even the Mailbox Servers. After installing the Mailbox Server role the Database Availability Group needs to be configured with the subnet of the 2nd datacenter and when done the Mailbox Server can be added to the Database Availability Group. Configure multiple Database Copies using the Exchange Management Console or the Exchange Management Shell and you're done.

Backup using Windows Server Backup is more challenging. As mentioned earlier only the Active Copy of an Exchange Server database can be configured in a backup, Passive Copies are not supported. When something happens in your Database Availability Group and another

server is hosting the Active Copy then you lose the ability to create a backup. Also when you have multiple servers hosting multiple Active Copies you have to configure Windows Server Backup on each Mailbox Server. I'm pretty sure that this will be far more convenient when the next version of Data Protection Manager becomes available. This is not only true for Data Protection Manager but for all 3rd-party backup application vendors like Symantec, UltraBac, CA, etc.

Conclusion

The latest update of Exchange Server 2010 is the Release Candidate version. This version is feature complete thus no new features will be added. Compared to the first Public Beta the most important additions are the archive and the backup feature. Also the stability and performance have been improved so far, it looks really good.

There are still some challenges, documentation that is not yet up-to-date but personally I think Microsoft is on-track with Exchange Server 2010. So, you better start looking and evaluating Exchange Server 2010. And if you're still on Exchange Server 2003 you may think about skipping Exchange Server 2007. Personally I think it's worth it!

Exchange backups on Windows Server 2008

24 August 2009

by Jaap Wesselius

> ***NTBackup lives no more, replaced by VSS backups. Unfortunately, "Windows Server Backup" didn't work with Exchange Sserver 2007 until now. Exchange Server 2007 Service pack 2 contains a backup plug-in that makes it possible. Although pretty limited compared to full backup application like Microsoft Data Protection Manager or Symantec Backup Exec it does what you want it to do: create backups. Jaap explains.***

In one of my earlier articles (Online Exchange Backups) I discussed the online backup options in Exchange Server 2007, streaming backups using the Backup API (Application Programming Interface) and snapshot backups using VSS. This article was written for Exchange Server 2007 running on Windows Server 2003.

Windows Server 2008 is a different story since it has a new feature called "Windows Server Backup" and NTBackup is discontinued. Windows Server Backup creates Volume Shadow Service (VSS) backups, but it is not Exchange aware. So for backing up your Exchange Server 2007 mailbox databases running on Windows Server 2008 you are dependent on 3rd-party products. No out-of-the-box solution like NTBackup in Windows Server 2003 is available in Windows Server 2008.

Now that Exchange Server 2007 Service Pack 2 is available, this has changed. Amongst other new features and functionality, Microsoft has added a new backup plug-in that makes it possible to create snapshot backups using VSS from Windows Server Backup. I'll show you how.

Windows Server Backup

Windows Server Backup is a feature in Windows Server 2008. You can install Windows Server Backup using the Server Manager, select Features and click "Add Features." Scroll down, select "Windows Server Backup Features" and click Install. If needed you can expand "Windows Server Backup Features" and select the "Command-line Tools" as well. When the setup has finished, the Windows Server Backup utility shows up in the Administrative Tools menu.

When you open Windows Server Backup nothing special appears. There's no indication that you can backup Exchange Server 2007.

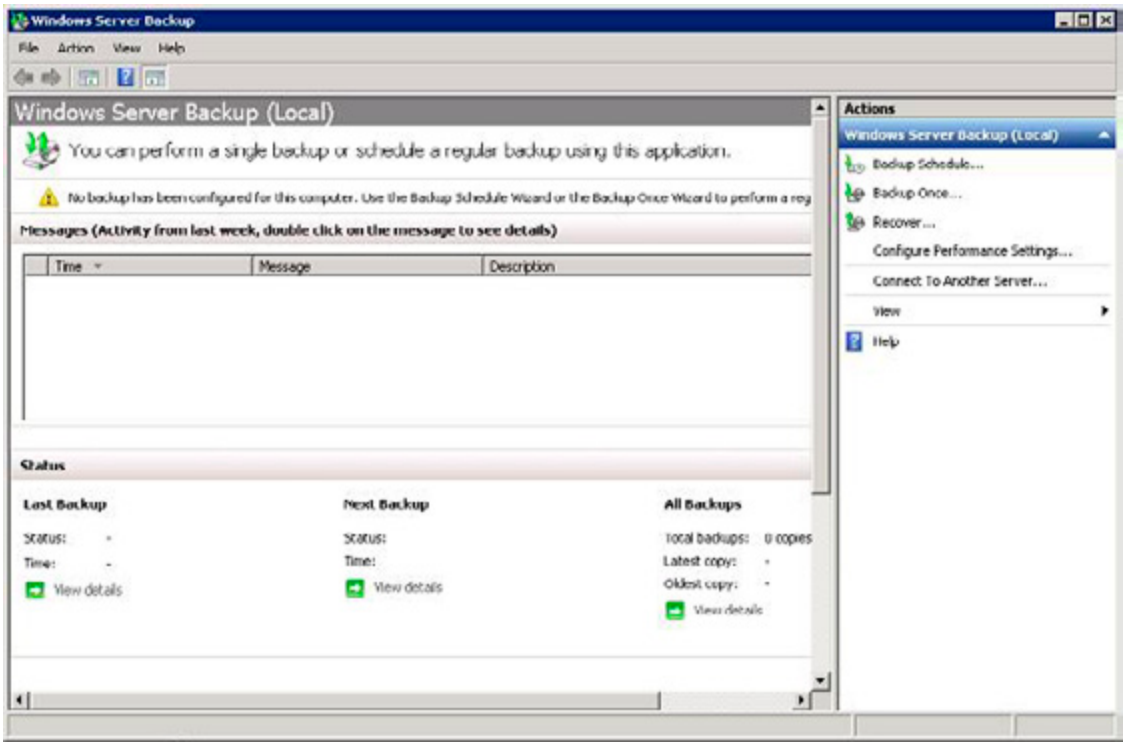

Figure 1. Windows Server Backup. There's no indication about the Exchange awareness.

Windows Server Backup can backup only on Volume level. When you have a default installation you can backup the entire system disk, including the C:\Program Files\Microsoft\Exchange directory. Since Windows Server Backup is aware of Exchange Server 2007 Service Pack 2 the Exchange Databases will be backed up, the database will be checked for consistency and the log files will be purged.

There might be scenarios where the database and the log files are placed on a separate disk. In this example I've located the database in G:\MDB1 and the accompanying log files in G:\LOGS1. The Public Folder database is located in G:\MDB2 and the accompanying log files in G:\LOGS2. This way you can backup only the database and the log files.

In the Windows Server Backup, in the Actions Pane click "Backup Once...." The Backup Once Wizard opens, click Next to continue. You can select a Full Server backup or a Custom backup. Select the Custom Backup if you only want to backup the Exchange database and the log files and click Next to continue.

In the Select Items window you can select the G:\ disk where the database and the log files are located. Remove the check at "Enable System Recovery" to unselect the System disk and click Next to continue.

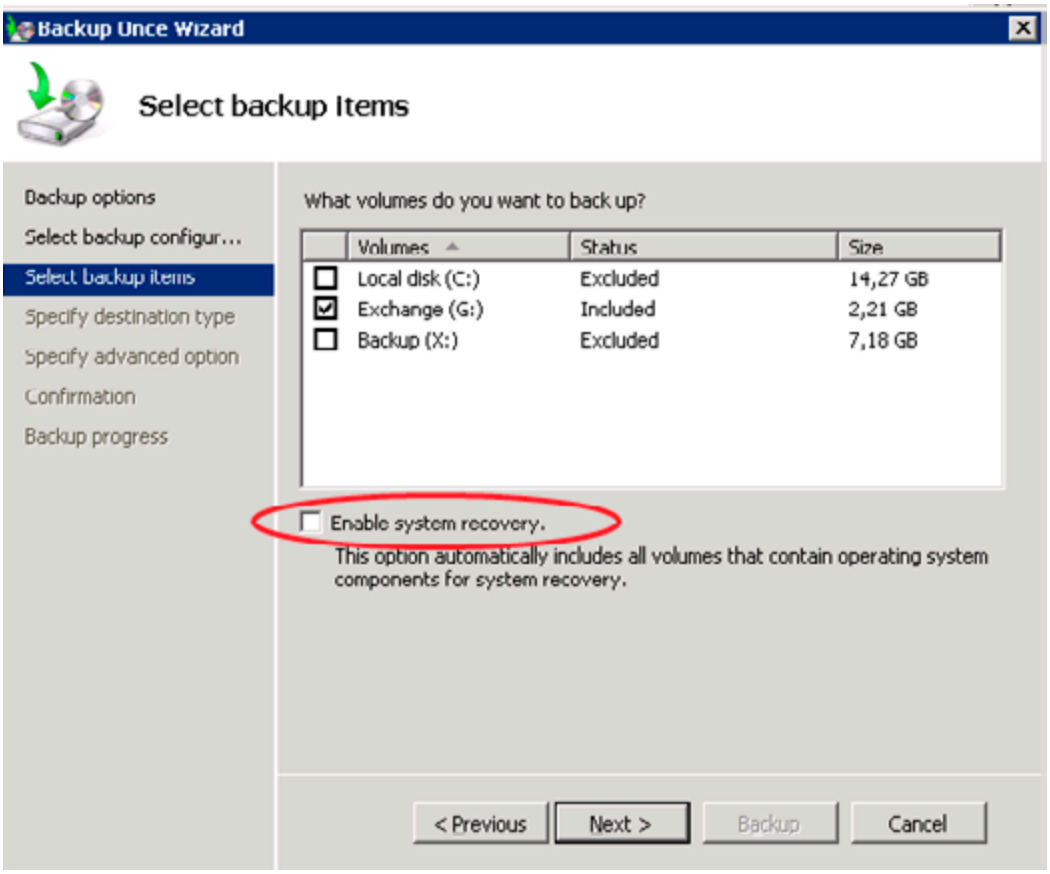

Figure 2. Remove the checkmark at "Enable System Recovery" to deselect the C:\ Drive.

The backup itself can be placed on any disk, except the disk that's being backed up and the System Disk. If you want to backup to disk you have to create an additional disk on the server. It's also possible to backup to a Remote Share. In this example I'll write the backup to another disk, click Next to continue.

Select the Backup destination (X: in this example) and click Next to continue.

It is possible to create a VSS Copy Backup or a VSS Full Backup. A VSS Copy Backup is a full backup of the database, but the header information will not be updated with backup information and the log files will not be deleted. You can select this option when you want to create a backup using Windows Server Backup, but you don't want to interfere with other backup solutions running on this particular server. If you select VSS Full Backup a full backup will be created, the database header information will be updated with the backup information and the log files will be purged (if the consistency check succeeds). In both options the database will be checked for consistency.

Select "VSS Full Backup" to create a full backup of the Exchange database and click Next to continue.

Check the confirmation page and click "Backup" to start the actual backup process.

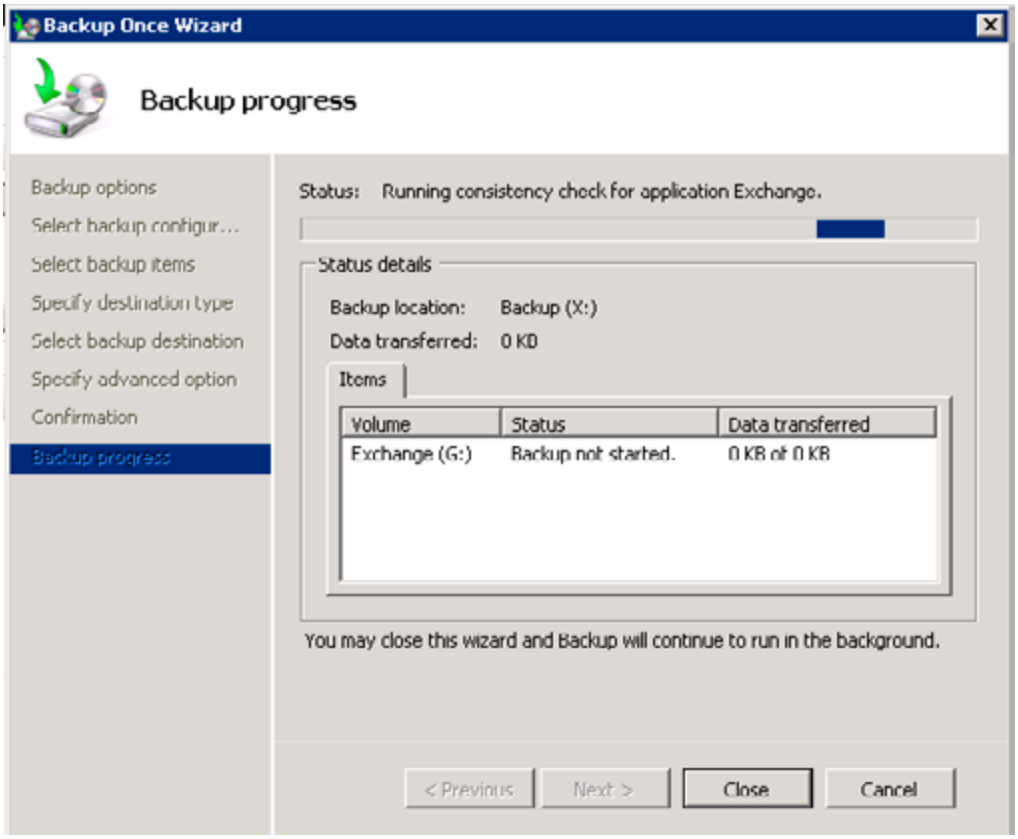

Figure 3. Notice the status bar, this is the only indication you're running an Exchange Server backup.

The only way to see if the actual Exchange backup is performed is by looking at the console. You can see the Status bar: "Running consistency check for application Exchange." You can close the Backup Once Wizard and the backup will continue to run.

If you check the Application log in the Event Viewer you'll see entries about the backup from the Shadow copy service and the Exchange VSS writer about the backup. Please check my other article http://www.simple-talk.com/exchange/exchange-articles/online-exchange-backups/ with detailed information regarding the VSS backup process.

When you check the header information of the Exchange database you'll find information about the backup as well:

```
G:\MDB1>eseutil /mh "mailbox database.edb"

Extensible Storage Engine Utilities for Microsoft(R) Exchange Server
Version 08.02
Copyright (C) Microsoft Corporation. All Rights Reserved.

Initiating FILE DUMP mode...
     Database: mailbox database.edb

    File Type: Database
   DB Signature: Create time:07/02/2009 12:38:52 Rand:172660 Computer:
     cbDbPage: 8192
      dbtime: 1172010 (0x11e22a)
      State: Clean Shutdown
  Log Required: 0-0 (0x0-0x0)
 Log Committed: 0-0 (0x0-0x0)
 Log Signature: Create time:07/06/2009 08:19:40 Rand:239950899 Computer:
   OS Version: (6.0.6001 SP 1)
```

```
Previous Full Backup:
    Log Gen: 208-227 (0xd0-0xe3) - OSSnapshot
      Mark: (0xE4,8,16)
      Mark: 07/20/2009 21:00:08

Previous Incremental Backup:
    Log Gen: 0-0 (0x0-0x0)
      Mark: (0x0,0,0)
      Mark: 00/00/1900 00:00:00

Operation completed successfully in 0.110 seconds.

G:\MDB1>
```

Note. This output has been edited for readability.

Windows Server Backup can only backup active Exchange Servers. This means that it is not aware of a Continuous Cluster Replication (CCR) environment and therefore you cannot backup the passive node of CCR Cluster. Microsoft Data Protection Manager (DPM) 2007 or other 3rd-party products like Backup Exec can backup the passive node and therefore reduce the load of the active node of the cluster.

Backup schedules

It becomes more interesting to use Windows Server Backup when you can schedule backups and create a backup once a day for example without any hassle.

Open the Windows Server Backup application and in the Actions pane select "Backup schedule…." Like the Backup Once wizard you can select a Full Server backup or a Custom Backup. If you select Custom Backup it is possible to select the disk containing the Exchange Server database and log files. When creating a backup schedule the system volume will be included by default, it is not possible to deselect this.

The next step is to specify the backup time, you can create a backup once a day or multiple times a day.

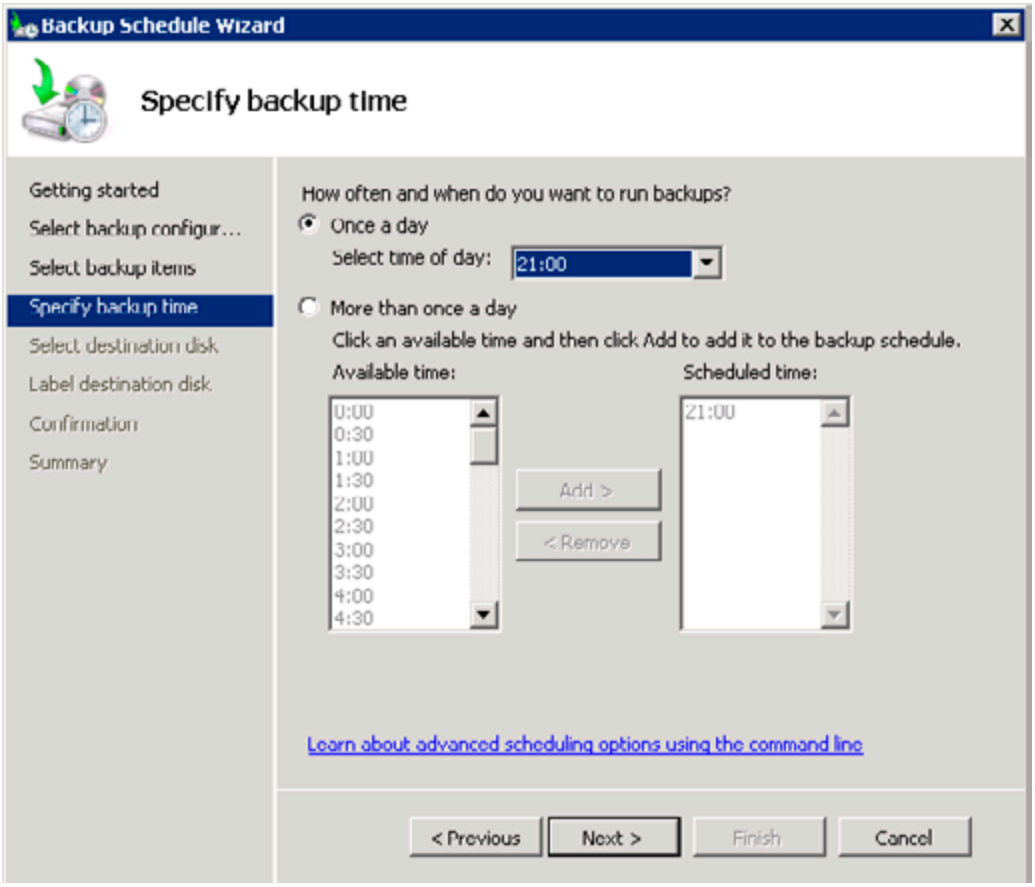

Figure 4. A backup will be created at 9pm. It is also possible to create multiple backups.

When selecting the destination disk it will be formatted by Windows Server Backup and all information on the disk will be lost. The drive doesn't get a drive letter during format, but you can still see it through Disk Management in the Server Manager.

Note

Although you can select "Always create an incremental backup" in the Windows Server Backup application it is not possible to do this for Exchange. You can only create full backups, either in the Backup Once and in the Backup Schedule option.

Restore the Exchange Backup

When you want to restore the previous backup open the Windows Server Backup and click on "Recover..." in the Actions Pane. You can select what data you want to recover, in this example select "This Server 2007MBX01." Click Next to continue.

If you have setup a backup schedule you can select a date and time of the backup you want to restore.

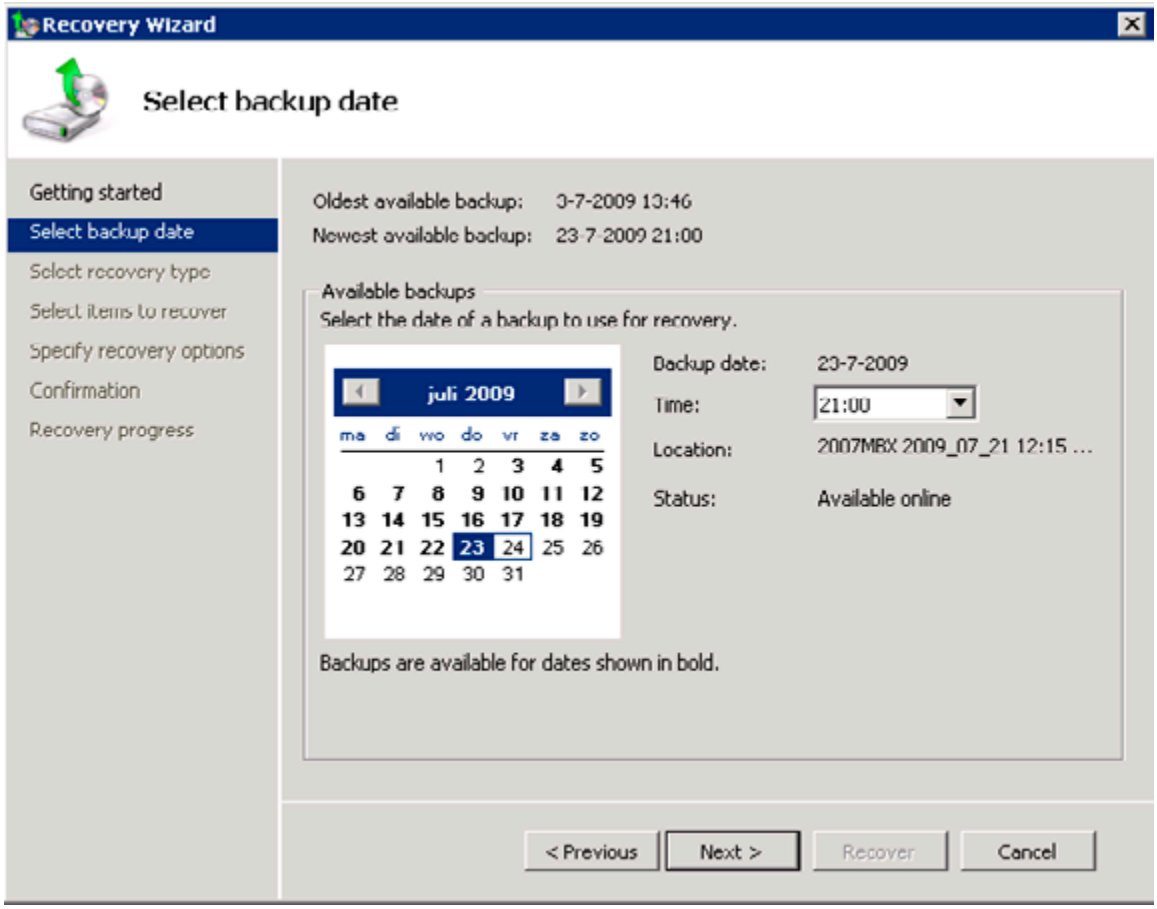

Figure 5. Select the backup you want to restore.

Click Next to continue.

In the Recovery Type window you can select the kind of information that needs to be restored. Since this is an Exchange database restore select "Applications" and click Next to continue.

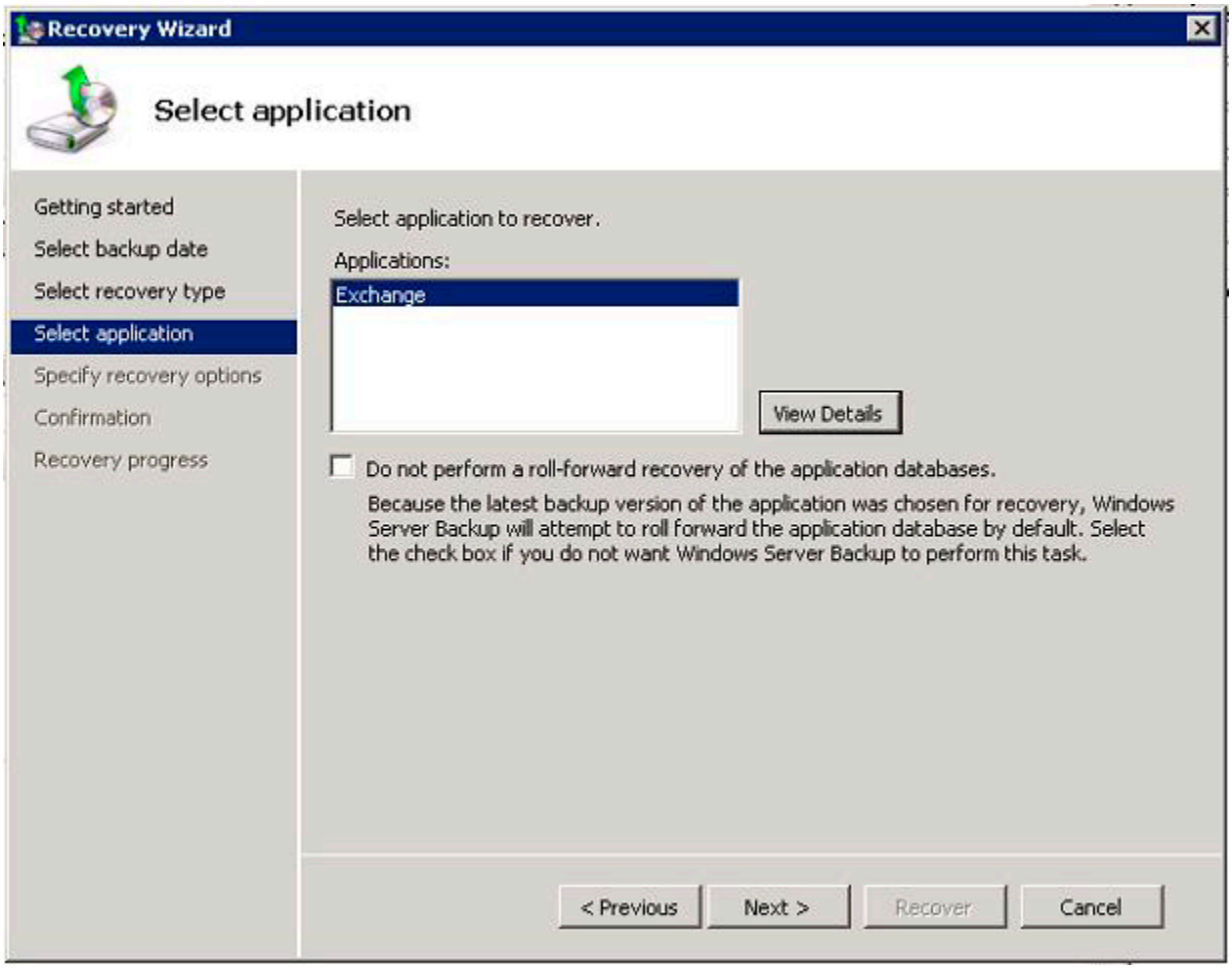

Figure 6. Select the application you want to restore information for. If you do not want to roll-forward log files tick the checkbox.

Select "Recover to original location" to restore the database to its original location. Do not forget to tick the checkbox "This database can be overwritten by a restore" in the database properties in Exchange Management Console. If you fail to do so the restore of the database will fail.

In the Confirmation window check your selections and click "Recover" to start the restore process.

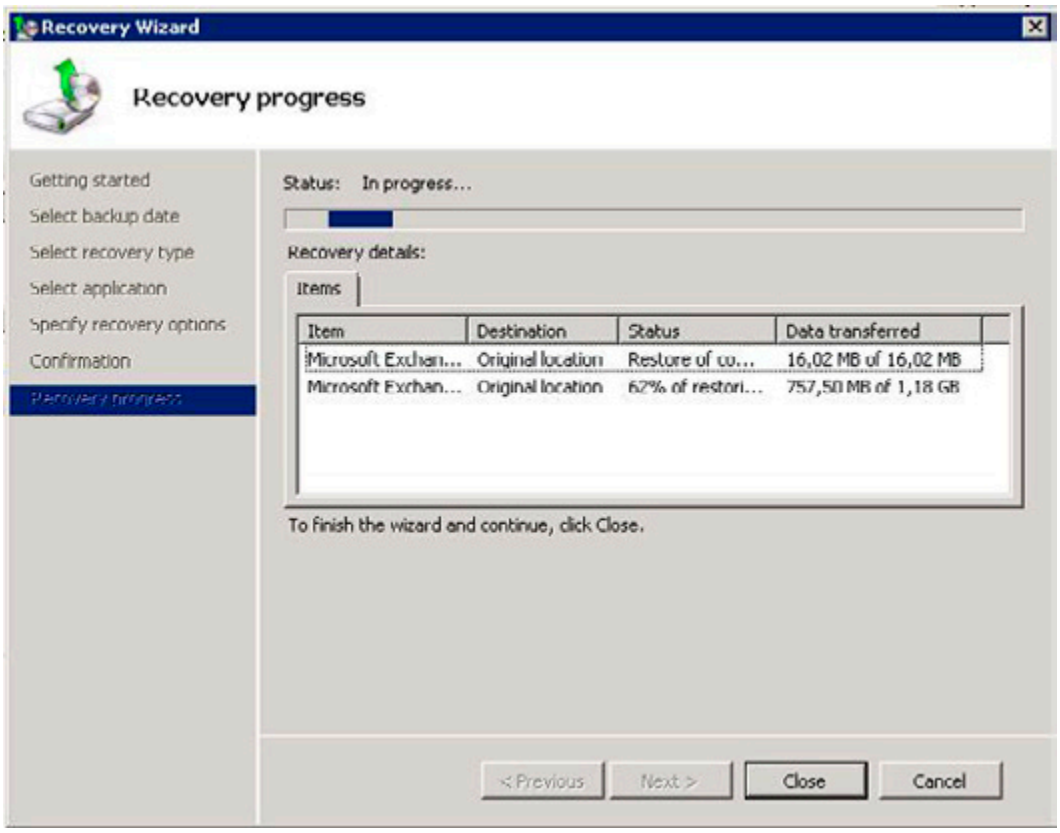

Figure 7. You can check the progress of the restore process.

Information stored in log files newer than the database (and thus not restored with the database) will be automatically rolled-forward. When possible the recovery process will automatically replay, or roll-forward all available log files from the restore point up to the latest possible point in time.

There's an easy way to check this. Create a full backup of your mailbox database. When the backup is successfully finished send a couple of message to yourself. Make them easy to identify. Dismount the database and restore the database from the backup. The messages you just sent are not in this backup, but in the log files written after the creation of the backup. These log files will be automatically replayed. When you logon to the mailbox after the restore of the database you'll find the messages again, even if they weren't in the actual backup.

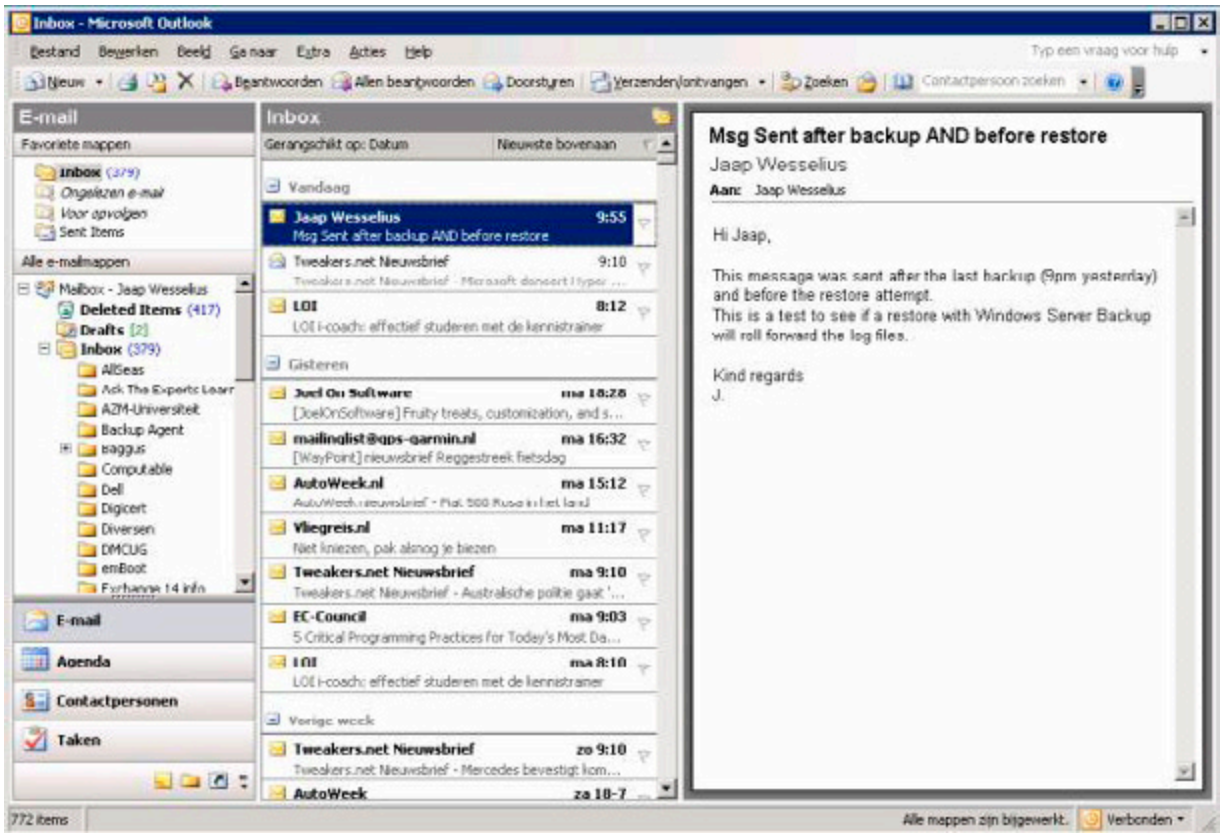

Figure 8. The last 3 messages were not part of the backup but were rolled-forward from the log files.

Conclusion

Exchange Server 2007 Service pack 2 will contain a backup plug-in that makes it possible to use Windows Server Backup in Windows Server 2008 to create Exchange Server backups. Although pretty limited compared to full backup application like Microsoft Data Protection Manager or Symantec Backup Exec it does what you want it to do: create backups. You can create backups manually or create a schedule for the backup application to run. It is a backup to disk solution, backup to tape is not available.

It can restore backups and it can roll forward log files after a restore, so the basic functionality is available. Nothing more, nothing less.

Note

Fellow MVP Michael B. Smith created a couple of scripts that are capable of creating backups of Exchange Server 2007 running on Windows Server 2008, without the need of Windows Server Backup. YOU CAN FIND HIS SOLUTION HERE.

Moving to Office Communications Server 2007 R2 – Part 2

24 August 2009

by Desmond Lee

In the second part of his article on Moving to Office Communications Server 2007 R2, Desmond looks into recommendations to standardize the client desktop to support the different R2 applications, other infrastructure concerns and migration of the remaining internal server roles before moving out to the Edge servers in the screened network.

In a nutshell, you will need to set up R2 server roles alongside the RTM counterparts to provide equivalent functionality. This is because not all RTM and R2 server roles can interoperate with one another directly without first sorting out certain changes which are vital for smooth operations. Therefore, the different RTM server roles should remain in operations until UC-enabled users are fully migrated and homed on new R2 server pools. Only then should you considering bringing in R2 specific server roles such as Group Chat and Outside Voice Control after the complete RTM infrastructure has been migrated to R2. In part 2 of this article (see Part 1), I shall take you through some of the highlights of how to implement an Enterprise consolidated topology from the inside-out in a side-by-side migration from OCS (Office Communications Server) 2007 RTM.

Back- to Front-End

Security

As part of security patch management best practice, do not forget to test and apply the most current patches, update packages and hotfixes prior to full scale production rollout. For example, you can download and install the July 2009 hotfix rollup packages for OCS 2007 R2 server roles (KB968802), Communicator 2007 R2 (KB969695) and R2 Attendant (KB970275). If you do not have an automated patch management system in-house, think about using Microsoft Update to deliver update packages to your R2 back-end servers as well as R2 clients. This functionality has been extended to the OCS family of products around the May 2009 timeframe. To maintain absolute control over patch delivery and actual patch application (what, when, why and how) – particularly crucial on the R2 back-end servers – manually configure the option to read "Notify me but don't automatically download or install them." Better still, deploy the free WSUS 3.0 SP1 solution with Active Directory Group Policy to regain centralized control of your patch management needs.

Before you start, check that you always login with an account that has membership in the Domain Admins, RTCUniversalServerAdmins as well as local Administrators group. You will find that not all but some of the steps in the R2 setup will fail if this condition is not fulfilled. With the named SQL database instance on the Back-End database ready, you can go ahead to run the Create Enterprise Pool task. This step should be carried out right on the back-end database server running x64 editions of SQL Server. For those of you who are still using the 32-bit version or SQL Server 2005 SP2 (in particular) as the Back-End database, you will have to install the SQL Server client tools on the new Front-End server before proceeding further. The database itself will be created for you when the R2 Enterprise pool is created.

A number of shared folders must already be prepared ahead of time – preferably on a dedicated file server – to support client and device updates as well as data files for meeting presentations, Address Book Service (ABS), etc. You may find that deploying them directly on a Front-End server may just work for you in your environment. Following that, you install the first R2 Front-End server to the pool and activate it. For administrators that are responsible for both the RTM and R2 infrastructure, it would be a good idea to use the same RTC* accounts in Active Directory. I suggest that you let the R2 setup program to create these service accounts (if they do not already exist)

so that the correct group membership and security settings can be applied. These service accounts do not have their "Password does not expire" option enabled by default. You can prevent a number of seemingly unexplained incidents that potentially arise after R2 is in production for a period of time by simply changing this initial setting (verify that this is in accordance with your corporate security policy though).

Subsequently, you assign the digital certificate that you prepared earlier to the Enterprise Edition server. Although the Web Components Server resides on the same physical machine as the Front-End server in a consolidated topology, you must make use of the Internet Information Services (IIS) Web Server Certificate Wizard to separately assign the certificate to it. If the design calls for the same FQDN to be deployed for the Enterprise Pool and Web Components, it is a simple matter of selecting the same certificate in the drop down box. This is possible by virtue of the saved certificate in the local computer certificate store.

Once you verified that AD has replicated changes in the domain, ascertain that Windows Firewall is already running before attempting to execute the Start Services Wizard. The big advantage of this is that R2 setup will automatically configure all the necessary exception rules in the firewall for you. In case Windows firewall is turned off or your organization has deployed other host-based firewall solutions as a standard, it will be necessary for you to make the adjustments manually. To identify the ports to open in the firewall, search through the installation log files for the string "firewall exceptions."

Communications Web Access

The optional CWA server role provides UC-enabled users the capability to utilize almost the entire feature set of the matching version of Front-End home server from a supported web browser (Internet Explorer, Mozilla Firefox, Apple Safari, etc.). There is no requirement to have a prior installation of the full MOC client application on the machine or device. One usage scenario that is gaining popularity is to initiate an IM session on your Apple iPhone 3Gs by simply using the in-built browser itself.

You must setup a new R2 CWA for users homed on R2 Front-End servers and retain the RTM version of CWA server as long as users remain on the older version. In order to maintain a consistent experience for all users, you can keep the current CWA URL by modifying the existing DNS record to point to the R2 CWA machine. The latter will then automatically redirect users to the corresponding version of CWA base on their actual home server. Similarly, you will find that you do not need to re-publish the URL in Active Directory as the `dialin` and `join` web pages for users to schedule and join conferences will remain unchanged i.e. https://im.swissitpro.ch/dialin and https://im.swissitpro.ch/join. The downside to this is that a user homed on RTM Front-End server will be prompted to sign in again.

You can use Table 1 as a guide to configure the appropriate internal DNS records to support such a scenario. Here, the CWA URL is identified as https://im.swissitpro.ch served by a server located behind the corporate firewall. This maps to a website or more commonly a CWA virtual server on IIS. The CNAME records `download` and `as` are essential to support the desktop sharing feature introduced in R2.

Fqdn	Dns record type	Remarks
cwa.swissitpro.ch	Host (A)	IP address of RTM CWA
cwar2.swissitpro.ch	Host (A)	IP address of R2CWA
im.swissitpro.ch	Canonical name (CNAME)	CWA URL, modify to reference cwar2.swissitpro.ch from cwa.swissitpro.ch
download.im.swissitpro.ch	Canonical name (CNAME)	references im.swissitpro.ch
as.im.swissitpro.ch	Canonical name (CNAME)	references im.swissitpro.ch

Table 1: Sample DNS records (CWA URL)

To improve the availability and capacity of CWA servers, you can deploy dedicated hardware load balancers (HLB) that are not shared with other R2 or RTM server roles, such as the Front-End Pool or Director arrays. In this case, you should edit the CNAME record for the CWA URL to resolve to the physical IP address of the load balancer.

One important fact to know is that it is no longer possible to add RTM version of the CWA server role once you get pass the R2 Active Directory preparation steps. To manage and administer the R2 release of CWA, you must also install the Communicator Web Access

snap-in (x64-only) separately after satisfying the requirements as discussed here. As a security best practice, you should create two virtual servers, one on each host computer, to segregate external and internal users from one another. Further, consider deploying a reverse proxy server in the perimeter network (DMZ) to publish the CWA virtual server that is designated to handle traffic from external users.

A commonly asked question concerns the collocation of the CWA with other R2 server roles on the same box. This is not an officially supported scenario (see Supported Server Role Collocation). Should you somehow decide to go ahead and setup CWA on the Front-End server for instance, you will have to explicitly resolve the conflict of both the CWA and Web Access Services that arise from using the same default TCP port 443.

The Director and Edge

The Director shields the CWA server and Front-End server pool from incoming traffic with another layer of security. This particular server role does not home any users and functions as a next-hop server. Internal users hitting the Director will be redirected to the respective home pool whereas traffic will be proxied for remote users coming through the Edge server. In an environment with more than one pool in production, deploying the Director will spare you from having to configure one of the pools to perform the proxy task and the risk of performance degradation (see Deploy a Director on how to do this). Existing RTM Director and Edge server roles can communicate with R2 server pools and clients. However, R2 editions of the Director and Edge server roles do not work with the RTM server pools. As a result, upgrade of these two server roles at the same time to support both RTM and R2 server pools may be essential early in your project plan to ease transition for users already homed on R2 servers.

DNS Records

According to published documentation, the Director (RTM) server role is required to support a mixed environment of RTM and R2 Enterprise Edition server pools in the same SIP domain. This setup facilitates automatic configuration and sign-in of internal MOC 2007 RTM and R2 clients simply by querying for the `_sipinternaltls._tcp` *DNS SRV record that points to a single pool of servers. For users homed on the R2 pool, they will automatically be redirected after being authenticated on the RTM pool.*

If you do not have or intend to introduce a Director (RTM) server role, you can work around this constraint by using manual client configuration or push out the settings through Group Policy for MOC R2 users. At a minimum, you will need to create new A host records in DNS describing the FQDN of each R2 pool. If you only have a single Enterprise Edition server without a hardware load balancer (HLB), point the A record to the IP address of the host machine. Otherwise, configure the A record with the virtual IP address of the HLB. Subsequently, you can update the DNS SRV to point to the new R2 Director once the RTM version is decommissioned. See DNS Requirements for Enterprise Pools and Standard Edition Servers for more information.

You should be aware of restrictions such as public IM and federation for side-by-side configurations in the same organization and SIP domain. Remote users with MOC R2 clients will not be able to communicate until the R2 versions of the Director and Edge server roles are operational. This is another good reason to hold back deployment of MOC R2 for users that need to consume any R2 services outside the corporate firewall in the absence of a VPN.

With or without hardware load balancer, RTM and R2 server roles cannot co-exist in the same Front-End Pool, CWA server pools, an Edge array or Director array. However, the same HLB can front RTM and R2 Enterprise Pools at the same time. As with the R2 Front-End configuration, the Access Edge, Web Conferencing Edge and Audio/Video Edge services now reside on the same physical machine in a consolidated Edge topology. In any case, only the Access Edge server role at the data center is active and responsible for SIP signaling traffic for the entire organization. This is regardless of how many branch office Edge servers you may have out there. On top of the classic DNS and certificate dependencies, I also want to draw your attention to how DNAT and SNAT could affect your firewall deployments (destination and source network address translation). See Planning for External User Access for more information on this topic.

Monitoring and Compliance

Because the Quality of Experience (QoE) Monitoring Server is deprecated in R2, this particular server role should not be removed from the environment if you need to continue collecting Call Detail Record (CDR) metrics from users that are still homed on RTM server pools. Tracking CDR collects statistical usage data without saving details of the actual conversation itself.

Note that in R2, both the QoE data collection and CDR services now reside on the new Monitoring Server role. Previously, the CDR service is an integral part of the Archiving and CDR Server (RTM). For legal and compliance requirements, your organization may be obliged to log the contents of every communication taking place in the network. As all Instant Messaging conversations must traverse through the users" home servers, you can enforce this transparently at the Archiving Server level. If you intend to start data capture from day one, you should already install, activate and configure the R2 Monitoring Server as well as R2 Archiving Server before you home any users on R2 Front-End pools.

You should double check that the pool is correctly associated with the deployed R2 Monitoring Server when you enable the CDR and QoE data collection option. Once again, you will have to maintain the existing Archiving and CDR Server (RTM) and build a different Archiving Server (R2) with its own SQL Server database (on the same or another dedicated machine) to facilitate logging conversations on the associated R2 Front-End pools. Additionally, you must manually install the Message Queuing (MSMQ) service on all R2 Front-End and Standard Edition servers, Monitoring Server and Archiving Server itself prior to deploying these 2 latter server roles. I propose that you consult the Archiving and Monitoring for Office Communications Server 2007 R2 for an overview of the archiving and monitoring architecture and lookout for the various agents and service components to make this work correctly.

The optional Monitoring Server Report Pack greatly simplifies generation of captured data and is highly recommended. To this end, you must install the SQL Server Reporting Services from the deployed version of SQL Server (2005 SP2 or 2008). For installations using SQL Server 2005 SP2, go ahead and apply both the KB942662 and KB940382 updates that are required for reporting to function correctly.

Enterprise Voice

The Quality of Experience Monitoring Server Administration Tool (RTM) does not have any built-in capability to associate a RTM Mediation server with a R2 Enterprise Pool. As such, you must take extra steps to realize this by applying KB956829 followed by the Microsoft provided script (Associate.vbs). The purpose is to enable QoE reports to be delivered to the R2 Monitoring Server that is bound to its R2 Standard or Enterprise Edition Pool in a coexistence environment. For this to work, UC-enabled users must be homed on the R2 Pool. Like the other RTM server roles, the QoE Monitoring Server can be decommissioned once you are completely migrated to the R2 counterparts.

For those of you who intend to integrate R2 with the Unified Messaging feature in Exchange Server 2007 SP1 or an existing PBX, you should closely follow the TechNet Library documentation entitled Deploying Enterprise Voice. The outlined steps apply equally regardless of whether you have a current RTM setup or not. Like the different R2 setup steps during installation, it is safe to re-run scripts and applications that you have executed previously e.g. from exchucutil.ps1, ExchUMutil.exe to OCSUMutil.exe. Obviously you should carefully review the installation logs for hints in resolving any warnings or errors that surface. Configuration information from location profiles, dial plans to Exchange AutoAttendant functionality should continue to work without modifications. Nevertheless, you will want to perform tasks such as defining a new Exchange UM IP Gateway object and point it to the new R2 Mediation Server to optimize traffic routing, for instance.

Conclusion

The principle behind a side-by-side migration is rather straightforward and can be tedious nonetheless. It mirrors that of a brand new R2 installation parallel to your production OCS 2007 RTM deployment. In this final installation of the article, you have learnt that to achieve this goal, you can essentially adopt the same set of official R2 deployment guidelines with a careful lookout for any gotchas.

In a typical inside-out migration, you start by deploying the Front-End Pool and gradually introduce additional internal R2 server roles from CWA, Monitoring, Archiving to Director that is appropriate to your organization. Along the way, you should validate and run functional tests to confirm that everything works as expected with the necessary adjustments. Until the supporting R2 Director and Edge server roles are in operations, UC-enabled users homed on R2 Front-End Pool will be blocked from access remotely.

As soon as the infrastructure has been fully migrated to R2, you should take the next step to move UC-enabled users to home on R2 server pools. To take immediate advantage of new R2 features, migration to the MOC R2 client is the next logical thing to do. Eventually, you can gradually decommission remaining RTM servers before planning on adding new R2 server roles to the equation. Taking your time to plan, test and execute the migration in manageable phases will go a long way to keep the project team and end users happy.

The Client Side Story

Besides Office Outlook 2007, you will find that the majority of users will primarily be using the MOC R2 and Office Live Meeting 2007 *clients to conduct their day to day business from instant messaging, presence, peer-to-peer video and voice call to web conferencing. With additional R2 productivity features, the new Office Communications Server 2007 R2 Conferencing Attendant and Office Communicator 2007 R2 Group Chat clients should be listed in your deployment plans as well. For those of you on the RTM version of MOC, you can directly upgrade to the R2 equivalent without first uninstalling the former. It is important to understand that users using the R2 version of MOC can only be homed on an R2 Enterprise pool or Standard Edition server, although users running MOC RTM can also be hosted on the latter two. In other words, you can first move users running MOC RTM to home on R2 pools and defer upgrading to the MOC R2 client to a later date.*
Taken together, all of the R2 client applications rely on different minimum versions of Windows service packs and supporting components to operate correctly. To take full advantage of the latest functionality, simplify rollout and on-going maintenance, I recommend that you consider standardizing at least on the following:
Windows XP SP3 *or* Windows Vista SP1

Microsoft Core XML Services (MSXML) 6.0 SP1 *(ships as part of XP SP3 and Vista SP1)*

.NET Framework 3.5 SP1

*.NET Framework 3.5 Service Pack 1 Update (*KB959209)

Conferencing Add-in for Microsoft Office Outlook

Microsoft Office Communicator 2007 R2 client

*Live Meeting 2007 client update package (*KB960165)

*Conferencing Add-in for Outlook update (*KB960164)

Office Communicator Web Access Plug-in (Desktop-sharing feature)

Conferencing Attendant application (optional for specific groups of users; requires the Office Communicator Web Access client to provide Dial-in Conferencing Web page for participants).

*In contrast with the back-end R2 server roles, Windows client computers are not required to be 64-bit capable to run any of the R2 client applications. Although Windows 7 as well as Service Pack 2 for Windows Vista have recently been released (*KB948465*), IT shops may still need to devote sufficient time to test for compatibility with their line-of-business applications. Besides, Vista SP1 is a prerequisite for SP2. Unless your Vista client desktops are already deployed with SP1, it makes sense to start off at this base level before pushing out SP2 within the next 24 months (see* Lifecycle Supported Service Packs *for more information).*

Monitoring and Scheduling Exchange 2007 Database Online Maintenance

25 August 2009

by Ben Lye

> ***To keep the Exchange database healthy, it is important to make sure that online maintenance and online defragmentation are running properly, or are at least conforming to Microsoft's recommendations. Ben shows how easy it is to automate the maintenance and monitor the defragmentation task.***

Exchange database maintenance is an important part of keeping Exchange healthy. Exchange Server automates online maintenance tasks and runs them based on the schedule specified by the Exchange administrator.

The online maintenance tasks are:

- Purge mailbox database and public folder database indexes.
- Maintain tombstones.
- Clean up the deleted items dumpster.
- Remove public folders that have exceeded the expiry time.
- Remove deleted public folders that have exceeded the tombstone lifetime.
- Clean up conflicting public folder messages.
- Update server versions.
- Check Schedule+ Free Busy and Offline Address Book folders.
- Clean up deleted mailboxes.
- Clean up reliable event tables.

For detailed information on these tasks you can refer to Microsoft TechNet: Maintaining Mailbox Databases and Public Folder Databases.

If one of these maintenance tasks is performed on a database then online defragmentation will be performed on that database.

If the online maintenance tasks are unable to complete in a single schedule window the tasks will be suspended and then resumed in the next window. In this way maintenance is guaranteed to eventually complete, however it is important to make sure that the scheduled maintenance windows are properly configured so that the maintenance tasks are able to complete regularly.

When scheduling online maintenance there are several guidelines to consider:

- Online maintenance should be scheduled for times when there is little activity on the database.

- Online maintenance must not run at the same time as a backup (online defragmentation cannot start while a backup is in progress).
- Online defragmentation should be able to complete at least once every two weeks.
- Online maintenance schedules for databases in the same storage group should not overlap (Microsoft recommends a 15 minute gap between maintenance schedules).

The default online maintenance schedule is nightly from 1 am to 5 am. To aid in customizing the online maintenance schedule, the Windows event log can be used to see how often online defragmentation is completing and Performance Monitor counters can be used to check the efficiency of online defragmentation. This data can be used to adjust the online maintenance schedule to give more or less time.

Event Log Entries for Online Defragmentation

There are five events relating to online defragmentation starting and stopping. The events are logged in the Application log with a source of "ESE" and a category of "Online Defragmentation'.

Event ID	Description
700	Online defragmentation is beginning a full pass
701	Online defragmentation has completed a full pass
702	Online defragmentation is resuming defragmentation
703	Online defragmentation has completed a resumed pass
704	Online defragmentation has been interrupted

Events 701 and 703 indicate a complete pass. In the case of event 703, completion of a resumed pass, the event text will include information about how long defragmentation took to complete, and how many times it was invoked.

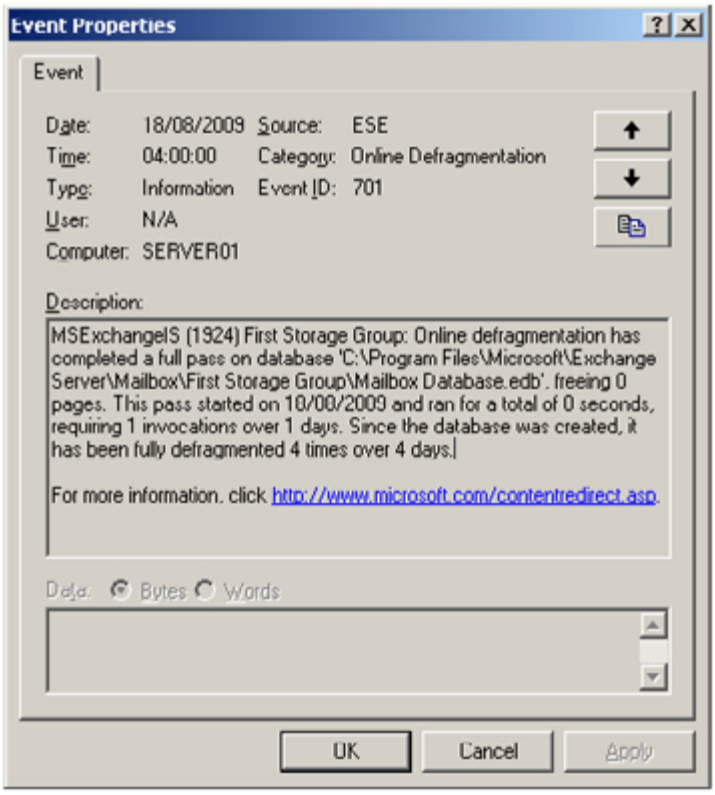

To make it easier to check how often online maintenance completes the Exchange Management Console and a PowerShell script can be used to parse the Application log.

This script will search the event log for online defragmentation messages for each database on the server on which it's run (or the clustered mailbox server if running on a cluster node) and return the amount of time taken to complete online defragmentation of each database.

```
# Script to check the status of Exchange database online defragmentation tasks
# Written by Ben Lye
#
# The script will parse the event log of the local machine looking for online
# defrag related messages. Messages are parsed to determine when online defrag
# last finished for each database and how long it took to complete.
#
# The script needs to be run on an Exchange 2007 mailbox server or mailbox
# cluster node If the script is run on a cluster node it should be the active
# node, or event log replication needs to be enabled (this is the default).


# The $records variable defines the number of events to retrieve from the log
# It can be increased or decreased according to the needs of a particular server.
# The script will run faster if fewer records are retrieved, but data may not be
# found for all databases.
$records = 10000

# Get the hostname
$hostname = Get-Content env:computername

# Check if the local machine is an Exchange mailbox server
```

```
$mbserver = Get-MailboxServer -Identity $hostname -ErrorAction SilentlyContinue

# Check if the local machine is a member of a mailbox cluster
$cms = Get-ClusteredMailboxServerStatus -ErrorAction SilentlyContinue

# Exit the script if the local machine is not a mailbox server or a CMS node
if (-not $mbserver -and -not $cms) {
   Write-Host "The machine $hostname is not a server an Exchange mailbox server." `
      -ForegroundColor Red
   Write-Host "This script must be run on a mailbox server or mailbox cluster node." `
      -ForegroundColor Red
   break
}

# Determine the server name to enumerate the databases
if ($cms) {
   # This server is a cluster node, the database server name is the name of the CMS
   $dbserver = $cms.ClusteredMailboxServerName
} else {
   # This server is a mailbox server - the database server name is the local hostname
   $dbserver = $hostname
}

# Get the mailbox databases from the server
$mbdatabases = Get-MailboxDatabase -Server $dbserver `
   | Sort-Object -Property Name

# Get the public folder databases from the server
$pfdatabases = Get-PublicFolderDatabase -Server $dbserver `
   | Sort-Object -Property Name

# Create an array for the databases
$databases = @()

# Check if mailbox databases were found on the server
If ($mbdatabases) {
   # Loop through the databases
   ForEach ($mdb in $mbdatabases) {
      # Create an object to store information about the database
      $db = "" | Select-Object Name,Identity,EdbFilePath,DefragStart,DefragEnd, `
         DefragDuration,DefragInvocations,DefragDays

      # Populate the object
      $db.Name = $mdb.Name.ToString()
      $db.Identity = $mdb.Identity.ToString()
      $db.EdbFilePath = $mdb.EdbFilePath.ToString()

      # Add this database to the array
      $databases = $databases + $db
   }
}

# Check if public folder databases were found on the server
If ($pfdatabases) {
   # Loop through the databases
   ForEach ($pfdb in $pfdatabases) {
```

```
      # Create an object to store information about the database
      $db = "" | Select-Object Name,Identity,EdbFilePath,DefragStart,DefragEnd, `
         DefragDuration,DefragInvocations,DefragDays

      # Populate the object
      $db.Name = $pfdb.Name.ToString()
      $db.Identity = $pfdb.Identity.ToString()
      $db.EdbFilePath = $pfdb.EdbFilePath.ToString()

      # Add this database to the array
      $databases = $databases + $db
   }
}

# Retrieve the events from the local Application log, filter them for ESE messages
$logs = Get-EventLog -LogName Application -Newest $records | `
   Where {$_.Source -eq "ESE" -and $_.Category -eq "Online Defragmentation"}

# Create an array for the output
$output = @()

# Loop through each of the databases and search the event logs for relevant messages
ForEach ($db in $databases) {

   # Create the search string to look for in the Message property of each log entry
   $s = "*" + $db.EdbFilePath + "*"

   # Search for an event 701 or 703, meaning that online defragmentation finished
   $end = $logs | where {
      $_.Message -like "$s" -and ($_.InstanceID -eq 701 -or $_.InstanceID -eq 703)
   } | select-object -First 1

   # Search for the first event 700 which preceeds the finished event
   $start = $logs | where {
      $_.Message -like "$s" -and $_.InstanceID -eq 700 -and $_.Index -le $end.Index
   } | select-object -First 1

   # Make sure we found both a start and an end message
   if ($start -and $end) {

      # Get the start and end times
      $db.DefragStart = Get-Date($start.TimeGenerated)
      $db.DefragEnd = Get-Date($end.TimeGenerated)

      # Parse the end event message for the number of seconds defragmentation ran for
      $end.Message -match "total of .* seconds" >$null
      $db.DefragDuration = $Matches[0].Split(" ")[2]

      # Parse the end event message for the number of invocations and days
      $end.Message -match "requiring .* invocations over .* days" >$null
      $db.DefragInvocations = $Matches[0].Split(" ")[1]
      $db.DefragDays = $Matches[0].Split(" ")[4]

   } else {

      # Output a message if start and end events weren't found
```

```
        Write-Host "Unable to find start and end events for database," $db.Identity `
           -ForegroundColor Yellow
        Write-Host "You probably need to increase the value of `$records." `
           -ForegroundColor Yellow
        Write-Host

    }
    # Add the data for this database to the output
    $output = $output + $db
}

# Print the output
$output
```

Microsoft recommends that in large organisations with large databases (150–200GB) and many storage groups (up to 20) on a single server online defragmentation should complete at least once every two weeks for each database. In smaller organizations with smaller databases it should complete more often.

In either case, if online defragmentation is completing within two days then it is probably safe to shorten the online maintenance window for the database. If defragmentation is not completing within 14 days the online maintenance window should be lengthened.

Performance Monitor Counters for Online Defragmentation

Exchange 2007 includes the following performance counters for monitoring online defragmentation:

- MSExchange Database ==> Instances\Online Defrag Average Log Bytes.
- MSExchange Database ==> Instances \Online Defrag Log Records/sec.
- MSExchange Database ==> Instances \Online Defrag Pages Dirtied/sec.
- MSExchange Database ==> Instances \Online Defrag Pages Preread/sec.
- MSExchange Database ==> Instances \Online Defrag Pages Read/sec.
- MSExchange Database ==> Instances \Online Defrag Pages Re-Dirtied/sec.
- MSExchange Database ==> Instances \Online Defrag Pages Referenced/sec.

Exchange 2007 Service Pack 1 adds these two additional counters.

- MSExchange Database ==> Instances \Online Defrag Pages Freed/Sec.
- MSExchange Database ==> Instances \Online Defrag Data Moves/Sec.

The two interesting counters are "Online Defrag Pages Read/sec" and "Online Defrag Pages Freed/Sec'. These two counters can be monitored during an online maintenance window and the average values compared to determine if the window should be increased or decreased.

If the ratio of Pages Read:Pages Freed is greater than 100:1 then the online maintenance window can be decreased, if the ratio is less than 50:1 then the maintenance window should be increased, and if the ratio is between 100:1 and 50:1 there is no need to change the window.

To use these counters extended ESE performance counters must be enabled, which is done by adding a new registry value.

Note

Incorrectly editing the registry can cause serious problems that may require you to reinstall your operating system. Problems resulting from editing the registry incorrectly may not be able to be resolved. Before editing the registry, back up any valuable data.

- Start the registry editor on your Exchange 2007 Mailbox server.
- Locate the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ESE\Performance registry key.
- Right-click Performance, select "New," and then select "DWORD value."
- Name the new DWORD value "Show Advanced Counters."
- Double-click "Show Advanced Counters."
- In the "Value data" field, enter 1.
- Close the registry editor.

Once the extended ESE performance counters are enabled Performance Monitor can be used to log the counter data to a file.

- Start Performance Monitor on the Exchange 2007 mailbox server (or the active cluster node in a CCR cluster) by clicking Start ▸ Programs ▸ Administrative Tools ▸ Performance.
- Expand "Performance Logs and Alerts" and select "Counter Logs."
- Right-click "Counter Logs" and select "New log settings."
- Enter a name for the new counter, such as "EDB Defrag," and click "OK."

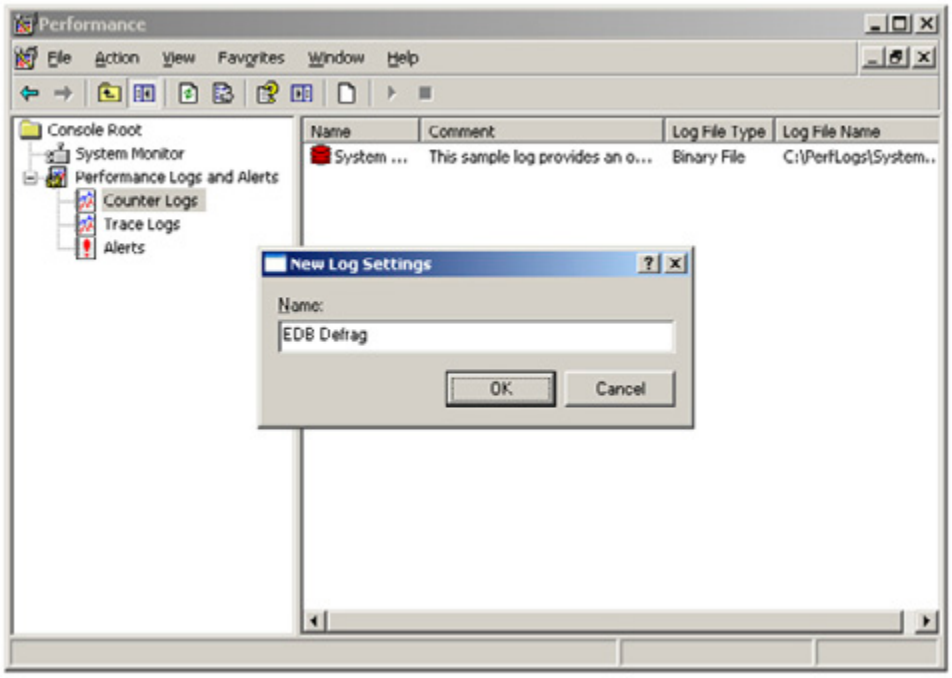

- Click the "Add Counters" button, and then select "MS Exchange Database ==> Instances" from the "Performance object" drop-down.
- Select the "Online Defrag Pages Freed/sec" and the "Online Defrag Pages Read/sec" counters from the list, select the Information Store storage groups, and click the "Add" button.

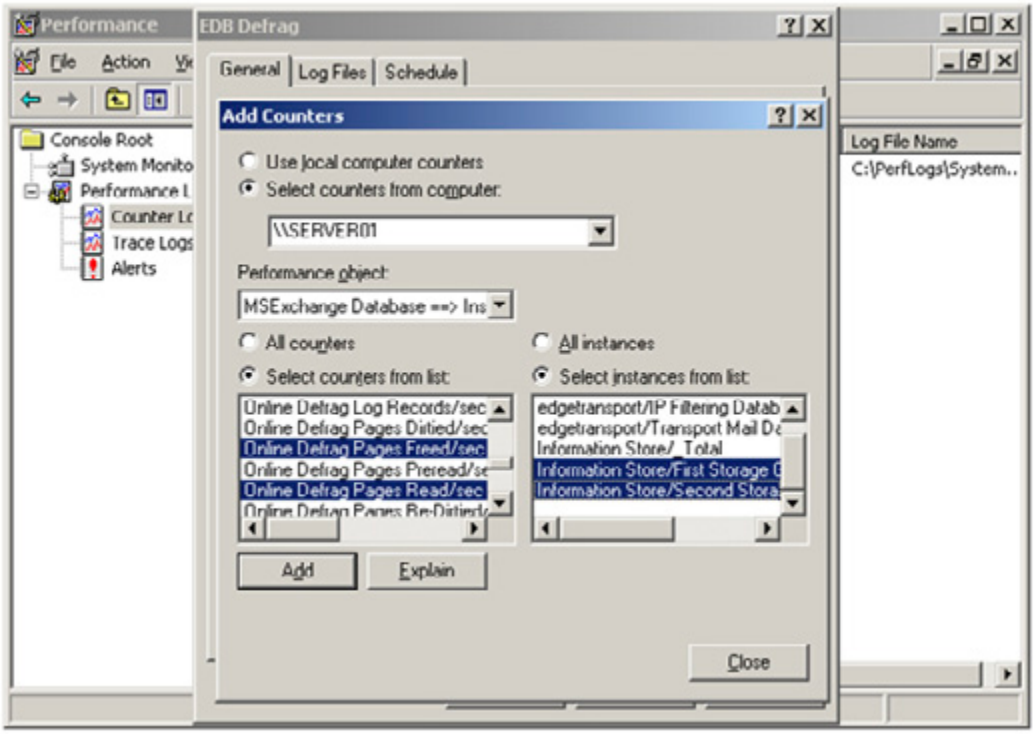

- Click "Close."
- On the "Log Files" tab change the "Log file type" to "Text File (Comma delimited)" and uncheck "End file names with."

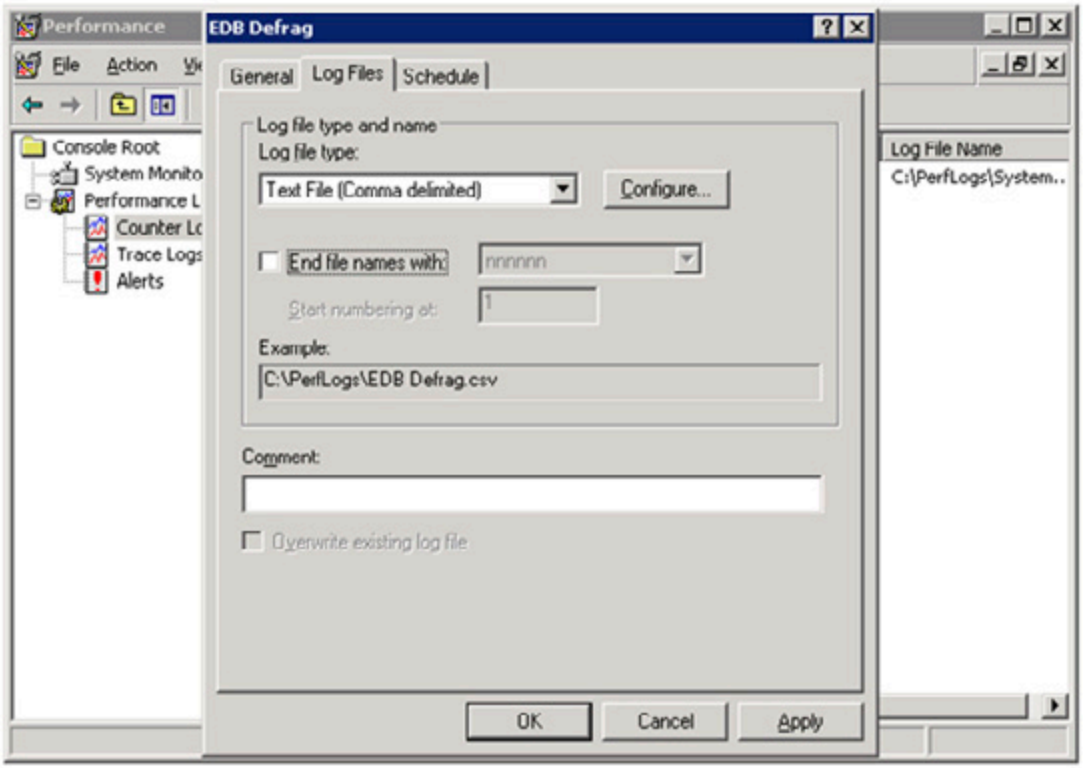

- On the schedule tab, set the start time to a few minutes before the start of the next run of online maintenance and the end time to a few minutes after the end.

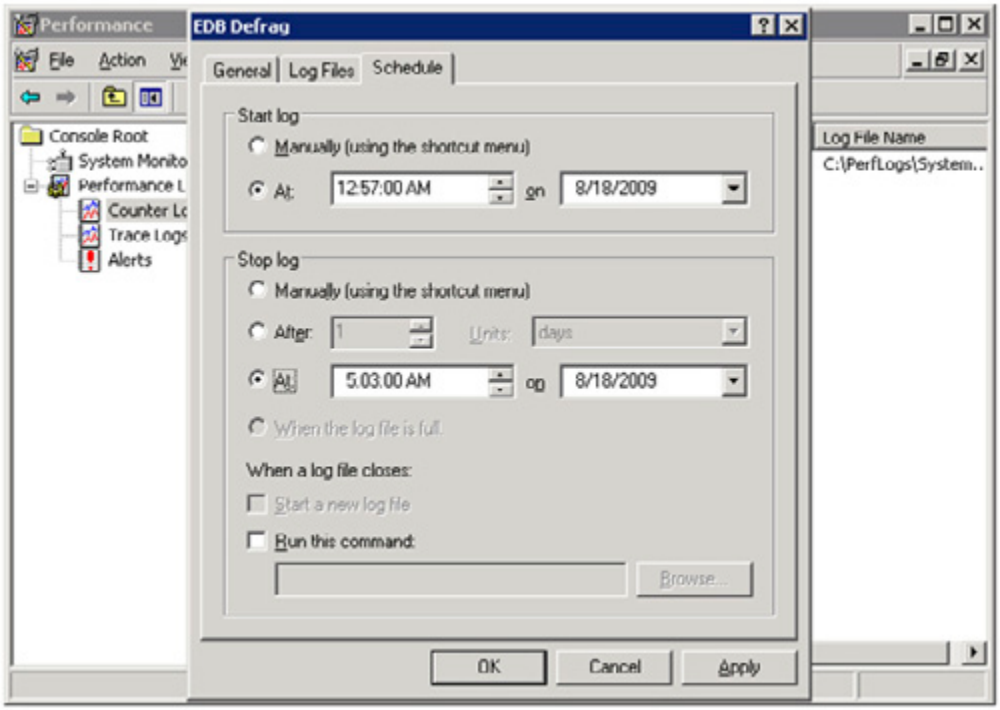

- Click "OK" and then close Performance Monitor.

The next time online maintenance runs the performance counter data will be gathered. This data can then be analyzed by looking at the average ratio of pages read to pages freed.

The resulting CSV file will look similar to this:

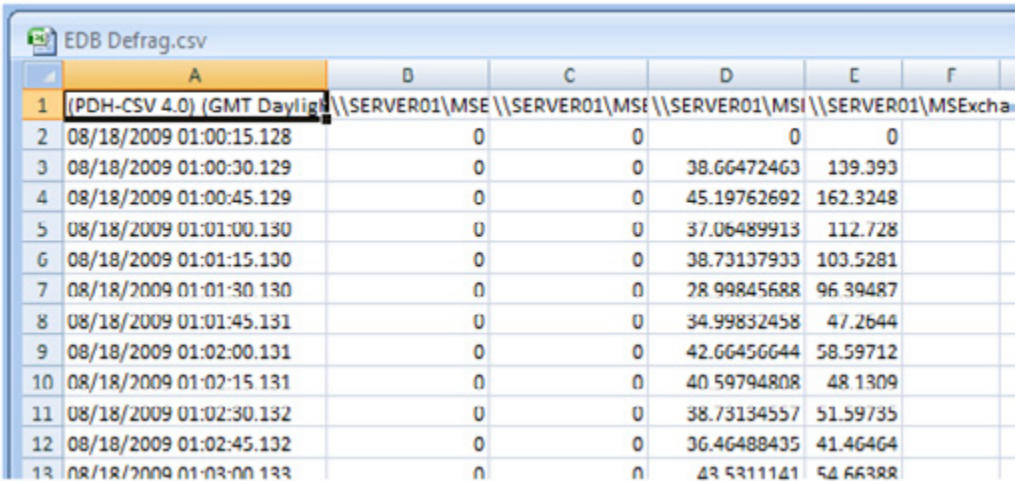

To determine the Pages Read:Pages Freed ratio the data is averaged over the duration of the maintenance window, and the average Pages Read/sec is divided by the average Pages Freed/sec. This gives the number of pages freed/sec for every one page read/sec.

For example, if the average Pages Read/sec is 545, and the average Pages Freed/Sec 6, the ratio is 90:1 and the online defragmentation window is appropriately set.

Setting the Online Maintenance Window

The online maintenance window is configured per database and can be set using either the Exchange Management Console or the Exchange Management Shell.

To configure the maintenance window using the console:

- Launch the Exchange Management Console.
- Expand "Server Configuration" and select the "Mailbox" node.
- Select the server where the database is located then select the database.
- Right-click the database and select "Properties."
- Select a predefined schedule from the "Maintenance Schedule" list or to create a custom schedule, select "Use Custom Schedule," and then click "Customize."
- Click OK to close the properties window.

Alternatively, use the **Set-MailboxDatabase** PowerShell cmdlet to set the maintenance window. This command will configure the maintenance window of the database "Mailbox Database" on server SERVER01 to run every day from 3 am to 4 am local time using the shell:

```
Set-MailboxDatabase -Identity "SERVER01\Mailbox Database" `
-MaintenanceSchedule Sun.03:00-Sun.04:00, Mon.03:00-Mon.04:00, `
Tue.03:00-Tue.04:00, Wed.03:00-Wed.04:00, Thu.03:00-Thu.04:00, `
Fri.03:00-Fri.04:00, Sat.03:00-Sat.04:00
```

With a relatively small amount of monitoring and analysis it is reasonably easy to ensure that online maintenance is running effectively and efficiently, helping to ensure that your Exchange databases stay in good shape.

The techniques I've shown here should help you to check that online maintenance and online defragmentation are running optimally, or are at least conforming to Microsoft's recommendations.

Exchange 2010 High Availability

18 September 2009

by Neil Hobson

In April 2009 Microsoft released a public beta of Exchange 2010, the latest and greatest version of a part of its unified communications family of products. Recently in August 2009, a feature complete Release Candidate version was released for public download. In this article Neil Hobson takes a look at some of the high availability features of Exchange 2010.

For many years Exchange only offered a high availability solution based on the shared storage model, whereby use of Microsoft clustering technologies protected against server-based failure but did nothing to protect against storage failures. Although there were improvements to this form of high availability in Exchange 2007, where it was known as a Single Copy Cluster (SCC), the real changes to high availability in Exchange 2007 came with the introduction of a technology known as continuous replication. With this technology, transaction logs are shipped from one copy of a database to another which allows an organization to deploy an Exchange high availability solution that also dealt with storage failure. This storage failure protection was available on a single server with the use of Local Continuous Replication (LCR) and was also available across servers with the use of Cluster Continuous Replication (CCR). Therefore, with LCR, CCR and SCC, Exchange 2007 administrators had three different high availability methods open to them. It was also possible to cater for site failure with an extension to the continuous replication technology that was known as Standby Continuous Replication. I won't go into detail on these Exchange 2007 solutions here as I've covered them in a previous article called Exchange 2007 High Availability here on Simple-Talk. However, the bottom line is that many organizations have deployed technologies such as CCR in order to provide high availability and technologies such as SCR to provide site resilience.

From my experiences, more organizations have deployed CCR in preference to SCC and it comes as no surprise to learn that SCC has been dropped entirely from Exchange 2010. As you will shortly see, the continuous replication technology lives on in Exchange 2010 but there are many changes in the overall high availability model.

With Exchange 2007 Service Pack 1, a full high availability solution was generally deployed using a total of four servers. Two servers were installed as a single CCR environment, giving high availability for the users" mailboxes. The other two servers were deployed as combined Hub Transport and Client Access Servers, and were configured as a load-balanced pair. The reason for this was simply that if the mailbox server role was clustered, it was not possible to implement the additional Exchange 2007 server roles, such as the Hub Transport and Client Access Server role, on the server running the mailbox role. For the larger enterprises, this wasn't an unreasonable approach but for the smaller organizations a total of four servers sometimes seemed to be overkill for an internal messaging system. To address this specific issue, Microsoft has designed Exchange 2010 such that all server roles can be fully redundant with as few as two servers, providing you have deployed an external load balancer for incoming Client Access Server connections. In other words, it's now possible to combine the mailbox server role with other roles such as the Hub Transport and Client Access Server role. Of course, larger organizations will still be likely to implement dedicated servers running the various server roles but this is something that will definitely help the smaller organizations to reduce costs. Remember, though, the external load balancer requirement for incoming Client Access Server connections.

With this in mind, let's get going and look at some of the high availability features of Exchange 2010. Don't forget that this is a high-level look at the new features; in later articles here on Simple-Talk, we'll be diving much more deeply into these features and how they work. Right now, the idea with this article is to get you to understand the concepts behind these new features and to allow you to do some initial planning on how you might use them in your organization.

Database Availability Groups

Perhaps one of the most important new terms to understand in Exchange 2010 is the *Database Availability Group* (*DAG*). The DAG is essentially a collection of as few as one (although two is the minimum to provide a high availability solution) and up to 16 mailbox servers that allow you to achieve high availability in Exchange 2010.

DAGs use the continuous replication technology that was first introduced in Exchange 2007 and are effectively a combination of Cluster Continuous Replication (CCR) and Standby Continuous Replication (SCR). DAGs make use of some of the components of Windows Failover Clustering to achieve high availability but to reduce overall complexity, these cluster elements are installed automatically when a mailbox server is added to a DAG and managed completely by Exchange. For planning reasons, it's important to understand that the DAG forms the boundary for replication in Exchange 2010. This is a key difference over SCR in Exchange 2007, where it was possible to replicate outside of a CCR environment to a standalone server in a remote data center. However, you should also be aware that DAGs can be split across Active Directory sites if required, meaning that DAGs can therefore offer high availability within a single data center as well as between different data centers.

An important component to a DAG is the file share witness, a term that you will be familiar with if you have implemented a CCR environment in Exchange 2007. Like its name suggests, the file share witness is a file share on a server outside of the DAG. This third server acts as the witness to ensure that quorum is maintained within the cluster. There are some changes to the file share witness operation as we shall discuss later in this section. When creating a DAG, the file share witness share and directory can be specified at the time; if they are not, default witness directory and share names are used. One great improvement over Exchange 2007 is that you do not necessarily need to create the directory and the share in advance as the system will automatically do this for you if necessary.

As with Exchange 2007, the recommendation from Microsoft is to use a Hub Transport server to host the file share witness so that this component will be under the control of the Exchange administrators. However, you are free to host the file share witness on an alternative server as long as that server is in the same Active Directory forest as the DAG, is not on any server actually in the DAG, and also as long as that server is running either the Windows 2003 or Windows 2008 operating system.

A DAG can be created via the *New-DatabaseAvailabilityGroup* cmdlet or via the New Database Availability Group wizard in the Exchange Management Console. The DAG must be created before any mailbox servers are added to it, meaning that effectively an empty container is created which is represented as an object in Active Directory. For example, Figure 1 shows a newly created DAG, called DAG1, in Active Directory as viewed using ADSIEdit.

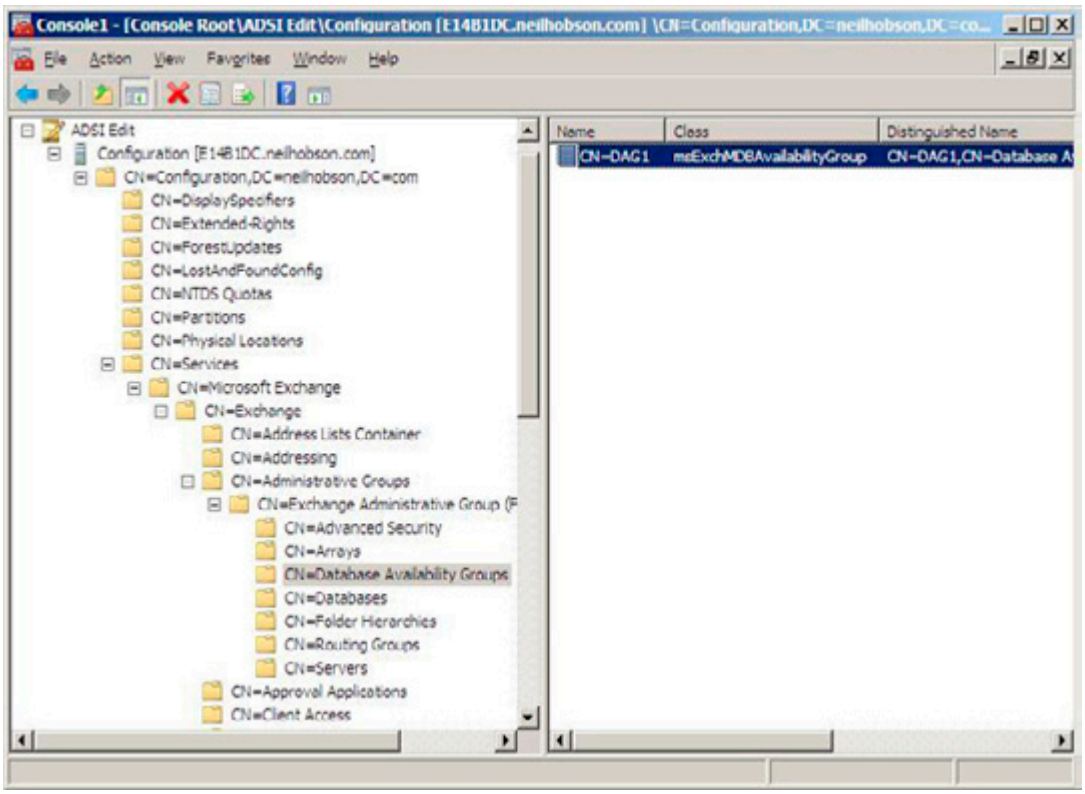

Figure 1: ADSIEdit DAG View.

You can see that a DAG has an object class of *msExchMDBAvailabilityGroup* and that the actual Database Availability Group container location is found under the Exchange 2010 administrative group container. Bringing up the properties of the DAG object in ADSIEdit reveals the important configuration items such as the file share witness share and directory names as you can see in Figure 2.

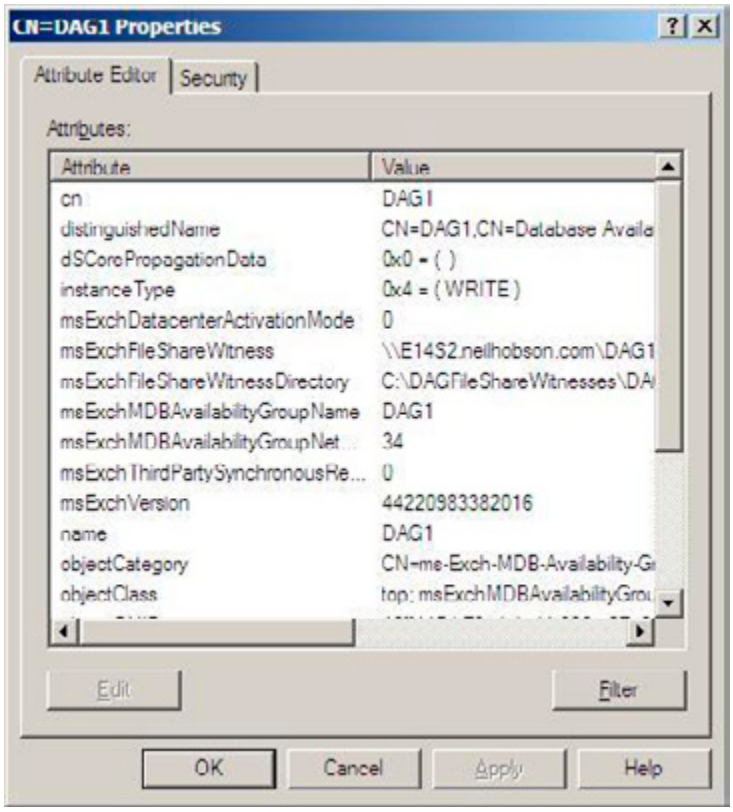

Figure 2: DAG Properties.

Once a DAG has been created, mailbox servers can be added to it as required. This is another simple process that can be achieved by right-clicking the DAG object in the Exchange Management Console and choosing the *Manage Database Availability Group Membership* option from the context menu. The corresponding Exchange Management Shell cmdlet is the *Add-DatabaseAvailabilityGroupServer* cmdlet. For example, to add the mailbox server called E14B1S1 to a DAG called DAG1, you'd run the following cmdlet:

```
Add-DatabaseAvailabilityGroupServer –Identity DAG1 `
  –MailboxServer E14B1S1
```

Since DAGs make use of several Windows Failover Clustering components, it comes as no surprise to see that the Enterprise Edition of Windows Server 2008 is required on mailbox servers that are added to a DAG, so do ensure that you take this into account when planning your Exchange 2010 implementation.

When creating a DAG, there are options around network encryption and compression that can be set. This is possible because Exchange 2010 uses TCP sockets for log shipping whereas Exchange 2007 used the Server Message Block (SMB) protocol. For example, it's possible to specify that the connections that occur using these TCP sockets are encrypted. Equally, it's also possible to decide that these same connections also use network compression.

Mailbox Servers and Databases

Inside each DAG there will normally exist one or more mailbox servers, although it is possible to create an empty DAG as discussed earlier within this article. On each mailbox server in the DAG, there will typically exist multiple mailbox databases. However, one of the key differences between Exchange 2010 mailbox servers and their Exchange 2007 counterparts is that Exchange 2010 mailbox servers can host active and passive copies of different mailbox databases; remember that in Exchange 2007, an entire server in a CCR environment, for example, was considered to be either active or passive. However, in Exchange 2010, the unit of failover is now the database and not the server, which is a fantastic improvement in terms of failover granularity. Consider the diagram below in Figure 3.

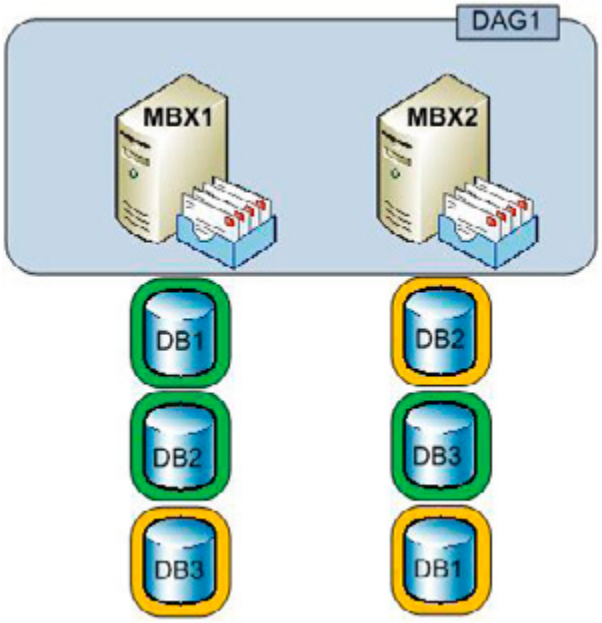

Figure 3: Database Copies.

In Figure 3, you can see that a DAG named DAG1 consists of two mailbox servers called MBX1 and MBX2. There are a total of three active mailbox databases, shown in green, across both servers and each active mailbox database has a passive copy, shown in orange, stored on the alternate server. For example, the active copy of DB1 is hosted on the server called MBX1 whilst the passive copy of DB1 is hosted on the server called MBX2. The passive copies of mailbox databases are kept up-to-date via log shipping methods in the same way that was used in Exchange 2007, such as between the two cluster nodes within a single Exchange 2007 CCR environment. As you might expect, the active copy of the mailbox database is the one which is used by Exchange. Within a DAG, multiple passive copies of a mailbox database can exist but there can only be a single active copy. Furthermore, any single mailbox database server in a DAG can only host 1 copy of any particular mailbox database. Therefore, the maximum possible number of passive copies of a mailbox database is going to be one less than the number of mailbox servers in a DAG, since there will always be one active copy of the mailbox database. For example, if a DAG consisted of the maximum of 16 mailbox servers, then there could be a maximum of 15 passive copies of any single mailbox database. However, every server in a DAG does **not** have to host a copy of every mailbox database that exists in the DAG. You can mix-and-match between servers however you wish.

As mentioned earlier in this section, the unit of failover in Exchange 2010 is now the database. However, if an entire mailbox server fails, all active databases on that server will need to failover to alternative servers within the DAG.

One other vital piece of mailbox database information that you should consider in your planning for Exchange 2010 is the fact that database names are now unique across the forest in which Exchange 2010 is installed. This could be a problem in organizations that have deployed Exchange 2007 with the default database name of *mailbox database*. Therefore, if you are going to be transitioning from Exchange 2007 to Exchange 2010 in the future, take time now to investigate your current database naming standards.

The Active Manager

At this point, you might be wondering how Exchange 2010 determines which of the mailbox databases is considered to be the active copy. To manage this, each mailbox server in a DAG runs a component called the Active Manager. Specifically, one mailbox server in the DAG will be the Primary Active Manager (PAM) whilst the remaining mailbox servers in the DAG will run a Secondary Active Manager (SAM). We will discuss the relationship between clients, Client Access Servers and the active copy of the mailbox database in the next section, as there are some significant changes in this area too. To view the Active Manager information, you can use the *Get-DatabaseAvailabilityGroup* cmdlet and pipe the results into the format-list cmdlet. In other words, you will need to run the following cmdlet:

```
Get-DatabaseAvailabilityGroup | fl
```

Some of the information returned with the *Get-DatabaseAvailabilityGroup* cmdlet references real-time status information about the DAG and one of the parameters returned is the *ControllingActiveManager* parameter. This parameter will show you which server is currently the PAM. It's the job of the PAM to decide which of the passive copies of the mailbox database should become the active copy in the event of an issue with the current active copy. In an environment consisting of many passive copies of mailbox databases, there will naturally be many choices of suitable mailbox databases available to the PAM. As might be expected, the PAM is able to determine the best copy of the mailbox database available for use and it does this via many different checks in order to minimize data loss. Each SAM also has an important part to play, as they inform other services within the Exchange 2010 infrastructure, such as Hub Transport servers, which mailbox databases are currently active.

Client Access Server Changes

In Exchange 2007, Outlook clients connect directly to the mailbox servers whilst other forms of client access, such as OWA, Outlook Anywhere, POP3, IMAP4 and so on, connect via a Client Access Server. The Client Access Server is then responsible for making the connection to the mailbox server role as required. In Exchange 2010, one other fundamental change over previous versions of Exchange is that Outlook clients no longer connect directly to the mailbox servers.

On each Client Access Server, there exists a new service known as the *RPC Client Access Service* that effectively replaces the RPC endpoint found on mailbox servers and also the DSProxy component found in legacy versions of Exchange. The DSProxy component essentially provides the Outlook clients within the organization with an address book service either via a proxy (pre-Outlook 2000) or referral (Outlook 2000 and later) mechanism. A likely high availability design scenario will therefore see a load-balanced array of Client Access Servers deployed, using technologies such as Windows Network Load Balancing or 3rd-party load balancers, which will connect to two or more mailbox servers in a DAG as shown in Figure 4.

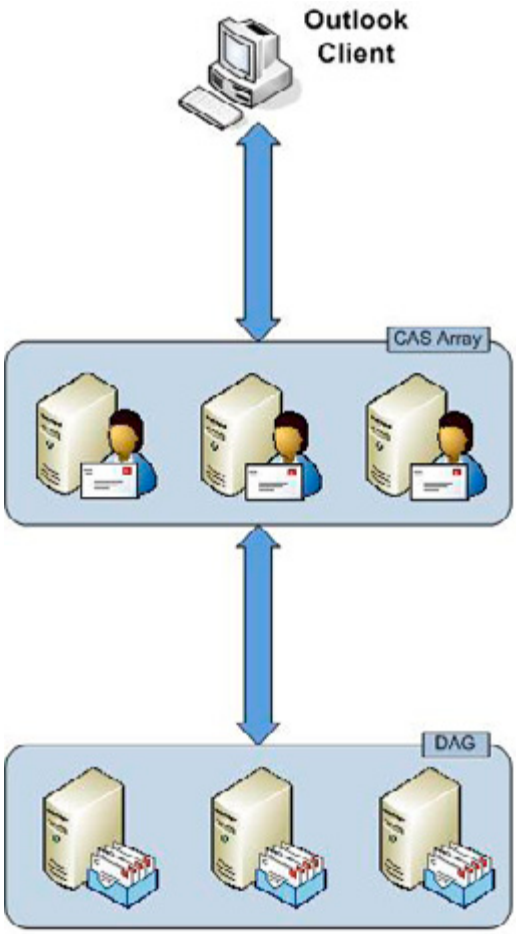

Figure 4: CAS Array.

When an Outlook client connects to an Exchange 2010 Client Access Server, the Client Access Server determines where the active copy of the user's mailbox database is located and makes the connection with the relevant server. If that particular mailbox database becomes unavailable, such as when the administrator wishes to take the database offline or it fails, one of the passive copies will become the new active copy as previously described within this article. It's the Client Access Server that loses the connection to the old active copy of the mailbox database; the actual connection from the client to the Client Access Server is persistent which is obviously good from a user experience point of view. Then, the Client Access Server will fail over to the new active mailbox database in the DAG as directed by the PAM.

Summary

In this article we've taken a high-level look at some of the new Exchange 2010 high availability features and how they come together to provide an overall high availability solution. If you're planning on looking at Exchange 2010, it makes sense to start understanding these new features and how they can benefit your organization. Also, there are other interesting features available in Exchange 2010 that further serve to increase the overall high availability and reliability of the messaging environment, such as shadow redundancy in the Hub Transport server role. In future articles here on Simple-Talk, we'll be covering these areas in much more detail.

Implementing Cluster Replication – Part 1

23 September 2009

by Brien Posey

Imagine that you're researching Continuous Cluster Replication, looking for a simple, direct guide to the basics of getting it set up and running. What do you do if everything you find is either too specialised, or goes into far more detail than you need? If you're Brien Posey, and you can't find that practical, to-the-point article, you roll up your sleeves and write it yourself. To kick off, he tackles the rather tedious task of setting up a Majority Node Set Cluster.

A few weeks ago, someone asked me to help them with an Exchange Server deployment. In order to make the deployment a bit more robust, they wanted to cluster their mailbox servers using Cluster Continuous Replication, or CCR. Although I had set up CCR a few times in the past, it had been a while since I had done one, so I decided to look online to re-familiarize myself with the process.

As I looked around on the Internet, I began to realize that there weren't any articles that met my needs. Some of the articles that I found only covered part of the process. Others covered the entire process but contained a lot more material than what I wanted. Nobody had published a simple article, written in plain English, which described the procedure of setting up CCR from start to finish. That being the case, I decided to write one myself.

Planning the Cluster

As you may already know, CCR is not a feature that you can enable on a whim. That's because CCR makes use of a majority node set cluster. You have to create the cluster at the operating system level. Only then can you install Exchange and create a clustered mailbox server. Since that is the case, I want to start out by showing you how to plan for and set up the cluster. In Part 2, I will show you how to install Exchange onto the cluster that you have created, and I will also show you how to perform a manual failover on the cluster.

The Server Hardware and Software

The first step in the planning process is to make sure that you have the required hardware. Generally speaking, the requirements for setting up CCR are roughly the same as for creating any other mailbox server. The biggest difference is that you will need two of everything. CCR does not make use of shared storage like a single copy cluster does, so you will need two servers, and each of the servers will have to have sufficient disk resources to accommodate the mailbox database and the transaction logs. As is the case with a non-clustered mailbox server, you should place the database and the transaction logs on separate disks in accordance with Microsoft's best practices for Exchange.

Your two servers don't have to be completely identical to each other, but they should at least have similar capabilities. Remember that at any time either one of the servers could be acting as the active cluster node, so it is important to make sure that both servers have hardware that is sufficient to host the mailbox server role in an efficient manner.

Although not an absolute requirement, I recommend installing two NICs in each of the servers. It is also worth noting that Microsoft won't support CCR unless you use two NICs in each node. One of the NICs will be used for communications with the rest of the network, and the other will be used for communications between cluster nodes. The installation steps outlined in this article assume that each server has two NICs.

Finally, you must ensure that you have the necessary software licenses. You will need two copies of Windows Server Enterprise Edition and two copies of Exchange 2007 Enterprise Edition (clustering is not supported in the standard editions), plus any necessary Client Access Licenses. For the purposes of this article, I will be using Windows Server 2008 and Exchange Server 2007 with SP1.

Other Cluster Resources

The next step in planning the cluster is to set aside the necessary names and IP addresses. Believe it or not, you are going to have to use five different names and seven different IP addresses.

The first step in setting up the cluster is going to be to install Windows on to both of the cluster nodes. At this point in the process your servers are not cluster nodes, but rather just a couple of Windows servers. Like any other Windows Servers, you are going to have to assign a name to both of the servers. You are also going to have to assign each of the servers an IP address. That accounts for two of the four names and two of the six IP addresses.

As you will recall, we have two NICs in each of the servers. You must assign an IP address to both of the secondary NICs. The IP addresses that you used for the primary NICs should fall into the same subnet as any other machines on the network segment. The addresses assigned to the secondary NICs should belong to a unique subnet, although it is permissible to connect a crossover cable directly between the two cluster nodes

So far we have used two of our four names, and four of our six IP addresses. The next name and IP address are assigned to the cluster. This name and IP address are used to communicate with the cluster as a whole, rather than with an individual cluster node.

Finally, you will need to set aside a name and an IP address to be assigned at the Exchange Server level to the clustered mailbox server. The IP address for the cluster and the IP address for the clustered mailbox server should fall into the same subnet as the other computers on the network segment. You will only be using the alternate subnet for cluster level communications between the two nodes over the secondary network adapters.

The table below should help to give you a clearer picture of how the various names and IP addresses will be used. You don't have to use the same names and addresses as I am. They are just samples:

	Name	Primary NIC	Secondary NIC
Cluster Node 1	Node1	192.168.0.1	10.1.10.11
Cluster Node 2	Node2	192.168.0.2	10.1.10.12
MNS Cluster	WinCluster	192.168.0.3	
Clustered Mailbox Server	Exch1	192.168.0.4	

Another thing that you are going to need before you begin creating the cluster is a cluster service account. You should create a dedicated account in your Active Directory domain, and set its password so that it never expires.

One last thing that I want to mention before I show you how to create the cluster is that the Clustered Mailbox Server role cannot be combined with any other Exchange Server roles. Since every Exchange 2007 organization requires at least one hub transport server and client access server, you will need at least one additional Exchange 2007 server in your organization.

Creating the Cluster

Now that I have explained what resources you will need I want to go ahead and show you how to create the cluster. This section assumes that you have already installed Windows Server 2008 onto both cluster nodes, assigned the appropriate names and IP addresses to the two nodes, and joined the nodes to your domain.

Configuring the First Cluster Node

Your cluster is going to consist of two separate nodes, and the setup procedure is going to be different for the two nodes. That being the case, I am going to refer to the first cluster Node as Node 1, and the second cluster node as Node 2. When the configuration process is complete, Node 1 will initially act as the active cluster node.

With that said, let's go ahead and get started by configuring Node 1. Begin the process by logging into the server with administrative credentials, and opening a Command Prompt window. Now, enter the following command:

```
Cluster /Cluster:<your cluster name> /Create /Wizard
```

For example, if you chose to call your cluster WinCluster, then you would enter:

```
Cluster: /Cluster:WinCluster /Create /Wizard
```

This command tells Windows that you want to create a new cluster named WinCluster, and that you want to use the wizard to configure the cluster.

Windows will now launch the New Server Cluster Wizard. Click Next to bypass the wizard's Welcome screen. Now, select your domain from the Domain drop down list. Verify that the cluster name that is being displayed matches the one that you typed when you entered the Cluster command, and then click Next.

Windows will now perform a quick check to make sure that the server is ready to be clustered. This process typically generates some warnings, but those warnings aren't usually a big deal so long as no errors are displayed. The warnings are often related to 1394 firewire ports being used as network interfaces, and other minor issues.

Click Next to clear any warning messages, and you will be prompted to enter the IP address that you have reserved for the MNS cluster. Enter this IP address into the space provided, and then click Next.

You will now be prompted to enter your service account credentials. Enter the required username and password, and click Next.

At this point, the wizard will display a summary screen. With virtually every other wizard that Microsoft makes, you can just take a second to verify the accuracy of the information, and then click Next. In this case though, you need to click the Quorum button instead. After doing so, you must set the quorum type to Majority Node Set. After doing so, click OK, followed by Next. When the wizard completes, click Next, followed by Finish. You have now created the first cluster node!

Adding the Second Node to the Cluster

Now that we have created our cluster, we need to add Node 2 to it. To do so, log into Node 2 as an administrator, and open a Command Prompt window. When the window opens, enter the following command:

```
Cluster /Cluster:<your cluster name> /Add /Wizard
```

For example, if you called your cluster WinCluster, you would enter this command:

```
Cluster /Cluster:WinCluster /Add /Wizard
```

Notice that this time we are using the /Add switch instead of the /Create switch because our cluster already exists.

Windows should now launch the Add Nodes Wizard. Click Next to clear the wizard's welcome screen. You must now select your domain from the Domain drop down list. While you are at it, take a moment to make sure that the cluster name that is being displayed matches what you typed.

Click Next, and you will be prompted to enter the name of the server that you want to add to the cluster. Enter the server name and click the Add button.

Click Next, and the wizard will perform a quick check to make sure that the server is ready to be added to the cluster. Once again, it is normal to get some warning messages. As long as you don't receive any error messages, you can just click Next.

At this point, you will be prompted to enter the credentials for the cluster's service account. After doing so, click Next.

You should now see the now familiar configuration summary screen. This time, you don't have to worry about clicking a Quorum button. Just click Next, and Windows will add the node to the cluster. When the process completes, click Next, followed by Finish.

Some Additional Configuration Tasks

Now that we have created our Majority Node Set Cluster, we need to tell Windows which NICs are going to be used for which purpose. To do so, select the Cluster Administrator console from Node 1's Administrative Tools menu. When the Cluster Administrator starts, take a moment to make sure that both of your cluster nodes are listed in the Cluster Administrator's console tree.

Now, navigate through the console tree to <your cluster name> | Cluster Configuration | Networks | Local Area Connection. This container should display IP addresses for both cluster nodes. Take a moment to verify that the addresses that are listed are the ones that fall into the same subnet as the other servers on the network segment. Now, right click on the Local Area Connection container, and choose the Properties command from the resulting shortcut menu. When the properties sheet opens, make sure that the Enable this Network for Cluster Use check box is selected. You must also select the Client Access Only (Public Network) option. When you have finished, click OK.

Now, we have to check the other network connection. To do so, navigate through the console tree to <your cluster name> | Cluster Configuration | Networks | Local Area Connection 2. Make sure that when you select the Local Area Connection 2 container, that the details pane displays both cluster nodes, and that the IP addresses that are listed are associated with the private subnet.

At this point, you must right click on the Local Area Connection 2 container, and select the Properties command from the resulting shortcut menu. When Windows opens the properties sheet for the connection, make sure that the Enable This Network for Cluster Use check box is selected. You must also select the Internal Cluster Communications Only (Private Network) option. When you are done, click OK and close the Cluster Administrator.

Creating a Majority Node Set File Share Witness

The problem with a Majority Node Set cluster is that the active node must be able to communicate with the majority of the nodes in the cluster, but there is no way to have a clear majority in a two node cluster. Windows can't allow a single node to count as the majority, because otherwise a failure of the communications link between the two cluster nodes could result in a split brains failure. This is a condition in which both nodes are functional, and each node believes that the other node has failed, and therefore tries to become the active node.
In order to prevent this from happening, we must create a Majority Node Set File Share Witness. The basic idea behind this is that we will create a special file share on our hub transport server. In the event of a failure, the share that we create will be counted as a node (even though it isn't really a node) in determining which cluster node has the majority of the node set.

To create the Majority Node Set File Share Witness, go to your hub transport server, open a Command Prompt window, and enter the following commands:

```
C:
CD\
MNS_FSW_CCR
Net Share MNS_FSW_CCR=C:\MNS_FSW_CCR /Grant:<your service account name>,Full
CACLS C:\MNS_FSW_CCR /G Builtin\Administrators:F <your service account>:F
```

When Windows asks you if you are sure, press Y.

What we have done is created a folder on our hub transport server named C:\MNS_FSW_CCR. We then created a share named MNS_FSW_CCR, and gave our service account full access to the share. Finally, we gave the built in Administrator account and the service account full access to the folder at the NTFS level.

Now, go to Node 1 and open a Command Prompt window. You must now enter the following commands:

```
Cluster <your cluster name> Res "Majority Node Set" /Priv MNSFileShare=\\<your hub transport
server's name>\MNS_FSW_CCR
Cluster <your cluster name> group "Cluster Group" /move
Cluster <your cluster name> group "Cluster Group" /move
Cluster <your cluster name> Res "Majority Node Set" /Priv
```

The first command in this sequence tells Windows to use the share that we have created as the Majority Node Set File Share Witness. When you enter this command, you will receive an error message telling you that the properties were stored, but that your changes won't take effect until the next time that the cluster is brought online.

The easiest way to get around this problem is to move the cluster group from Node 1 to Node 2 and then back to Node 1. That's what the second and third commands in the sequence above accomplish for us.

The last command in the sequence above simply causes Windows to display the private properties for the Majority Node Set. The first line of text within the list of properties should make reference to the share that we have created for our Majority Node Set File Share Witness. This confirms that the cluster is using our Majority Node Set File Share Witness.

Conclusion

As you can see, creating a Majority Node Set Cluster can be a bit tedious. In Part 2 of this series, we will wrap things up by installing Exchange Server onto our cluster and then working through the failover procedure.

The Active Directory Recycle Bin in Windows Server 2008 R2

23 September 2009

by Jonathan Medd

> ***It has always been a curse as well as a blessing that Active Directory has allowed the rapid removal of whole branches. Until now, administrators have looked in vain for an "undo" function after having accidentally deleted an entire division of their company. At last, with Windows Server 2008 R2, comes a way to rollback changes, as long as you are handy with Powershell. Jonathan Medd explains.***

Since Active Directory was included as part of Window Server 2000, administrators have often asked for a simple way to roll back mistakes, whether that is the incorrect deletion of the wrong user account to the accidental removal of thousands of objects by deleting an OU. Before the release of Windows Server 2008 R2 there were a number of ways using built-in or third-party methods to restore Active Directory objects, but typically they were not as quick or complete as say retrieving a deleted email or file.

Microsoft has included with their release of Windows Server 2008 R2 the facility, under the correct conditions, to enable a Recycle Bin for Active Directory and allow simple restoration of objects which have been erroneously removed. In this article we will briefly cover some of the options prior to 2008 R2 and then examine how to enable the new Recycle Bin and restore objects from it.

Pre-Windows Server 2008 R2

The 2008 R2 Recycle Bin for Active Directory is a great motivating point for upgrading your forest and domain(s) to the latest version, but this is not always a quick process in many enterprises so it is worth knowing what options are available prior to this version. Like many things it's a lot better to examine and plan for possible resolutions before a significant mistake happens that you need to deal with. Retrieving Active Directory objects typically falls into two available categories, authoritative restore from a backup or tombstone reanimation.

Authoritative Restore

The Microsoft KB article 840001(http://support.microsoft.com/kb/840001) details how to perform the restoration of a user account using a system state backup of a domain controller. Typically, you would use a global catalog so that you can also restore all group membership information.

Tombstone Reanimation

The above article also details how to recover an account when you don't have a system state backup by using tombstone reanimation which was introduced with Windows Server 2003 – you can retrieve objects from the Deleted Objects container where they are kept after deletion until their tombstone period expires. Obviously regular system state backups of Active Directory are critical for your full disaster recovery procedures, but taking advantage of tombstone reanimation means you can get objects back quicker than having to go through the full authoritative restore process.

You could use the procedure in the article which utilises the ldp.exe tool, but there are other methods around which you may find simpler.

- The article itself links to a Sysinternals tool, **ADRestore** (http://technet.microsoft.com/en-us/sysinternals/bb963906.aspx), which is a command line tool for reanimating objects.

- The free **ADRestore.Net**, a GUI tool made by Microsoft PFE Guy Teverovsky. http://blogs.microsoft.co.il/blogs/guyt/archive/2007/12/15/adrestore-net-rewrite.aspx.

- Quest produces a freeware product **Object Restore for Active Directory**, an easy to use GUI tool. http://www.quest.com/object-restore-for-active-directory/ (Note: there is a commercial version with more features, **Recovery Manager for Active Directory**.)

- Quest also produces a cmdlet library for managing Active Directory with Windows PowerShell (http://www.quest.com/powershell/activeroles-server.aspx). As of version 1.2 a number of the cmdlets had a Tombstone parameter added to them so that a search of objects would also include items which have been tombstoned. These results could then be piped through to the new cmdlet Restore-QADDeletedObject to undelete the object represented by the tombstone. For instance the command **Get-QADUser –Tombstone -LastChangedOn ((Get-Date).adddays(-1)) | Restore-QADDeletedObject** would restore all user accounts deleted yesterday.

The drawback with tombstone reanimation is that because most of the object's attributes are removed at the time of the object's deletion, a restored object using this method requires many properties of the account, such as address fields and group membership, to be manually repopulated. Whilst this is obviously preferable to re-creating an account from scratch it does not make for a quick overall process. However, you will at least get back the objectGUID and objectSid attributes which means there would be no need to re-configure a user's workstation profile.

The original release of Windows Server 2008 introduced snapshot backups for Active Directory. You can take point-in-time snapshots of your Active Directory with the **NTDSUTIL** command line utility which utilizes Volume Shadow Copy to provide a snapshot. It is then possible to mount this snapshot using different ports on the same domain controller as the live Active Directory database and use standard tools to compare the two. This could really make the tombstone reanimation a lot simpler because after restoring the object you could view two versions of Active Directory Users and Computers side by side and view the properties of the restored object from a previous time, so making it simpler to repopulate properties.

The Directory Service Comparison Tool (http://lindstrom.nullsession.com/?page_id=11) takes advantage of these snapshots and makes the repopulation process more streamlined.

For those with Microsoft Exchange messaging environments, once you have the Active Directory account back, you can use the **Reconnect Mailbox** feature within Exchange to tie the restored account back up with the mailbox. This is of course providing you have a similar tombstone retention period for mailboxes that you do for AD accounts.

Active Directory Recycle Bin

The real reason you decided to read this article though was not so that we could spend time going over all the possible options for how you can piece together restored AD objects, but rather to find out how the Recycle Bin is going to make your life as an Active Directory administrator easier without necessarily the need for these different tools. The key differences from previous versions of Windows Server are that by default you get all of the attributes back and the tools to use are PowerShell cmdlets, which are quickly becoming a more

essential part of every Windows administrator's standard toolkit.

Firstly though the Active Directory Recycle Bin is not enabled by default and has certain domain and forest wide requirements before it can be enabled.

- Firstly, all domain controllers within the Active Directory forest must be running Windows Server 2008 R2.
- Secondly, the functional level of the Active Directory forest must be Windows Server 2008 R2.

Naturally organizations are typically cautious when upgrading Active Directory and these types of infrastructure projects don't tend to happen quickly, but the Recycle Bin could be one of the features which gives you more weight behind a decision. You should also be aware though that enabling the Recycle Bin is a onetime only move, there's no easy way to disable it again, so careful consideration of this decision must be taken.

It's worth noting that if you are making a fresh forest install of Windows Server 2008 R2 the Active Directory schema will already include all of the necessary attributes for the Recycle Bin to function. If however you are upgrading your domain controllers from previous versions of Windows Server then you will need to run the well known procedure of `adprep /forestprep` and `adprep /domainprep` (for each domain) and possibly `adprep /domainprep /gpprep` (for Group Policy preparation)

before you can introduce Windows Server 2008 R2 domain controllers into the environment.

So let's go ahead and run through all the steps we need to get the Recycle Bin enabled. Firstly, ensure that all of your domain controllers are running Windows Server 2008 R2 and then we need to use PowerShell; the great news with Windows Server 2008 R2 is that version 2 of PowerShell is installed by default and is placed directly on your taskbar.

After you have installed Active Directory Domain Services the Active Directory specific cmdlets are available to use via a module; modules essentially are the evolution of snapins from version 1 of PowerShell. To access these cmdlets you can either open the Active Directory specific version of the PowerShell console from the Administrative Programs menu, or the method I would prefer, use the **Import-Module** cmdlet.

Tip

You could add the following expression to your PowerShell profile so that the cmdlets are available every time you open PowerShell

```
PS> Import-Module activedirectory
```

Once complete all of the Active Directory cmdlets will be at your fingertips. As previously discussed we now need to get the functional level of the forest up to the level of Windows Server 2008 R2. The most common way to do this previously was through Active Directory Domains and Trusts.

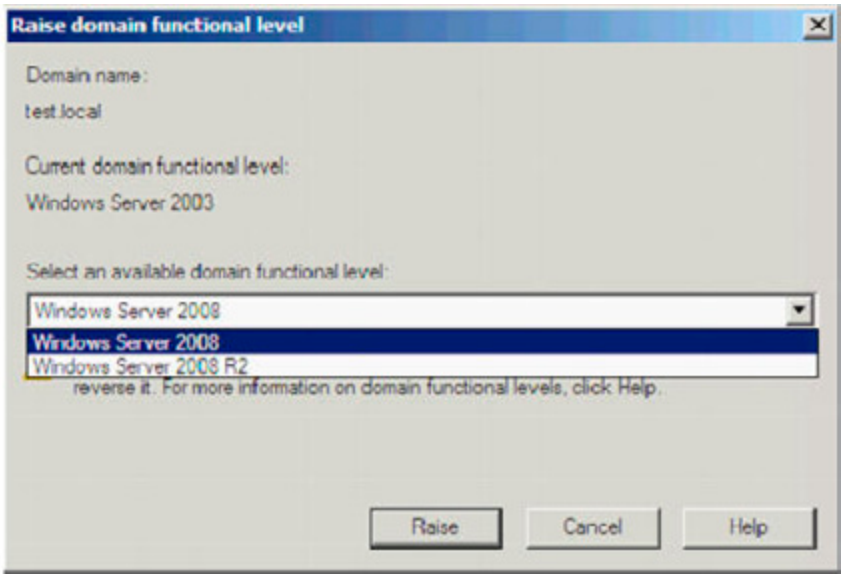

Now though we can do this through PowerShell. The **Get-ADForest** cmdlet will return information about your forest and the **Set-ADForestMode** cmdlet will enable you to raise the current functional level – since it is such a significant change to your environment you will be prompted to confirm that you wish to go ahead.

```
PS> Get-ADForest | Set-ADForestMode -ForestMode Windows2008R2Forest
```

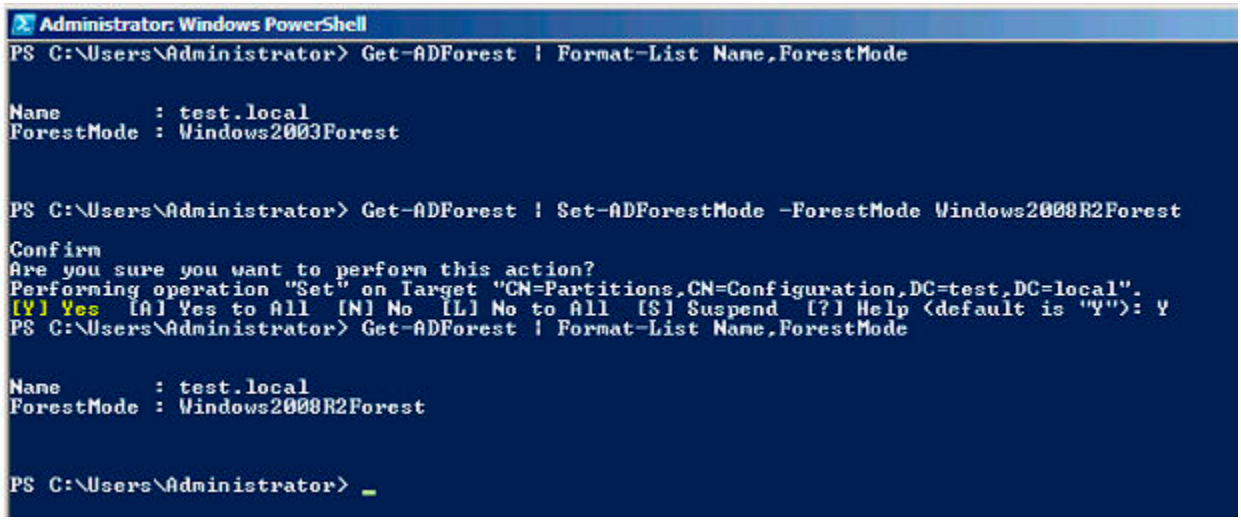

Now that our forest is at the correct functional level we can enable the Recycle Bin, to do so we use the **Enable-ADOptionalFeature** cmdlet. This must be either run on the DC with the Domain Naming Master FSMO role or directed at that server with the **–server** parameter. Again you will be prompted to confirm your command since the action is irreversible.

```
PS> Enable-ADOptionalFeature 'Recycle Bin Feature' -Scope ForestOrConfigurationSet -target 'test.
local'
```

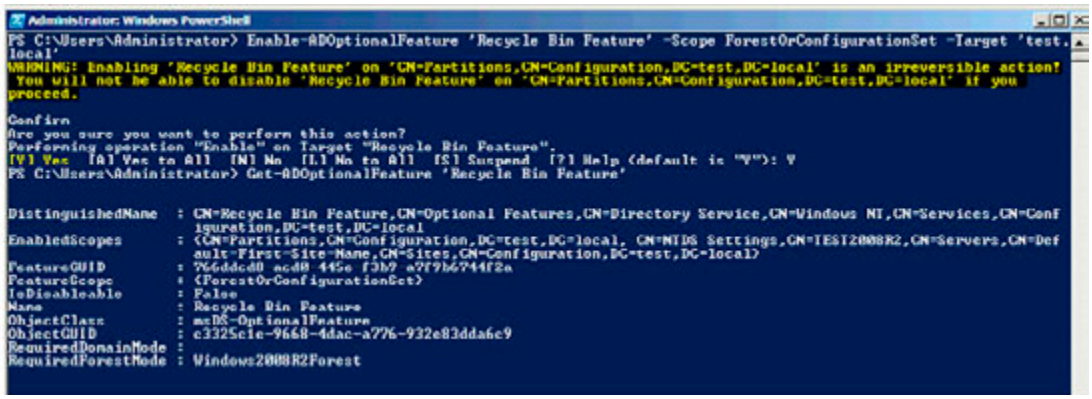

Now that we have the Recycle Bin enabled it's time to go check out how we recover some deleted objects. In this environment we have a very simple AD structure with a couple of test accounts to illustrate the example.

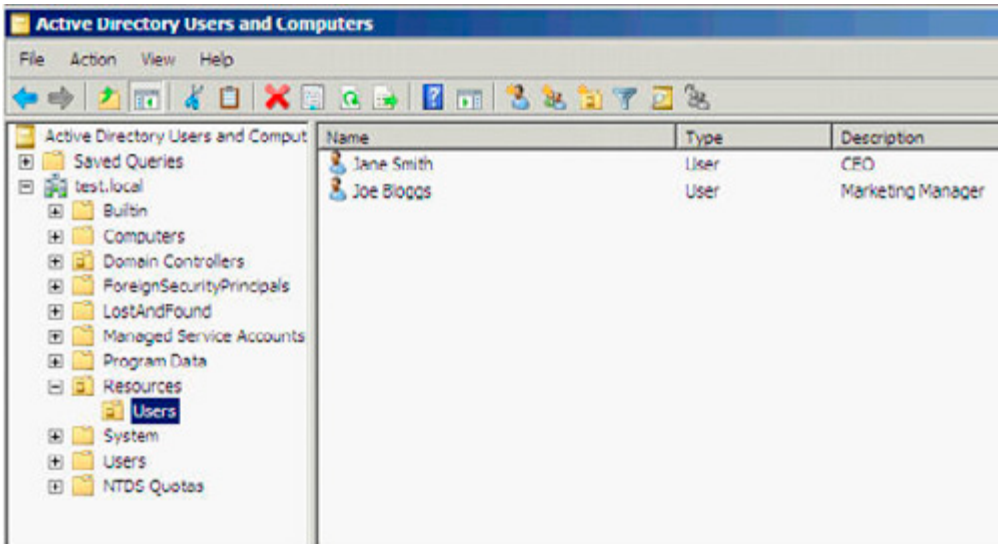

Let's take the situation where an administrator accidently deletes the **Users** OU. One of the most common reasons this can happen is because it is actually possible to delete OU's from the Group Policy Management tool, not just Active Directory Users and Computers – so an administrator might think they are removing a GPO and in a bad moment delete the wrong item and remove a whole OU. The administrator is prompted for what they are about to do, but I have seen it happen more than once!

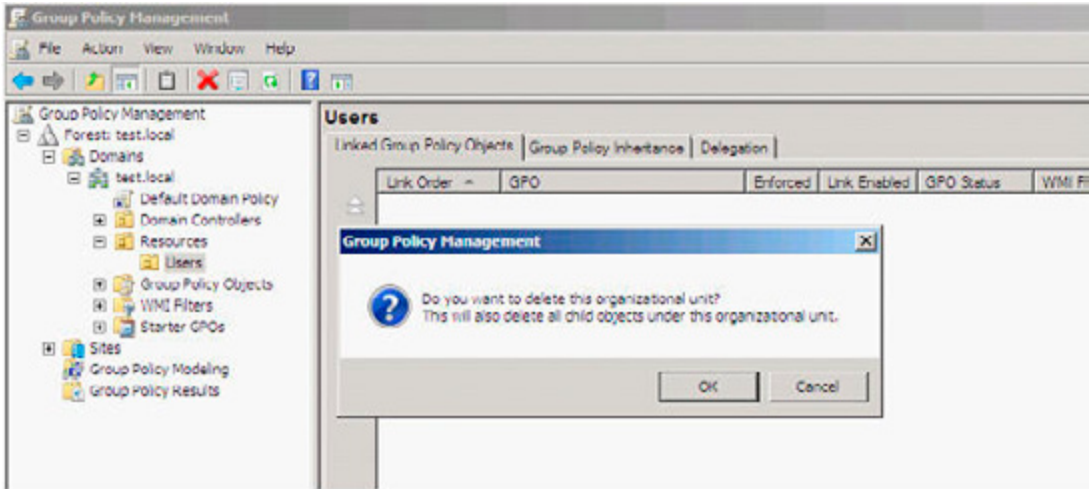

The initial release of Windows 2008 Server actually included a new checkbox "Protect object from accidental deletion'. In the example of the OU below any attempt to delete the OU will be met with an **Access is denied** response and the administrator will actually have to remove the tick from that checkbox before the OU can be deleted.

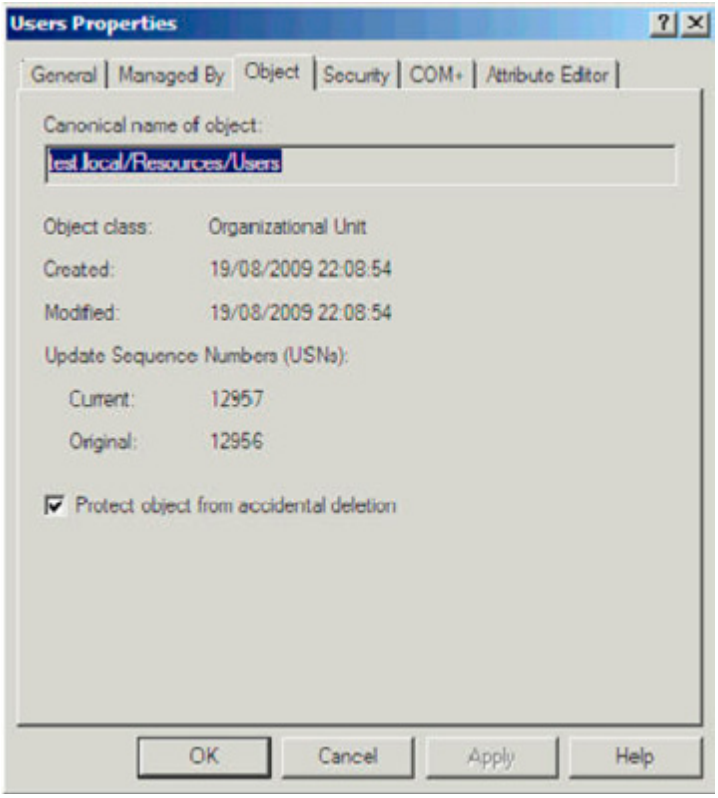

However, what you would naturally expect to happen as a consequence of the **Protect object from accidental deletion** would be any user or computer account created in that protected OU would also be supported by the same mechanism. Unfortunately by default they are not, so as a good practise you would either need to build that into your account creation process or programmatically check and set that checkbox on all accounts in the OU on a regular basis.

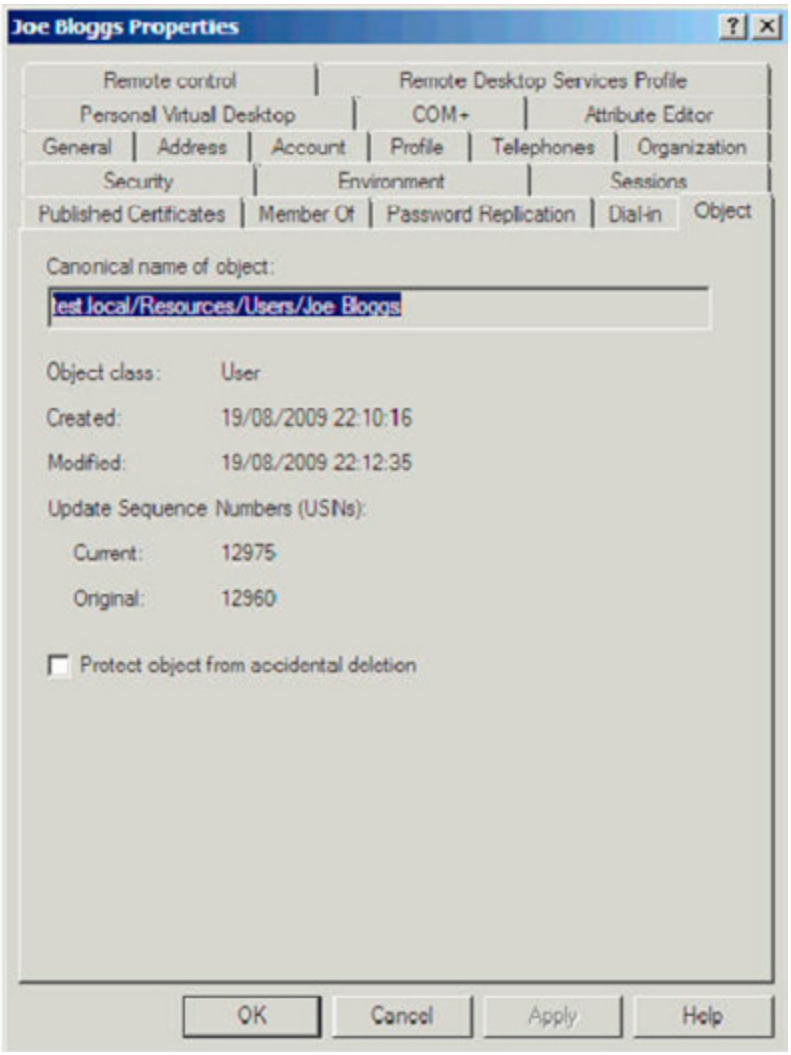

Consequently, in the above example if we accept the warning to delete the OU we are greeted with an **Access is denied** message since the OU has protection set.

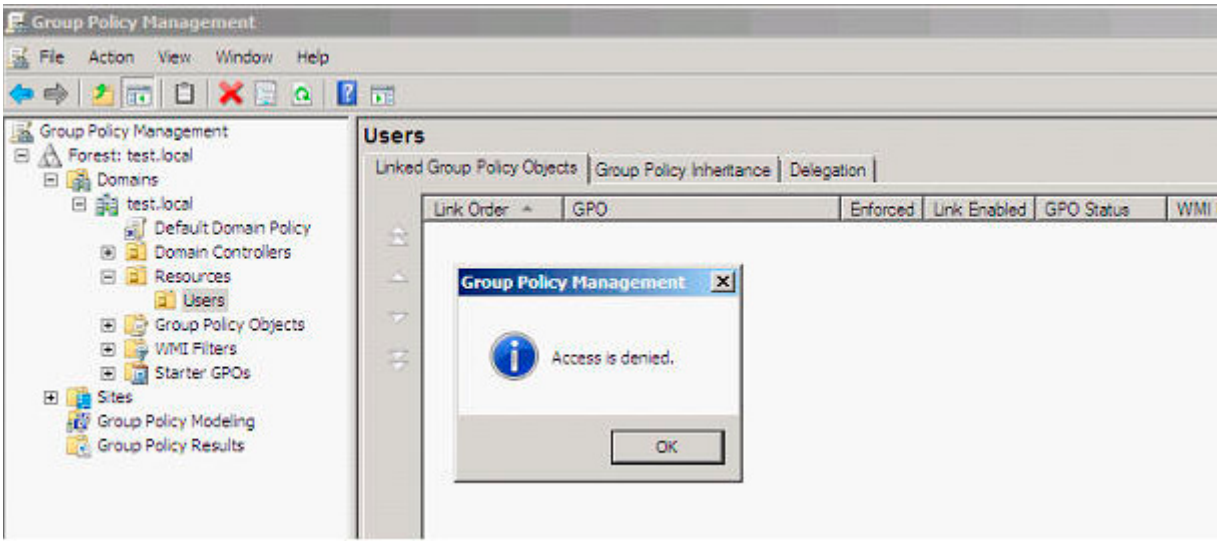

So we were saved from deleting the OU, but all of the unprotected child objects were deleted.

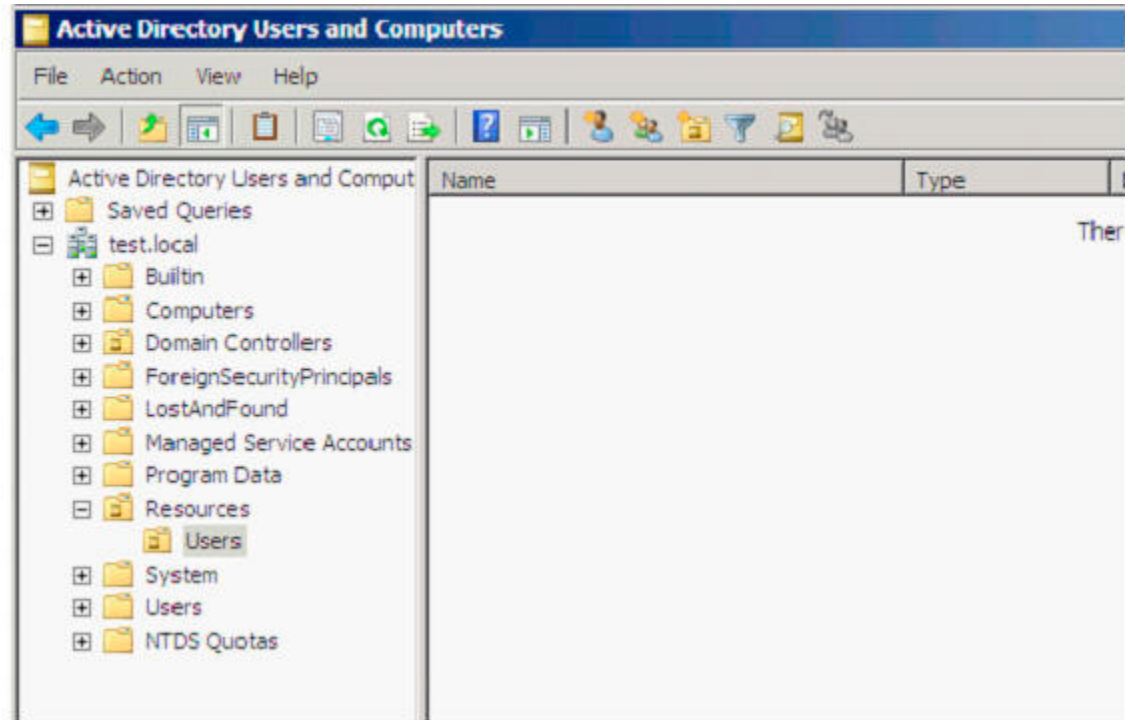

(For the purposes of this article I now remove the **Users** OU by first clearing the checkbox for protecting the object from accidental deletion.)

We can browse the current contents of the Active Directory Recycle Bin using the **Get-ADObject** cmdlet, directing it at the **Deleted Objects** container and using the **–includeDeletedObjects** parameter.

```
PS> Get-ADObject -SearchBase "CN=Deleted Objects,DC=test,DC=local" -ldapFilter "(objectClass=*)"
-includeDeletedObjects | Format-List Name,ObjectClass,ObjectGuid
```

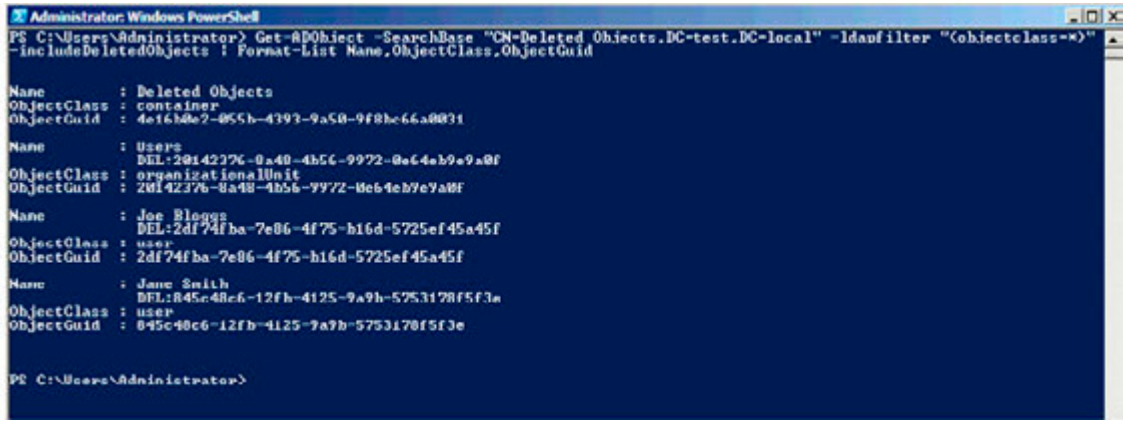

We can see from the resultant output that we have both the **Users** OU in there and the two user accounts. So let's try restoring one of the user accounts back, to do so we need the **Restore-ADObject** cmdlet and supply the ObjectGuid property of the user account.

```
PS> Restore-ADObject -identity 2df74fba-7e86-4f75-b16d-5725ef45a45f
```

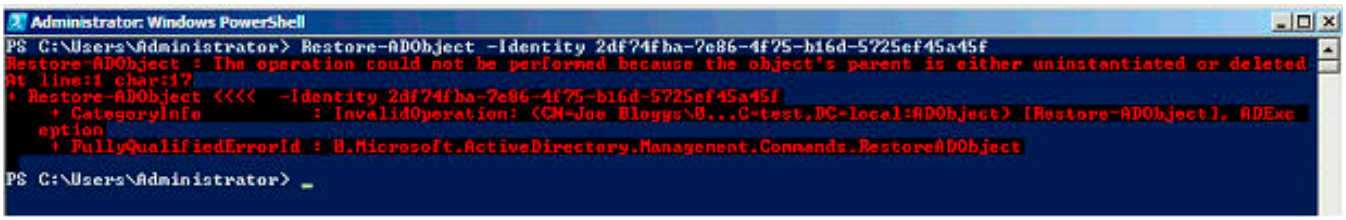

Oh dear, it failed to restore, but PowerShell tells us that it failed because the object's parent no longer exists either, i.e. we need to first restore the **Users** OU. (**Note**: an alternative would be to use the -targetpath parameter and re-direct the restore to a different OU.)

To restore the **Users** OU we can use the same cmdlet (**Restore-ADObject**) as to restore users, just supply the ObjectGuid of the OU.

```
PS> Restore-ADObject -identity 20142376-8a48-4b56-9972-0e64eb9e9a0f
```

The **Users** OU returns.

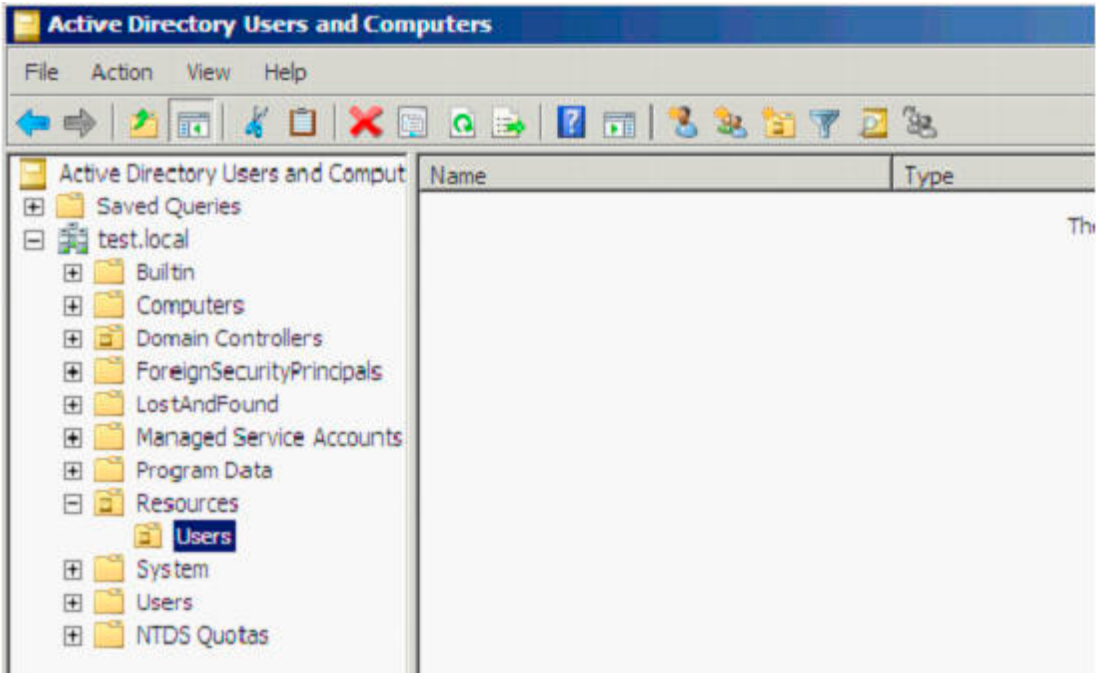

Now we just need to get those user accounts back. Rather than have to type out the ObjectGuid for each account we wish to restore we can instead create a search which will match all of the accounts we wish to restore and then use the PowerShell pipeline to send those results to the **Restore-ADObject** cmdlet.

```
PS> Get-ADObject -ldapFilter "(lastKnownParent=OU=Users,OU=Resources,DC=test,DC=local)"
-includeDeletedObjects | Restore-ADObject
```

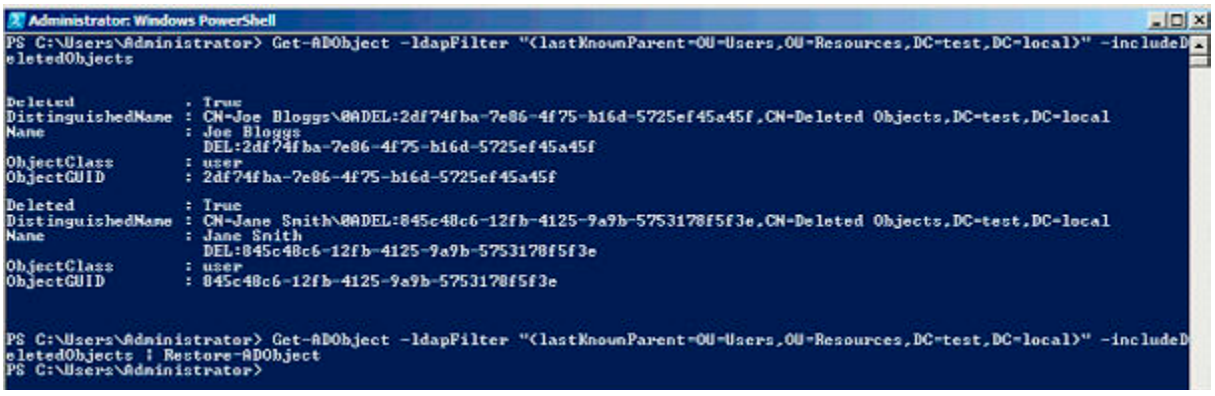

The user accounts are back in the **Users** OU.

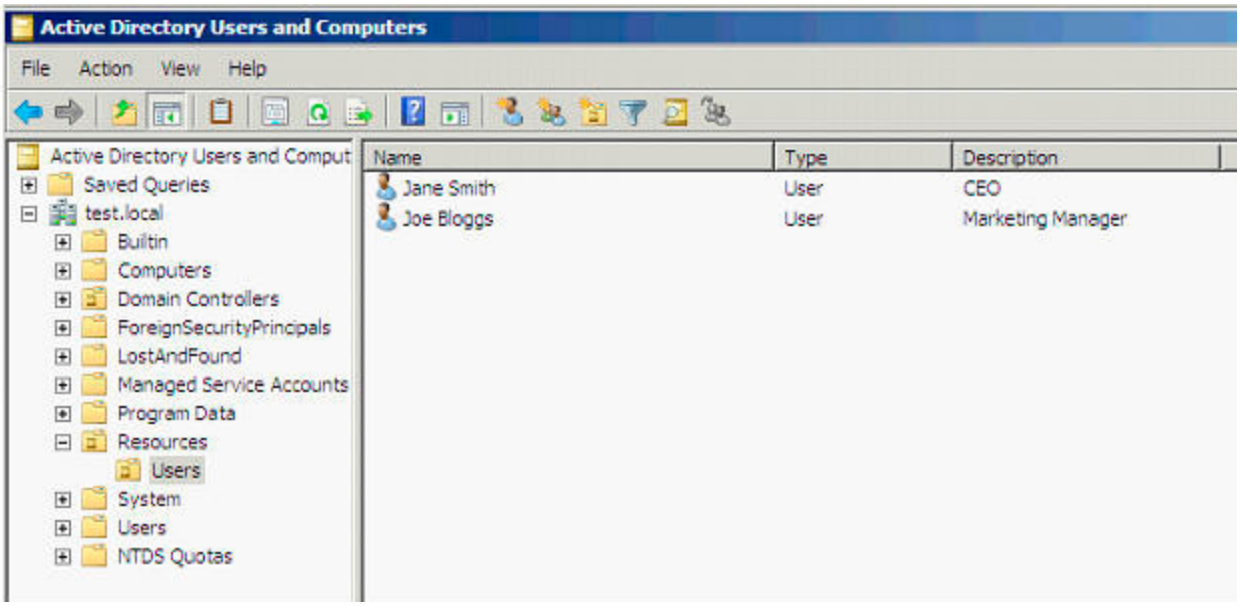

If we check the properties of the account we can confirm that, different from tombstone reanimation, we get all of the properties back.

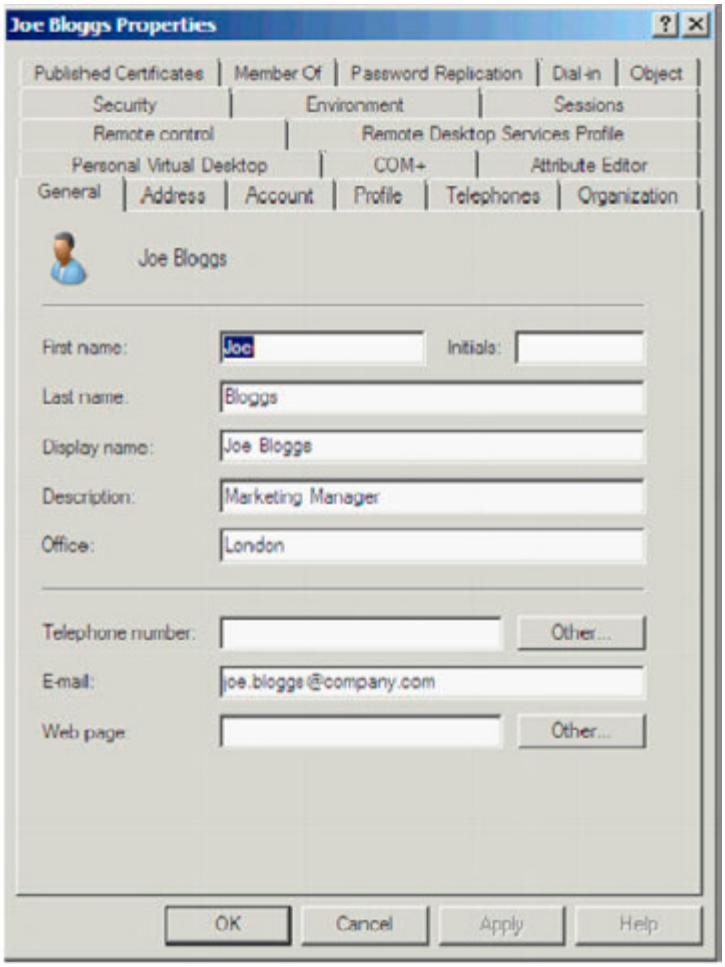

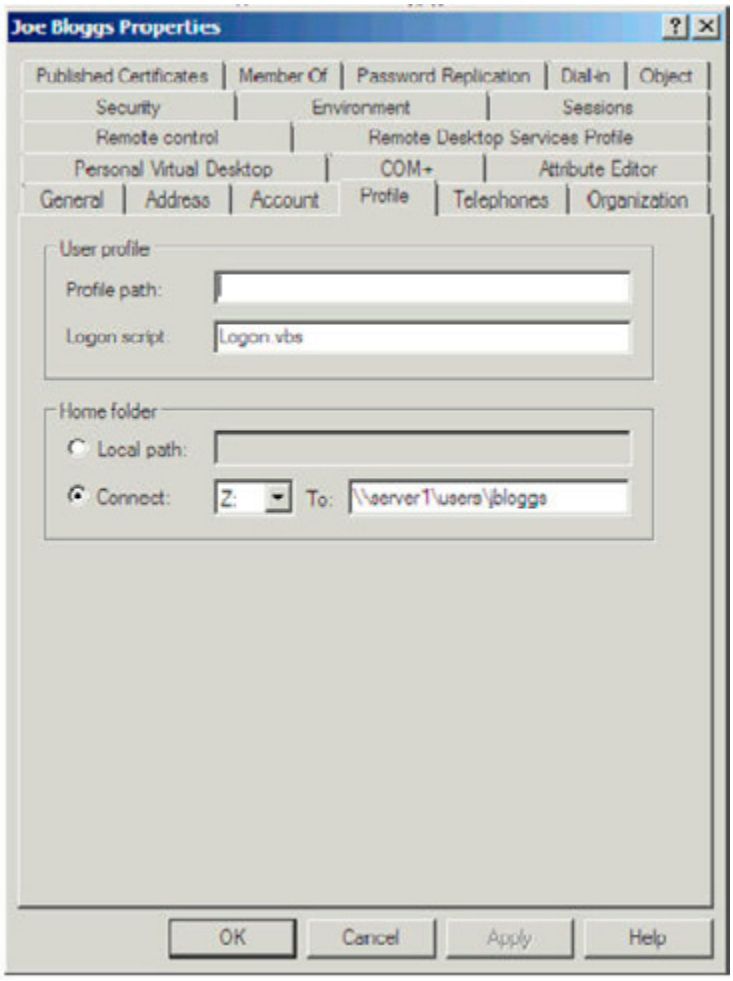

Active Directory Recycle Bin PowerPack for PowerGUI

Although the Recycle Bin is a great new feature within Windows Server 2008 R2 Microsoft is already getting feedback that there is no GUI for managing it. Whilst a lot of administrators are comfortable with PowerShell, some may still prefer to use a GUI based management tool for these tasks. Fortunately a great tool to plug this gap has already been provided by the community; PowerShell MVP Kirk Munro has created the Active Directory Recycle Bin PowerPack for PowerGUI (http://www.powergui.org/entry.jspa?categoryID=21&externalID=2461). This free tool has bundled up scripts using the previously demonstrated Active Directory PowerShell cmdlets and provides a graphical front end for administration.

Simply download the PowerGUI tool plus the Active Directory Recycle Bin PowerPack and import it into PowerGUI. Open up the PowerPack and you will have a graphical view of the current contents of the Recycle Bin with the ability to drill down through Organisational Units. Options for restoring single items or recursively are provided in the **Actions** column as well as alternate restoration paths and emptying items from the Recycle Bin.

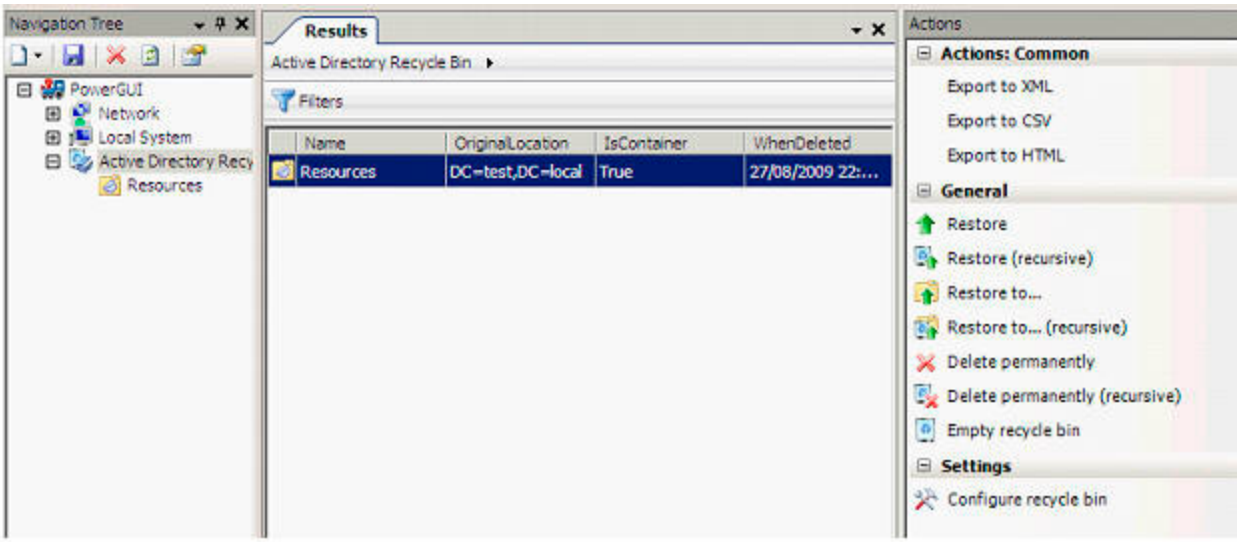

It is also possible to use the **Configure recycle bin** action to set the values for **DeletedObjectLifetime**, the amount of days objects reside in the Recycle Bin, and **TombstoneLifetime**, the amount of days objects can be restored using Tombstone Reanimation after they have left the Recycle Bin. In Windows Server 2008 R2 both of these values default to 180 days, in some earlier versions of Windows Server this value was 60 days and if you upgrade those domain controllers it will remain the same so you may wish to change the values – you can use the **Modify** action to do this.

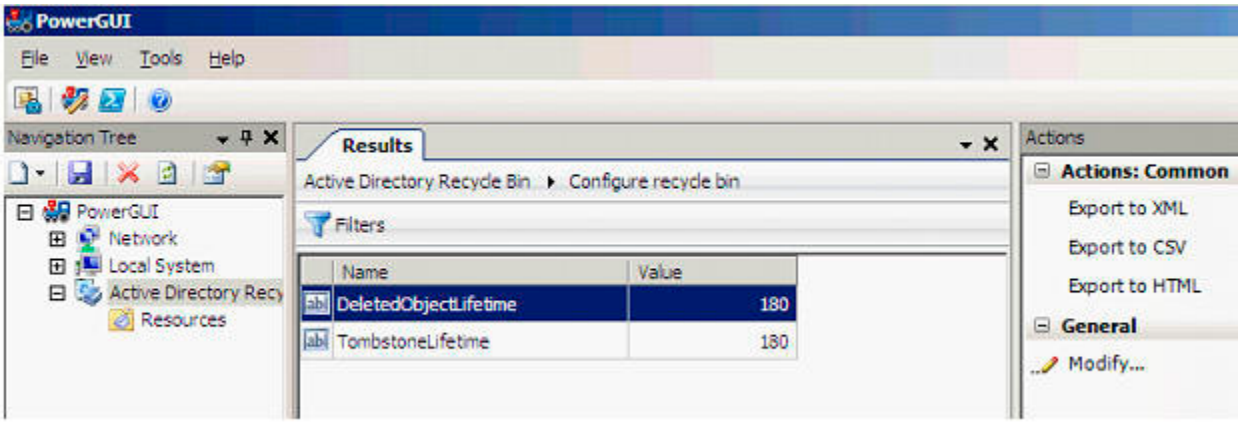

For this example I have deleted from Active Directory the **Resources** and **Users** containers and the two user accounts which you can see nicely in the below screenshot using PowerGUI.

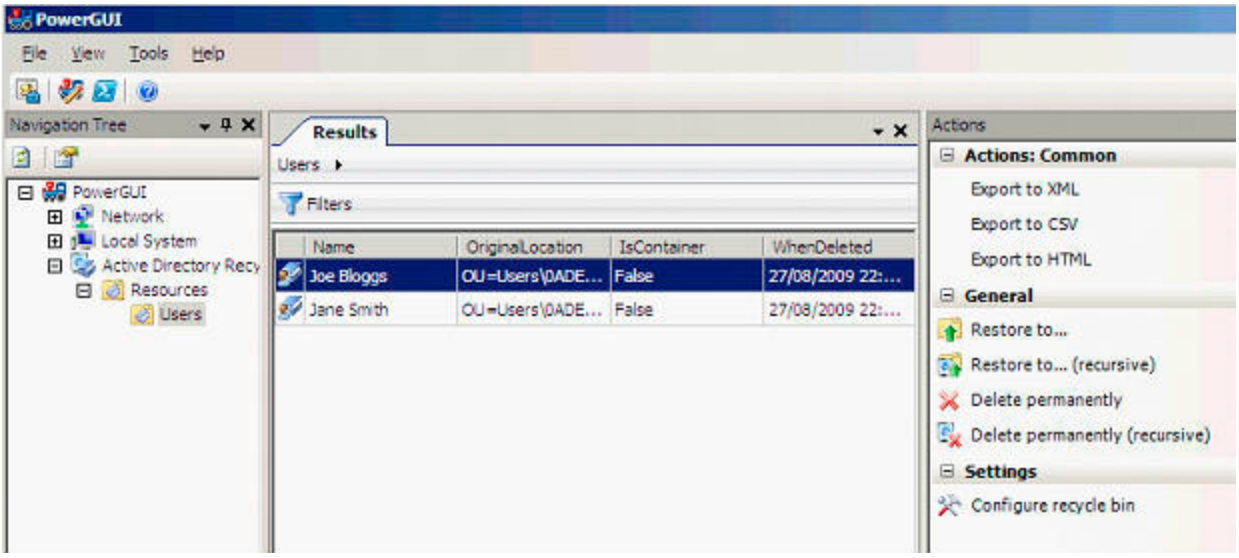

This time we will restore the account **Joe Bloggs**, but to an alternative location using the **Restore to....** Action. (Remember: this is done in PowerShell using the **–targetpath** parameter of the **Restore-ADObject cmdlet**) Simply input the path to the Organisational Unit you wish to restore the object to. In this example we use the default **Users** container as the target location.

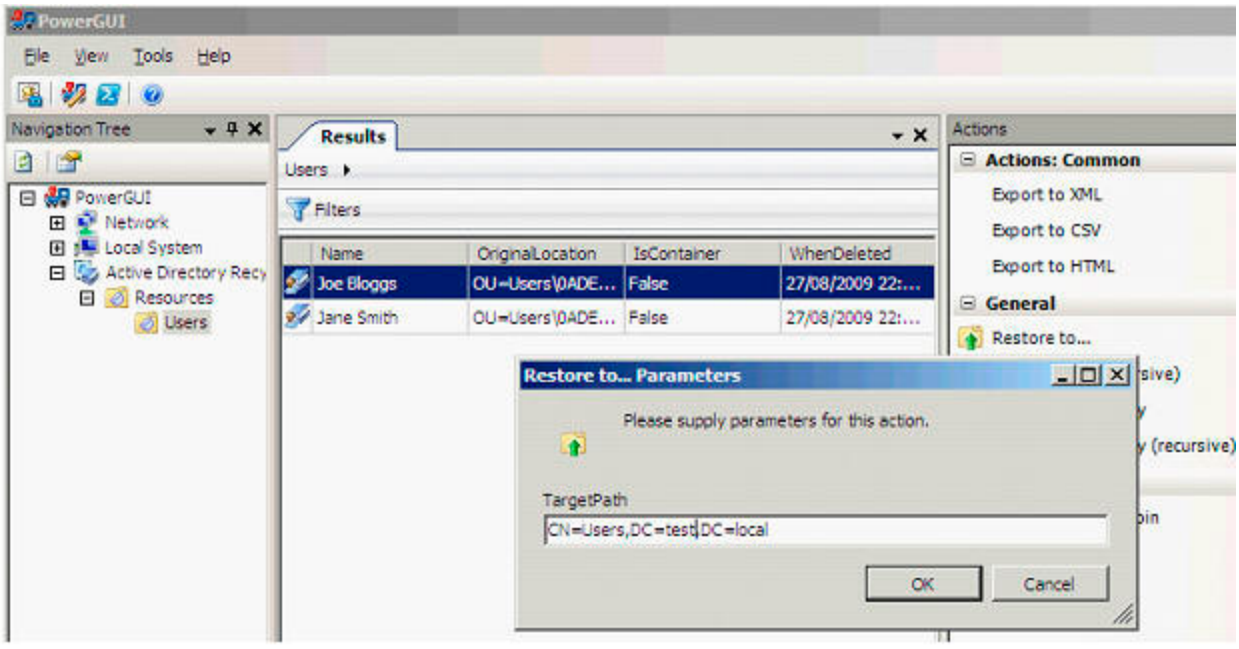

The user has been restored to the alternate location; this is particularly useful if we did not wish to bring back the entire OU(s) as we did previously.

If however, you do wish to bring back the contents of an entire OU and everything below it there is an action, **Restore (recursive)**.

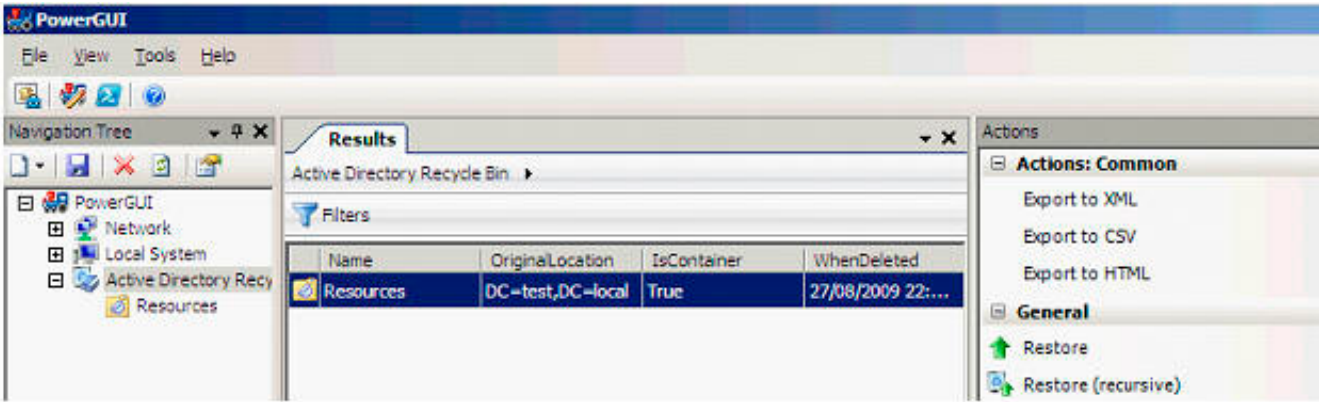

Using the **Restore (recursive)** action in this scenario brings back both the **Resources** and **Users** OU's as well as the single account remaining in it, **Jane Smith**.

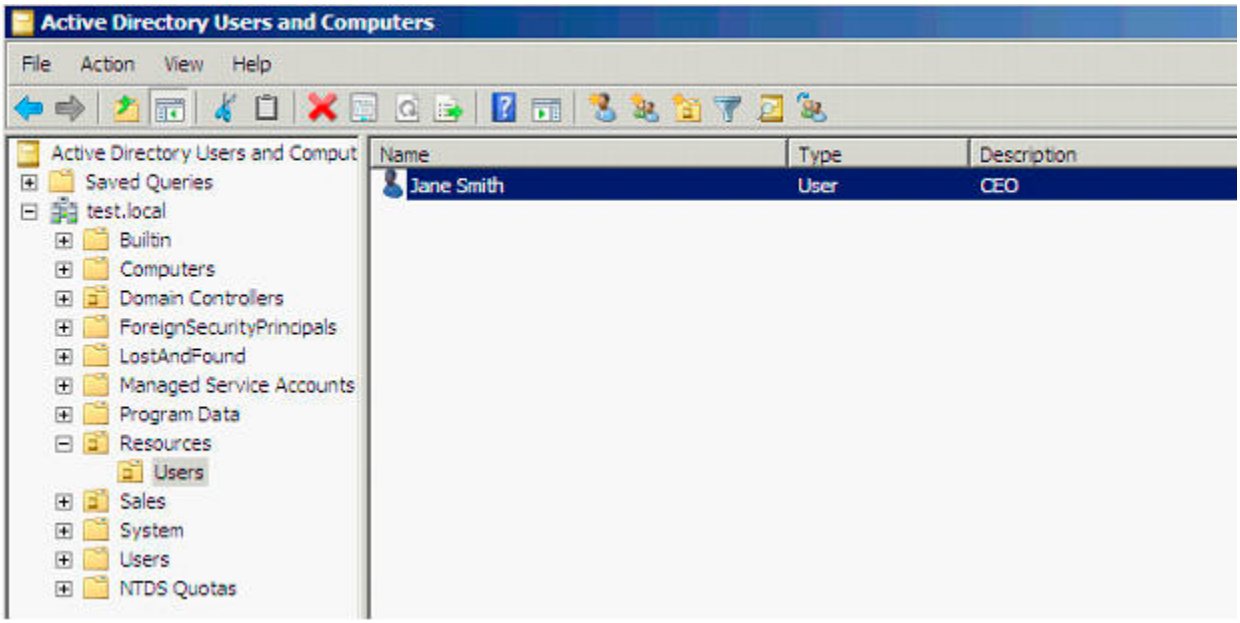

Hopefully in a future release of Windows Server this functionality will be provided out of the box, the most natural home would be a viewable container within Active Directory Users and Computers, until then the Recycle Bin PowerPack for PowerGUI will prove very useful.

Summary

One of the most requested features for a long time with Active Directory has been a Recycle Bin. Microsoft has finally delivered this with the release of Windows Server 2008 R2. It may not be a feature that enterprises get to use for a little while given the system requirements of all 2008 R2 Domain Controllers and your Active Directory Forest at 2008 R2 functional level, but it could be one of those compelling reasons that enables you to pursue an upgrade.

Administration is via the new Active Directory PowerShell cmdlets which Microsoft is using to provide a consistent command line interface across all of their products. Although currently there is no native GUI for these administration tasks, the Active Directory Recycle Bin PowerPack for PowerGUI enables administrators to leverage the underlying PowerShell functionality and provide a graphical interface for carrying out these tasks.

Using Group Policy to Restrict the use of PST Files

08 October 2009

by Ben Lye

> ***Outlook PST files are a problem for Exchange users, and give no benefits over Exchange mailboxes. You have to use them on a local drive, they are difficult to back up, and tricky for the administrator to manage. It is possible to use Exchange Group Policy settings to limit the use of PST files, and thereby alleviate some of the difficulties they cause.***

Outlook Personal Folder Files (PST files) are a data store that can be used by all versions of Outlook to store e-mail data. PST files have long been seen as a way to archive mail out of an Exchange mailbox, often to get the mailbox under a quota limit. However, the use of PST files causes Exchange administrators some serious pain when it comes to managing e-mail.

Briefly, these are some of the problems with PST files:

Microsoft does not support using PST files over a LAN or WAN network (KB267019). Using PST files located on network shares can slow down Outlook and can cause corruption of the PST file.

Anti-virus countermeasures cannot be implemented on PST files as easily as Exchange Server mailbox databases.

It is difficult to accurately report on PST file use, making reporting on organisational mail storage and planning for future growth difficult.

Managing content of PST files is difficult. Exchange Server provides tools to manage the content of mailboxes (such as Messaging Records Management) and to export or remove data from mailboxes (such as the Export-Mailbox cmdlet) but there are no such tools to manage the content of PST files.

Local PST files are difficult to back up, making them vulnerable to data loss.

Fortunately for the Exchange administrator, it is possible to restrict the ability to use PST files. There are two settings available, and both can be applied using Group Policy registry changes – **PSTDisableGrow** and **DisablePST**. PSTDisableGrow prevents new data being added to existing PST files, and DisablePST prevents users creating or opening PST files altogether.

Description	Registry Path	Registry Value
Disable PST files	HKEY_CURRENT_USER\Software\Policies\Microsoft\Office\12.0\Outlook	DisablePST
Prevent PST file growth	HKEY_CURRENT_USER\Software\Policies\Microsoft\Office\12.0\Outlook\PST	PstDisableGrow

Table - PST Restriction Registry Values

Note that the registry paths are specific to Outlook versions "12.0" refers to Outlook 2007, the registry path for Outlook 2003 would be ".\..\Office\11.0\Outlook" and so on.

In an environment where PST files already exist, these settings can be applied separately or together to phase out their use. The first step could be to implement restrictions on the growth of PST files using PSTDisableGrow which would allow users to access existing data but not allow it to be added to. Subsequently, all PST file use could be disabled by implementing DisablePST.

In a new Exchange environment, or one where PST files are not used (and the Exchange administrator wants to keep it that way), the DisablePST setting can be applied on its own to stop users being able to add PST files to Outlook.In any Exchange environment it is probably worth considering implementing a server-side archiving solution before disabling PST files. Server-side archiving has many benefits compared to PST files, and as many users are determined to keep large quantities of historic e-mail it is better to have a managed solution than unmanaged ad-hoc PST file use – a scenario often know as "PST hell."

If you are ready to disable PST file use the settings can be applied to Outlook 2007 with Group Policy using the Office 2007 Group Policy Administrative Templates.

Applying PST Group Policy for Outlook 2007

Download the Office 2007 ADM Templates and extract the files
http://www.microsoft.com/downloads/details.aspx?FamilyID=92d8519a-ef43-4aee-8f7a-e4bbaebaf3e7&displaylang=en

- Launch the Group Policy Management Console, click Start ▸ Administrative Tools ▸ Group Policy Management.
- Expand the Forest, Domains, and domain containers then select "Group Policy Objects."

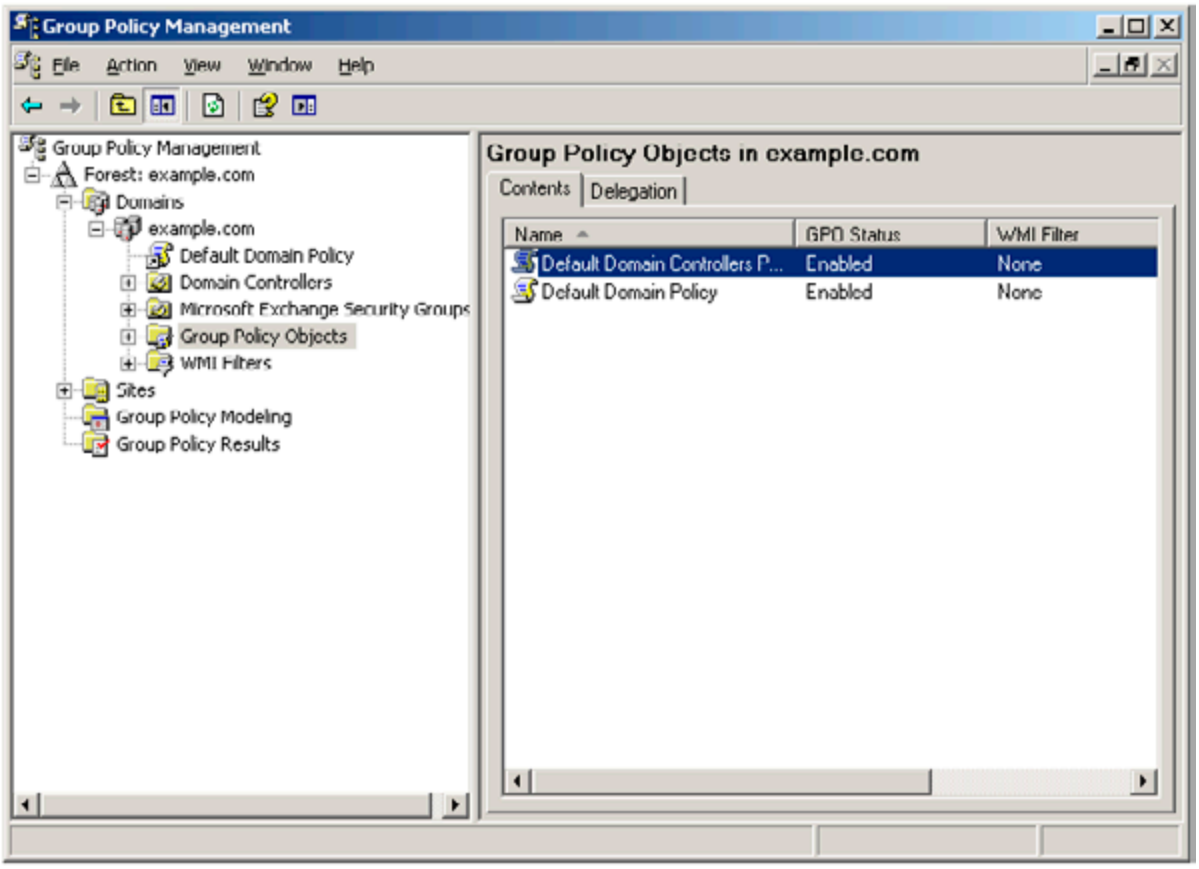

- Right-click "Group Policy Objects" and select "New." Give the new GPO a name, for example "PST Policy," and click OK. (Skip this step if you want to add these settings to an existing GPO.)
- Right-click "PST Policy" (or the existing policy you wish to edit)and choose "Edit."

- Expand "User Configuration," right-click "Administrative Templates" and choose "Add/Remove Templates."
- Click the "Add" button and browse to the location of the files extracted in step 2. Open the "ADM" folder and the appropriate language subfolder (en-us for English), select the file named "outlk12.adm" and click "Open."
- Click "Close" to close the "Add/Remove Templates" dialogue box.
- Expand "User Configuration\Administrative Templates\Microsoft Office Outlook 2007\Miscellaneous\PST Settings."

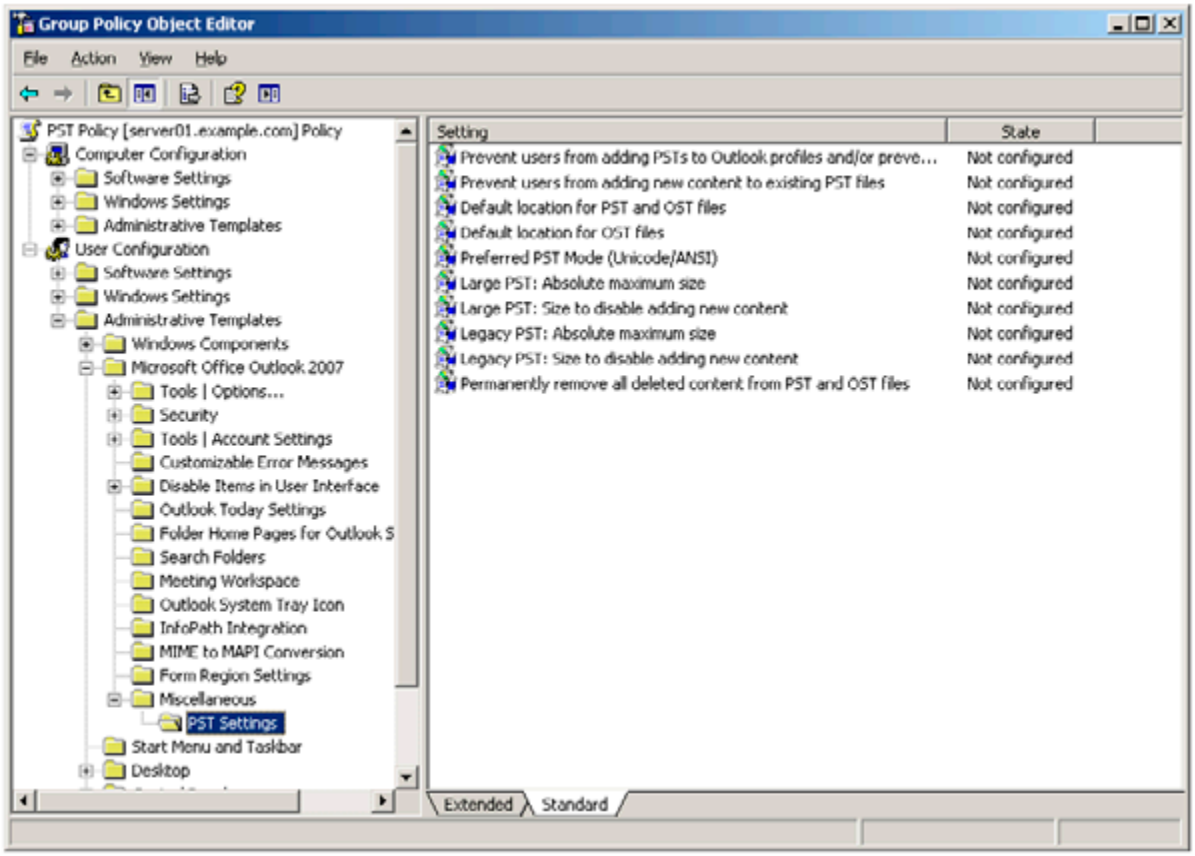

- To implement the DisablePST restriction enable the "Prevent users from adding PSTs to Outlook profiles..." setting and set the option to "Only Sharing-Exclusive PSTs can be added." This will allow PST files for application such as SharePoint lists, but will prevent user-created PST files from being added.

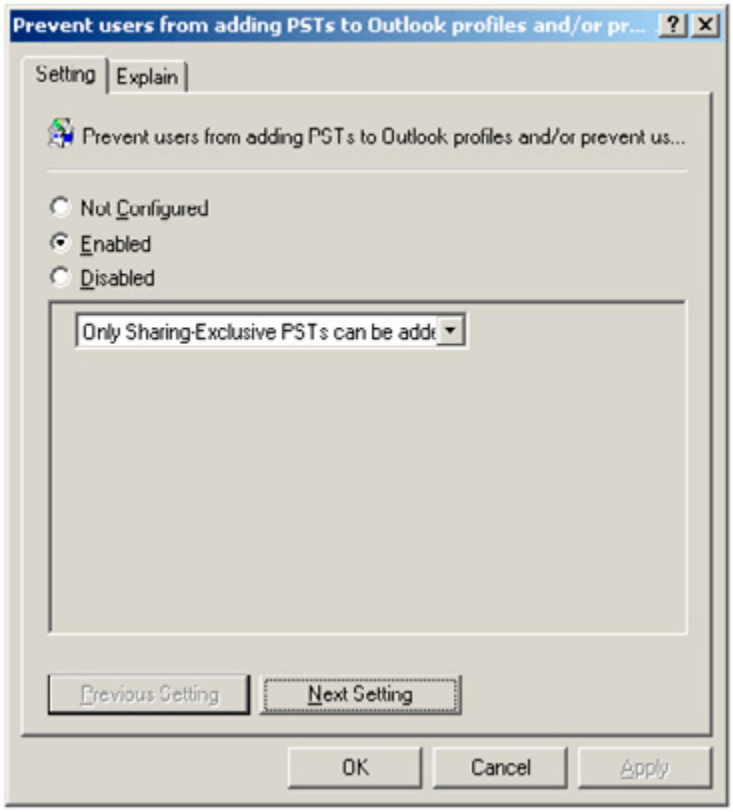

- To implement the PSTDisableGrow restriction enable the "Prevent users from adding new content to existing PST files." setting.

If the PSTDisableGrow setting is implemented users will still be able to create and open PST files, but they will not be able to add any data to any PST files. If they try they will receive this error message:

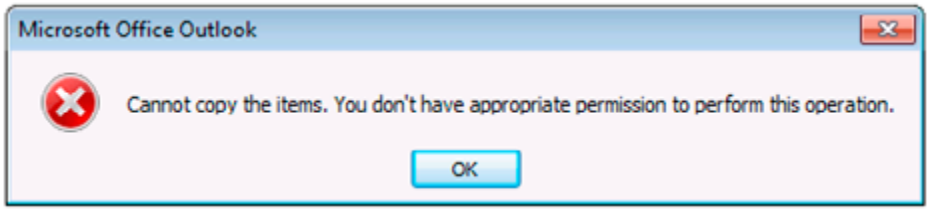

If the DisablePST setting is implemented the user will see changes in the Outlook user interface. While any PST files which were already loaded will remain part of the profile, the options to create new PST files or to open any other existing PST files will no longer be in the menu. Archive options will also be removed.

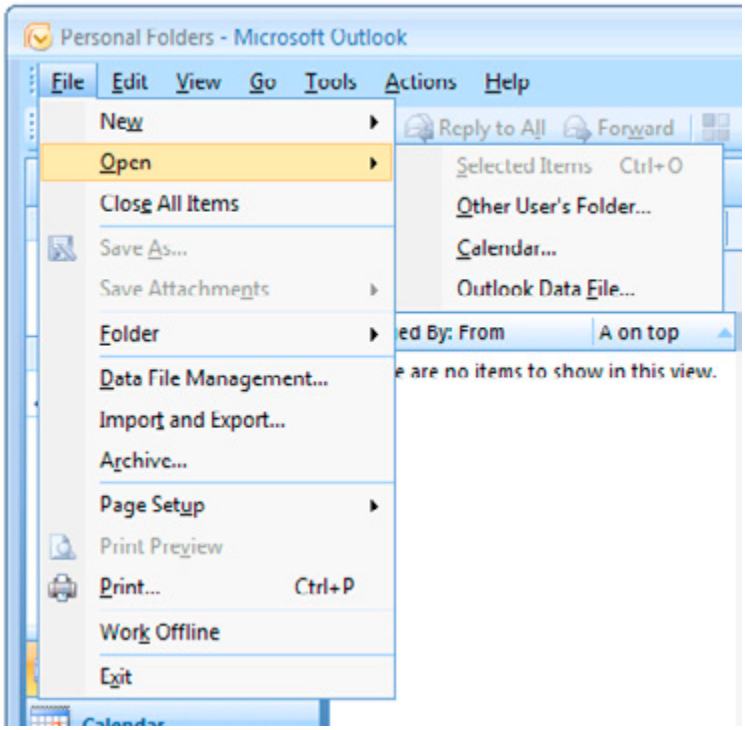

DisablePST not implemented.

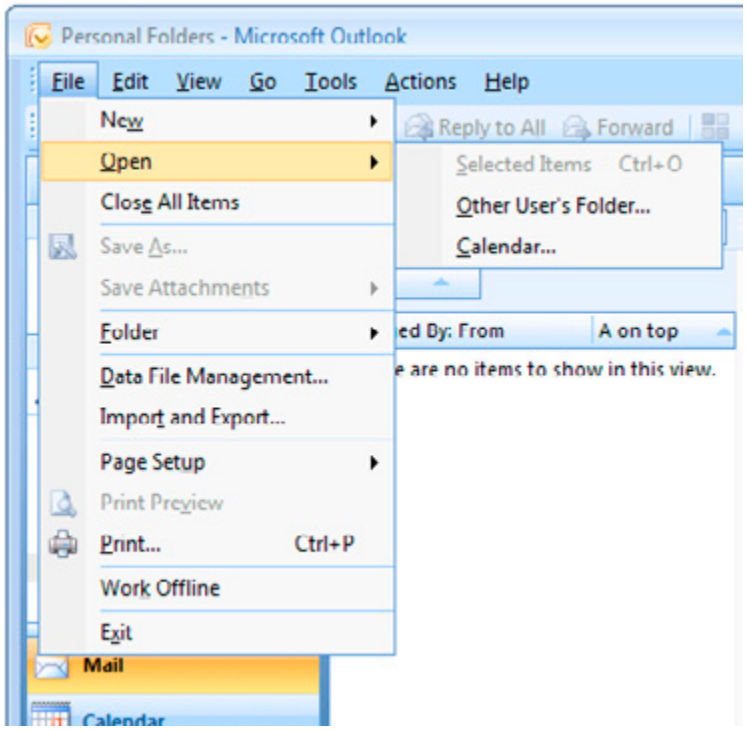

Disable PST implemented.

PST files can be a headache for Exchange administrators, but they don't have to be. With easily-applied Group Policy settings the use of PST files can be limited, and the problems they cause can be eradicated.

Introduction to Exchange Server 2010

22 October 2009

by Jaap Wesselius

What's new in Exchange Server 2010 and what features from Exchange Server 2007 have been deprecated? What has been done to further integrate Exchange with Active Directory? In an extract from his new book, Jaap gives a rooftop view of Exchange's new features and their significance, and spells out the main reasons why it is worth upgrading.

First things first – let's cover some basic background: Exchange Server 2010 is an e-mail and calendaring application that runs on Windows Server 2008 and, like its predecessor Exchange Server 2007, can also integrate with your phone system. It is the seventh major version of the product and, while not revolutionary, it does include some important changes and lots of small improvements over Exchange Server 2007.

The scalability of Exchange Server 2010 has improved, especially when compared to the complex storage requirements of Exchange Server 2007. The user experience has also improved in Outlook Web App, and a lot of complex issues have seen solved, or the complexity has been removed, to make the administrator's life much easier.

In this article I will give a brief overview of what's changed in Exchange Server 2010, what the new features are, what features have been removed, and how it makes your life as an Exchange administrator easier.

Under The Hood: What's changed?

- By far the most important change with respect to Exchange Server 2007 is the new Database Availability Group. This will allow you to create multiple copies of an Exchange Server database within your organization, and you are no longer bound to a specific site (like in Exchange Server 2007), but can now stretch across multiple sites. Microsoft has also successfully transformed Cluster Continuous Replication and Stand-by Continuous Replication into a new "Continuous Availability" technology.

- While on the topic of simplifying, a lot of SysAdmins were having difficulties with the Windows Server fail-over clustering, so Microsoft has simply "removed" this from the product. The components are still there, but they are now managed using the Exchange Management Console or Exchange Management Shell.

- With the new Personal Archive ability, a user can now have a secondary mailbox, acting as a personal archive - this really is a .PST killer! You now have the ability to import all the users" .PST files and store them in the Personal Archive, and using retention policies you can move data from the primary mailbox to the archive automatically, to keep the primary mailbox at an acceptable size, without any hassle.

- To deal with ever-growing storage requirements, Microsoft also made considerable changes to the underlying database system. All you will need to store your database and log files with Exchange Server 2010 is a 2 TB SATA (or other Direct Attached Storage) disk. As long as you have multiple copies of the database, you're safe! And the maximum supported database size? That has improved from 200 GB (in an Exchange Server 2007 CCR environment) to 2 TB (in a multiple database copy Exchange Server 2010 environment). If you haven't yet considered what your business case will look like when upgrading to Exchange Server 2010, bear in mind that this will truly safe a tremendous amount of storage cost - and that's not marketing talk!

- Installing Exchange 2010 is not at all difficult, and configuring a Database Availability Group with multiple copies of the Mailbox Databases is just a click of the mouse (you only have to be a little careful when creating multi-site DAGs). Even installing Exchange Server 2010 into an existing Exchange Server 2003 or Exchange Server 2007 environment is not that hard! The only thing you have to be aware of is the additional namespace that shows up. Besides the standard namespace like webmail.contose.com and Autodiscover.contoso.com, a third namespace shows up in a coexistence environment: legacy.contoso.com. This is used when you have mailboxes still on the old (i.e. Exchange Server 2003 or Exchange Server 2007) platform in a mixed environment.

- Lastly, for a die-hard GUI administrator it might be painful to start managing an Exchange environment with the Exchange Management Shell. Basic management can be done with the graphical Exchange Management Console, but you really do have to use the Shell for the nitty-gritty configuration. The Shell is remarkably powerful, and it takes quite some getting used to, but with it you can do fine-grained management, and even create reports using features like output-to-HTML or save-to-.CSV file. Very neat!

Getting Started

Exchange Server 2010 will be available in two versions:

- **Standard Edition**, which is limited to hosting 5 databases.

- **Enterprise Edition**, which can host up to 100 databases.

However, the available binaries are identical for both versions; it's the license key that establishes the difference in functionality. Exchange Server 2010 is also only available in a 64-Bit version; there is absolutely no 32-bit version available, not even for testing purposes. Bear in mind that, as 64-Bit-only software, there's no Itanium version of Exchange Server 2010.

Exchange Server 2010 also comes with two Client Access License (CAL) versions:

- **Standard CAL** – This license provides access to e-mail, calendaring, Outlook Web App and ActiveSync for Mobile Devices.

- **Enterprise CAL** – This is an additive license, and provides Unified Messaging and compliance functionality, as well as Forefront Security for Exchange Server and Exchange Hosted Filtering for anti-spam and anti-virus functionality.

This is not a complete list. For more information about licensing, check the Microsoft website at http://www.microsoft.com/exchange.

What's been removed from Exchange Server 2010?

As always, as new features come, old features go. There are inevitably a few that have found themselves on the "deprecated list" this time around, and so will not be continued in Exchange Server 2010 and beyond. Since this is a much shorter list than the "new features," we'll start here:

There are some major changes in Exchange Server clustering: in Exchange Server 2007 you had **LCR** (Local Continuous Replication), **CCR** (Cluster Continuous Replication) and **SCR** (Standby Continuous Replication) - three different versions of replication, all with their own management interfaces. All three are *no longer available* in Exchange Server 2010.

Windows **Server Fail-over Clustering** has been removed in Exchange Server 2010. Although seriously improved in Windows Server 2008, a lot of Exchange Administrators still found the fail-over clustering complex and difficult to manage. As a result, it was still prone to error and a potential source of all kinds of problems.

Storage Groups are no longer available in Exchange Server 2010. The concepts of a database, log files and a checkpoint file are still there, but now it is just called a Database. It's like CCR in Exchange Server 2007, where you could only have one Database per Storage Group.

Due to major reengineering in the Exchange Server 2010 databases, the **Single Instance Storage** (SIS) is no longer available. This means that when you send a 1 MB message to 100 recipients, the database will potentially grow by 100 MB. This will surely have an impact on the storage requirements in terms of space, but the performance improvements on the Database are really great. I'll get back on that later.

What's new in Exchange Server 2010?

Exchange Server 2010 contains a host of improvements and a lot of new features, as well as minor changes and improvements. Over the coming sections, I'll provide an overview of the most significant updates and additions.

Outlook Web App

The most visible improvement for end-users is Outlook Web App (previously known as Outlook Web Access). One of the design goals for the Outlook Web App was a seamless cross-browser experience, so users running a browser like Safari, even on an Apple MacBook, should have exactly the same user experience as users running Internet Explorer.

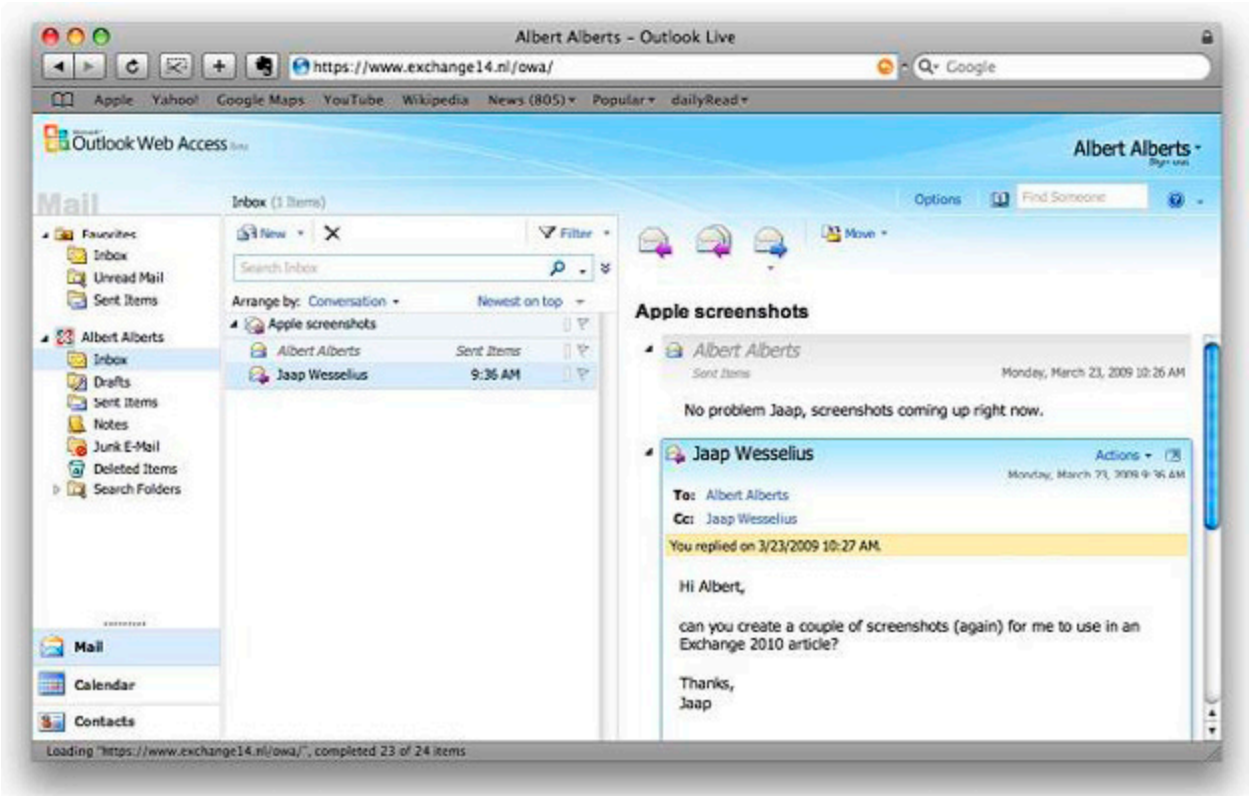

Figure 1. Outlook Web App running on an Apple MacBook using a Safari browser!

Outlook Web App offers a very rich client experience and narrows the gap between a fully-fledged Outlook client and Outlook Web Access. To reinforce that experience, a lot of new features have been introduced. To name a few: Favorites, Search Folders, attaching messages to messages, integration with Office Communicator, a new Conversation View (which works very well!), integration with SMS (text) messages and the possibility to create Outlook Web Access policies, which give the Exchange organization administrator the ability to fine tune the user experience. The Web App is a feature which you will find mentioned throughout the book.

High Availability

The Exchange Server 2007 Cluster Continuous Replication (CCR) and Standby Continuous Replication (SCR) features are now combined into one new feature called **database availability**.

Database copies exist just as in an Exchange Server 2007 CCR environment and are created in a "Database Availability Group," but it is now possible to create multiple copies. The replication is not on a server level as in Exchange Server 2007 but on a database level, which gives the Exchange administrator much more fine control and granularity when it comes to creating a high available Exchange organization. The servers in such a Database Availability Group can be at the same location, or other locations to create an offsite solution. There's also no longer any need to install the Microsoft Cluster Service (MSCS) before setting up the Database Availability Group, as all cluster operations are now managed by Exchange.

Exchange Core Store functionality

Compared to Exchange Server 2003, Exchange Server 2007 dramatically decreased the I/O on the disk subsystem (sometimes by 70%). This was achieved by increasing the Exchange database page size from 4KB to 8KB and by using the 64-Bit operating system. The memory scalability of the 64-Bit platform makes it possible to use servers with huge amounts of memory, giving them the opportunity to cache information in memory instead of reading and writing everything to the disk.

One of the design goals of Exchange Server 2010 was to use a single 1TB SATA disk for the mailbox database *and* its log files. Another goal was to allow multi GB mailboxes without any negative performance impact on the server. To make this possible, the database schema in Exchange Server 2010 has now been flattened, making the database structure used by the Exchange Server *much* less complex than it was in Exchange Server 2007 and earlier. As a result, the I/O requirements of an Exchange Server 2010 server can be up to *50% less* than for the same configuration in Exchange Server 2007.

As a result of the flattened database schema, Microsoft has removed Single Instance Storage (SIS) from Exchange Server 2010, but the improvements in performance are much more significant, and more-than-adequate compensation for the (comparatively minor) loss of SIS.

Microsoft Online Services

Microsoft is gradually moving "into the cloud." Besides an Exchange Server 2010 implementation on premise, it is now also possible to host mailboxes in a datacenter; you can host your mailboxes with your own ISP, or with Microsoft Online Services.

Exchange Server 2010 can be 100% on premise, 100% hosted, or it can be a mixed environment, with some percentage of your mailboxes hosted and the rest on premise. This is, of course, fully transparent to end users, but it has its effects on the administration. Instead of managing just one, on-site environment, you'll have to manage the hosted organization as well. This is can all be handled through Exchange Server 2010's Exchange Management Console, where you can connect to multiple forests containing an Exchange organization.

New Administration Functionality

As a consequence of the major changes made to the High Availability features of Exchange Server 2010, the Exchange Management Console has also changed rather significantly.

Due to the new replication functionality, the Mailbox object is no longer tied to the Exchange Server object, but is now part of the Exchange Server 2010 organization. Also, since the concept of Storage Groups is no longer relevant, their administration has been removed from both the Exchange Management Console and the Exchange Management Shell. PowerShell cmdlets like New-StorageGroup, Get-StorageGroup, and so on, have also all been removed, although the options of these cmdlets have been moved into other cmdlets, like database-related cmdlets.

Speaking of which, Exchange Server 2010 also runs on top of **PowerShell Version 2**. This version not only has a command line interface (CLI), but also an Interactive Development Environment (IDE). This enables you to easily create scripts and use variables, and you now have an output window where you can quickly view the results of your PowerShell command or script.

In addition to PowerShell V2, Exchange Server 2010 also uses **Windows Remote Management** (WinRM) Version 2. This gives you the option to remotely manage an Exchange Server 2010 server without the need to install the Exchange Management Tools on your workstation, and even via the Internet!

One last small but interesting new feature is "**Send Mail**," allowing you to send mail directly from the Exchange Management Console - ideal for testing purposes.

Exchange Control Panel

It is now possible to perform some basic Exchange management tasks using the options page in Outlook Web Access; not only on the user's own properties, but also at an organizational level. With this method, it is possible to create users, mailboxes, distribution groups, mail-enabled contact, management e-mail addresses, etc.

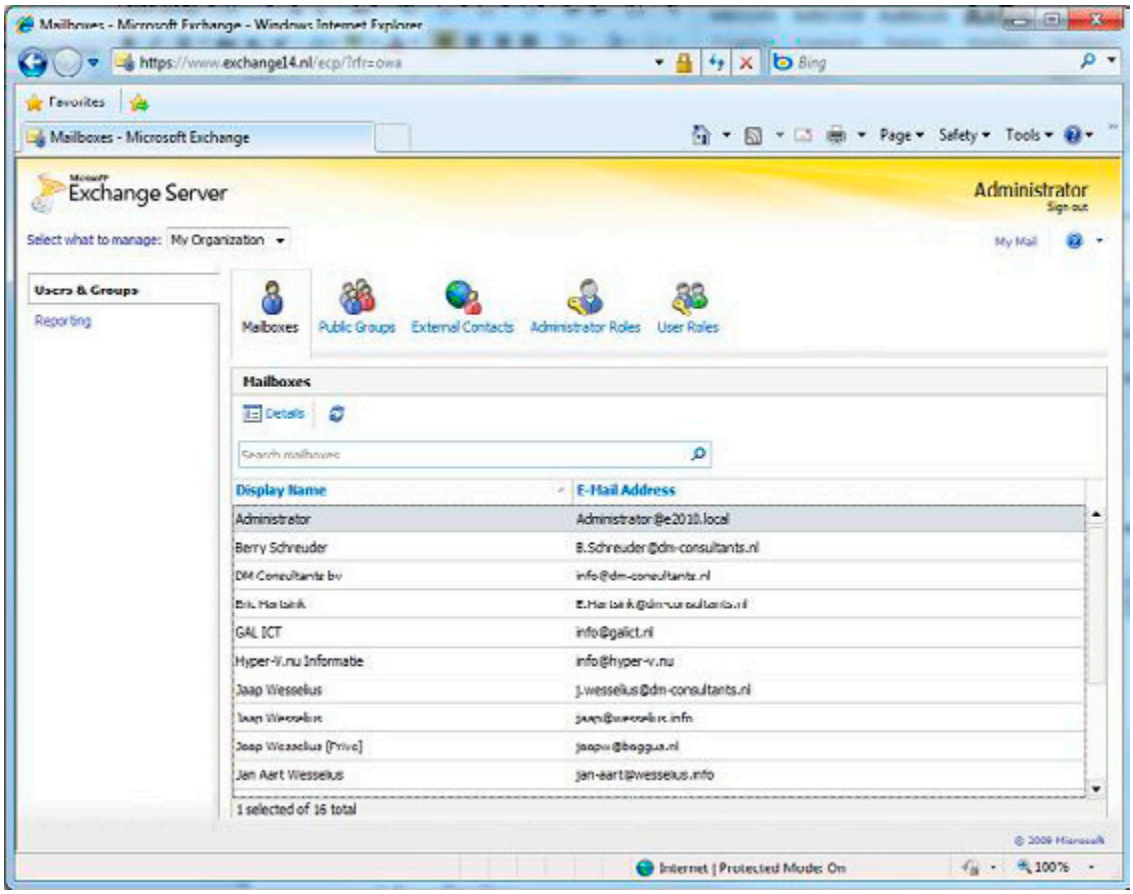

Figure 2. The Exchange Control Panel for basic management functions.

Active Directory Rights Management

Active Directory Rights Management Service lets you control what users can do with E-mail and other documents that are sent to them. It is possible, for example, for classified messages to disable the "Forward" option to prevent messages being leaked outside the organization. With Exchange Server 2010, new features have been added to the Rights Management Services, such as:

- **Integration with Transport Rules** - a template for using RMS to protect messages over the Internet.
- **RMS protection for voice mail messages** coming from the Unified Messaging Server Role.

Active Directory is discussed throughout this book, as the Exchange Server 2010 has a much close relationship with AD than previous versions of Exchange Server.

Transport and Routing

With Exchange Server 2010 it is possible to implement **cross premises message routing**. When using a mixed hosting environment, Exchange Server 2010 can route messages from the datacenter to the on-premise environment with full transparency.

Exchange Server 2010 also offers (at last) **enhanced disclaimers**, making it possible to add HTML content to disclaimers to add images, hyperlinks, etc. It is even possible to use Active Directory attributes (from the user's private property set) to create a personal disclaimer.

To create a highly available and reliable routing model, the Hub Transport Servers in Exchange Server 2010 now contain **Shadow Redundancy**. A message is normally stored in a database on the Hub Transport Server and, in Exchange Server 2007, the message is deleted as soon as it is sent to the next hop. In Exchange Server 2010, the message is *only* deleted after the next hop reports a successfully delivery of the message. If this is not reported, the Hub Transport Server will try to resend the message.

For more High Availability messaging support, the messages stay in the transport dumpster on a Hub Transport Server, and are only deleted if they are successfully replicated to all database copies. The database on the Hub Transport Server has also been improved on an ESE level, resulting in a higher message throughput on the transport level.

Permissions

Previous versions of Exchange Servers relied on delegation of control via multiple Administrative Groups (Specifically, Exchange Server 2000 and Exchange Server 2003) or via Group Membership. Exchange Server 2010 now contains a **Role Based Access Model (RBAC)** to implement a powerful and flexible management model.

Messaging Policy and Compliance

As part of a general compliance regulation, Microsoft introduced the concept of Managed Folders in Exchange Server 2007, offering the possibility to create some sort of compliancy feature. This has been enhanced with new interfaces in Exchange Server 2010, such as the option of tagging messages, cross mailbox searches and new transport rules and actions.

Mailbox Archive

Exchange Server 2010 now contains a personal archive; this is a secondary mailbox connected to a user's primary mailbox, and located in the same Mailbox Database as the user's primary mailbox. Since Exchange Server 2010 now supports a JBOD (Just a Bunch of Disks) configuration this isn't too big a deal, and the Mailbox Archive really is a great replacement of (locally stored) .PST files.

Unified Messaging

The Exchange Server 2010 Unified Messaging Server Role integrates a telephone system, like a PABX, with the Exchange Server messaging environment. This makes it possible to offer Outlook Voice Access, enabling you to interact with the system using your voice, listen to voice mail messages, or have messages read to you. Exchange Server 2010 offers some new functionality like **Voicemail preview**, **Messaging Waiting Indicator**, integration **with text (SMS) messages**, additional **language support** etc. Unified Messaging is, unfortunately, a little outside the scope of this book, so you won't find me going into too much detail later on.

Exchange Server 2010 and Active Directory

As far as Active Directory is concerned, its minimum level needs to be on a Windows Server 2003 level, both for the domain functional level as well as the forest functional level. This might be confusing, since Exchange Server 2010 only runs on Windows Server 2008 or Windows Server 2008 R2, but that's just the actual server which Exchange Server 2010 is running on!

The Schema Master in the forest needs to be Windows Server 2003 SP2 server (Standard or Enterprise Edition) or higher. Likewise, in each Active Directory Site where Exchange Server 2010 will be installed, there must be *at least* one Standard or Enterprise Windows Server 2003 SP2 (or higher) server configured as a Global Catalog server.

From a performance standpoint, as with Exchange Server 2007, the ratio of 4:1 for Exchange Server processors to Global Catalog server processors still applies to Exchange Server 2010. Using a 64-Bit version of Windows Server for Active Directory will naturally also increase the system performance.

> ***Note***
>
> *It is possible to install Exchange Server 2010 on an Active Directory Domain Controller. However, for performance and security reasons it is recommended not to do this, and instead to install Exchange Server 2010 on a member server in a domain.*

Active Directory partitions

A Windows Server Active Directory consists of one forest, one or more domains and one or more sites. Exchange Server 2010 is bound to a forest, and therefore one Exchange Server 2010 Organization is connected to one Active Directory forest. The actual information in an Active Directory forest is stored in three locations, also called partitions:

- **Schema partition** – this contains a "blue print" of all objects and properties in Active Directory. In a programming scenario this would be called a class. When an object, like a user, is created, it is instantiated from the user blueprint in Active Directory.
- **Configuration partition** – this contains information that's used throughout the forest. Regardless of the number of domains that are configured in Active Directory, all domain controllers use the same Configuration Partition in that particular Active Directory forest. As such, it is replicated throughout the Active Directory forest, and all changes to the Configuration Partition have to be replicated to all Domain Controllers. All Exchange Server 2010 information is stored in the Configuration Partition.
- **Domain Partition** – this contains information regarding the domains installed in Active Directory. Every domain has its own Domain Partition, so if there are 60 domains installed there will be 60 different Domain Partitions. User information, including Mailbox information, is stored in the Domain Partition.

Delegation of Control

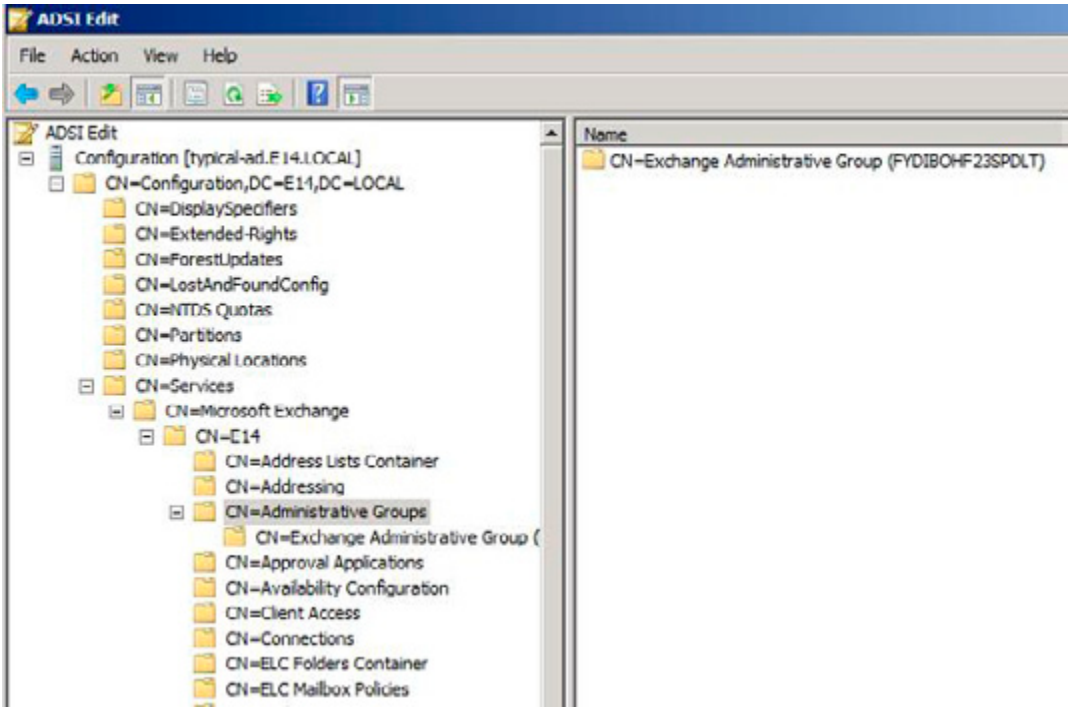

Figure 3. The Configuration partition in Active Directory holds all information regarding Exchange Server 2010 in an Administrative Group.

In Exchange Server 2003 the concept of "Administrative Groups" was used to delegate control between different groups of administrators. A default "First Administrative Group" was created during installation, and subsequent Administrative Groups could be created to install more Exchange 2003 servers and delegate control of these servers to other groups. The Administrative Groups were stored in the Configuration Partition so all domains and thus all domain controllers and Exchange servers could see them.

> *Just shift all letters in the word* **FYDIBOHF23SPDLT** *to the left and you get* **EXCHANGE12ROCKS.**

Exchange Server 2007 used Active Directory Security Groups for delegation of control, and only one Administrative Group is created during installation of Exchange Server 2007, called "Exchange Administrative Group (FYDIBOHF23SPDLT)." All servers in the organization are installed in this Administrative Group. Permissions are assigned to Security Groups and Exchange administrators are member of these Security Groups.

Exchange Server 2010 uses the same Administrative Group, but delegation of control is not done using Active Directory Security Groups, as Microsoft has introduced the concept of "Role Based Access Control" or RBAC.

Active Directory Sites

Exchange Server 2010 uses Active Directory Sites for routing messages. But what is an Active Directory site?

When a network is separated into multiple physical locations, connected with "slow" links and separated into multiple IP subnets, then in terms of Active Directory we're talking about sites. Say, for example, there's a main office located in Amsterdam, this has an IP subnet of 10.10.0.0/16. There's a Branch Office located in London, and this location has an IP subnet of 10.11.0.0/16. Both locations have their own Active Directory Domain Controller, handling authentication for clients in their own subnet. Active Directory site links are created to control replication traffic between sites. Clients in each site use DNS to find services like Domain Controllers in their own site, thus preventing using services over the WAN link.

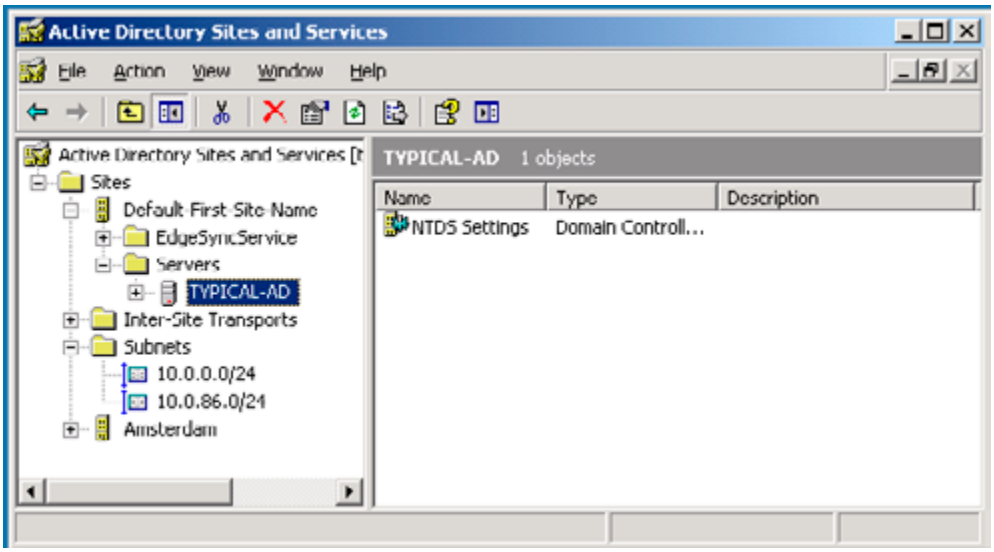

Figure 4. Two subnets in Active Directory, one for the main office and one for the Amsterdam Datacenter.

Exchange Server 2010 uses Active Directory sites for routing messages between sites. Using our current example, if there is an Exchange Server 2010 Hub Transport Server in Amsterdam and an Exchange Server 2010 Hub Transport Server in London, then the IP Site Links in Active Directory are used to route messages from Amsterdam to London. This concept was first introduced in Exchange Server 2007, and nothing has changed in Exchange Server 2010.

Exchange Server 2003 used the concept of Routing Groups, where Active Directory already used Active Directory Sites; Active Directory Sites and Exchange Server Routing Groups are not compatible with each other. To have Exchange Server 2003 and Exchange Server 2010 work together in one Exchange organization, some special connectors have to be created - the so called Interop Routing Group Connector.

Exchange Server coexistence

It is very likely that large organizations will gradually move from an earlier version of Exchange Server to Exchange Server 2010, and Exchange Server 2010 can coexist, in the same forest, with (both) Exchange Server 2007 and Exchange Server 2003. It is also possible to move from a mixed Exchange Server 2003 and Exchange Server 2007 environment to Exchange Server 2010.

Please note that it is not possible to have a coexistence scenario where Exchange Server 2000 and Exchange Server 2010 are installed in the same Exchange Organization. This is enforced in the setup of Exchange Server 2010. If the setup detects an Exchange Server 2000 installation the setup application is halted and an error is raised.

Integrating Exchange Server 2010 into an existing Exchange Server 2003 or Exchange Server 2007 environment is called a "transition" scenario. A "migration" scenario is where a new Active Directory forest is created where Exchange Server 2010 is installed. This new Active Directory forest is running in parallel to the "old" Active Directory with a previous version of Exchange Server. Special care has to be taken in this scenario, especially when both organizations coexist for any significant amount of time. Directories have to be synchronized during the coexistence phase, and the free/busy information will need to be constantly synchronized as well, since you'll still want to offer this service to users during the coexistence period.

This is a typical scenario when 3rd-party tools like Quest are involved, although it is not clear at the time of writing this book how Quest is going to deal with Exchange Server 2010 migration scenarios.

Exchange Server 2010 Server roles

Up until Exchange Server 2003, all roles were installed on one server and administrators were unable to select which features were available. It was possible to designate an Exchange 2000 or Exchange 2003 server as a so called "front-end server," but this server was just like an ordinary Exchange server acting as a protocol proxy. It still had a Mailbox Database and a Public Folder database installed by default.

Exchange Server 2007 introduced the concept of "server roles" and this concept is maintained in Exchange Server 2010. The following server roles, each with a specific function, are available in Exchange Server 2010:

- Mailbox Server (MB) role.
- Client Access Server (CAS) role.
- Hub Transport Server (HT) role.
- Unified Messaging Server (UM) role.
- Edge Transport Server (Edge) role.

These server roles can be installed on dedicated hardware, where each machine has its own role, but they can also be combined. A typical server installation, for example in the setup program, combines the Mailbox, Client Access and Hub Transport Server role. The Management Tools are always installed during installation, irrespective of which server role is installed.

By contrast, the Edge Transport Server role cannot be combined with *any* other role. In fact, the Edge Transport Server role cannot even be part of the (internal) domain, since it is designed to be installed in the network's Demilitarized Zone (DMZ).

There are multiple reasons for separating Exchange Server into multiple server roles:

- **Enhanced scalability** – since one server can be dedicated for one server role, the scalability profits are huge. This specific server can be configured and optimized for one particular Role, resulting in a high performance server.
- **Improved security** – one dedicated server can be hardened for security using the Security Configuration Wizard (SCW). Since only one Server Role is used on a particular server, all other functions and ports are disabled, resulting in a more secure system.
- **Simplified deployment and administration** – a dedicated server is easier to configure, easier to secure and easier to administer.

I will explain each server role in detail, in the following sections.

Mailbox Server role

The Mailbox Server role is the heart of your Exchange Server 2010 environment. This is where the Mailbox Database and Public Folder Database are installed. The sole purpose of the Mailbox Server role is to host Mailboxes and Public Folders; nothing more. In previous versions of Exchange Server, including Exchange Server 2007, Outlook clients using MAPI still connected directly to the Mailbox Server Role, but with Exchange Server 2010 this is no longer the case. MAPI clients now connect to a service called "MAPI on the Middle Tier" (MoMT), running on the Client Access Server. The name MoMT is still a code name and will have been changed before Exchange Server 2010 is released.

The Mailbox Server Role does not route any messages, it only stores messages in mailboxes. For routing messages, the Hub Transport Server role is needed. This latter role is responsible for routing all messages, even between mailboxes that are on the same server, and even between mailboxes that are in the same mailbox database.

For accessing mailboxes, a Client Access Server is also always needed; it is just not possible to access any mailbox without a Client Access Server.

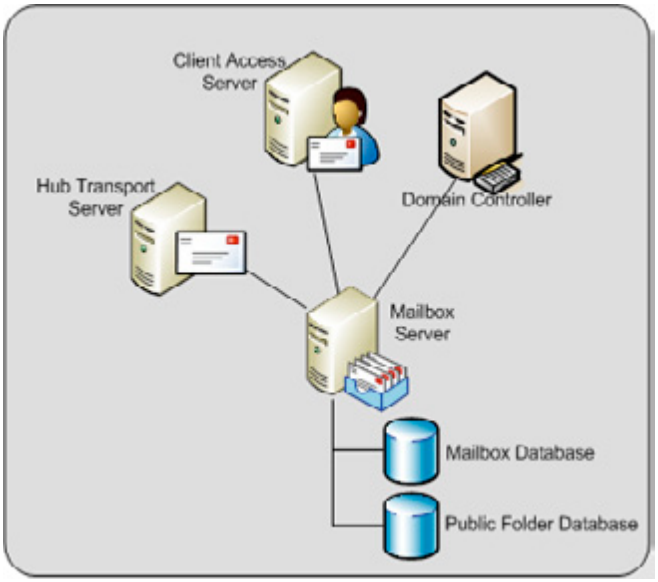

Figure 5. The Mailbox Server role is hosting Mailboxes and Public Folders.

Note that Internet Information Server is needed on a Mailbox Server role in order to implement the Role Based Access Control model (RBAC) even if no client is accessing the Mailbox Server directly.

As I mentioned, Storage Groups no longer exist in Exchange Server 2010, but mailboxes are still stored in databases, just like in Exchange Server 2007. Although rumors have been circulating for more than 10 years that the database engine used in Exchange Server will be replaced by a SQL Server engine, it has not happened yet. Just as in earlier versions of Exchange Server, the Extensible Storage Engine (ESE) is still being used, although major changes have been made to the database and the database schema.

By default, the first database on a server will be installed in the directory:

```
C:\Program Files\ Microsoft\Exchange Server\V14\Mailbox\Mailbox Database <<identifier>>
```

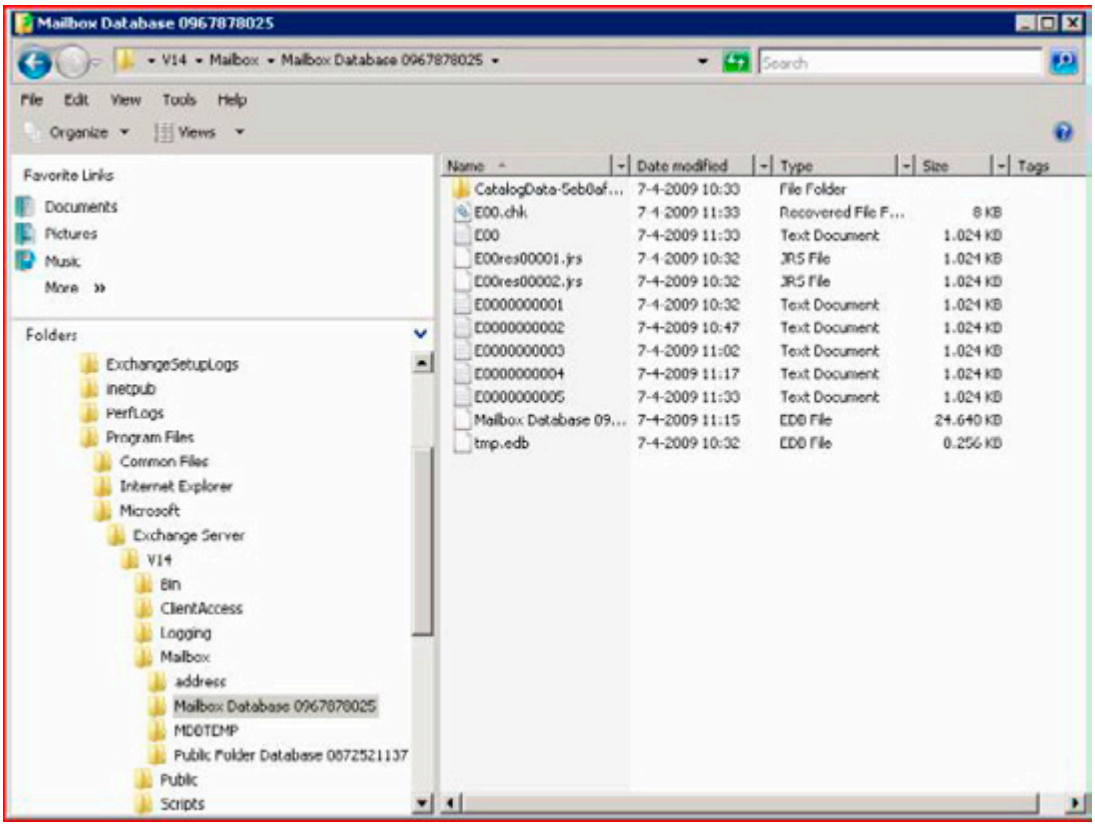

Figure 6. The default location for the Mailbox Databases and the log files.

The <<identifier>> is a unique number to make sure that the Mailbox Database name is unique within the Exchange organization.

It is a best practice, from both a performance and a recovery perspective, to place the database and the accompanying log files on a dedicated disk. This disk can be on a Fiber Channel SAN, an iSCSI SAN, or on a Direct Attached Storage (DAS) solution. Whilst it was a design goal to limit the amount of disk I/O to a level that both the database and the log files could be installed on a 1TB SATA disk, this is only an option if Database Copies are configured and you have at least two copies of the Mailbox Database, in order to avoid a single point of failure.

Client Access Server role

The Client Access Server XE "Client Access Server" role offers access to the mailboxes for all available protocols. In Exchange Server 2003, Microsoft introduced the concept of "front-end" and "back-end" servers, and the Client Access Server role is comparable to an Exchange Server 2003 front-end server.

All clients connect to the Client Access Server and, after authentication, the requests are proxied to the appropriate Mailbox Server. Communication between the client and the Client Access Server is via the normal protocols (HTTP, IMAP4, POP3 and MAPI), and communication between the Client Access Server and the Mailbox Server is via Remote Procedure Calls (RPC).

The following functionality is provided by the Exchange Server 2010 Client Access Server:

- HTTP for Outlook Web App.
- Outlook Anywhere (formerly known as RPC/HTTP) for Outlook 2003, Outlook 2007 and Outlook 2010.
- ActiveSync for (Windows Mobile) PDAs.
- Internet protocols POP3 and IMAP4.
- MAPI on the Middle Tier (MoMT).
- Availability Service, Autodiscover and Exchange Web Services. These services are offered to Outlook 2007 clients and provide free/busy information, automatic configuration of the Outlook 2007 and Outlook 2010 client, the Offline Address Book downloads and Out-of-Office functionality.

Note

SMTP Services are not offered by the Client Access Server. All SMTP Services are handled by the Hub Transport Server.

At least one Client Access Server is needed for each Mailbox Server in an Active Directory site, as well as a fast connection is between the Client Access Server and the Mailbox Server. The Client Access Server also needs a fast connection to a Global Catalog Server.

The Client Access Server should be deployed on the internal network and NOT in the network's Demilitarized Zone (DMZ). In order to access a Client Access Server from the Internet, a Microsoft Internet Security and Acceleration (ISA) Server should be installed in the DMZ. All necessary Exchange services should be "published" to the Internet, on this ISA Server.

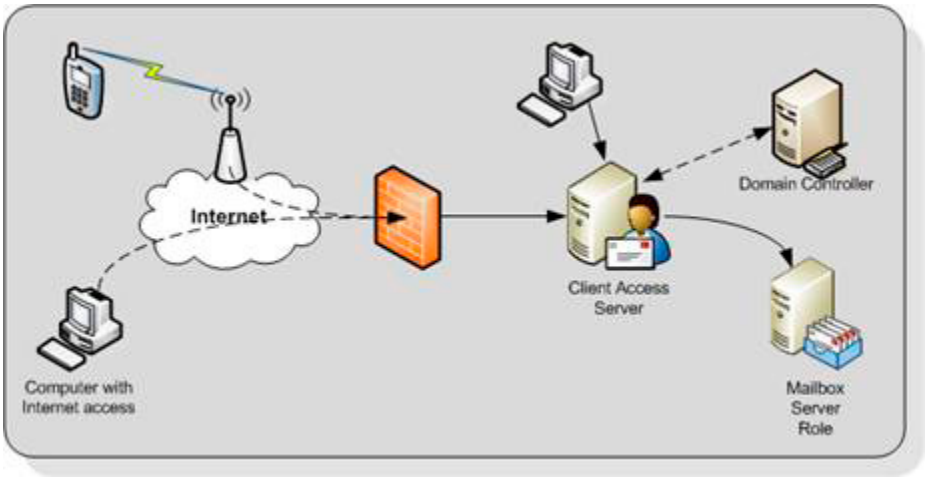

Figure 7. The Client Access Server is responsible for providing access to (Internet) clients. The ISA Server is not in this picture.

Hub Transport Server role

The Hub Transport Server role is responsible for routing messaging not only between the Internet and the Exchange organization, but also between Exchange servers within your organization.

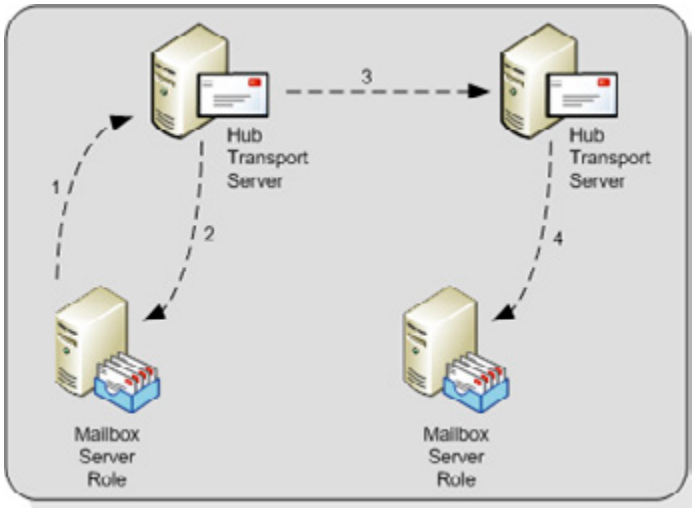

Figure 8. The Hub Transport Server is responsible for routing all messages

All messages are always routed via the Hub Transport Server role, even if the source and the destination mailbox are on the same server and even if the source and the destination mailbox are in the same Mailbox Database. For example in Figure 8:

- Step 1: A message is sent to the Hub Transport Server.
- Step 2: A recipient on the same server as the sender means the message is sent back.
- Step 3: When the recipient is on another mailbox server, the message is routed to the appropriate Hub Transport Server. This is then followed by...
- Step 4: The second Hub Transport Server delivers the message to the Mailbox Server of the recipient.

The reason for routing all messages through the Hub Transport Server is simply compliancy. Using the Hub Transport Server, it is possible to track all messaging flowing through the Exchange organization and to take appropriate action if needed (legal requirements, HIPAA, Sarbanes-Oxley etc.). On the Hub Transport Server the following agents can be configured for compliancy purposes:

- Transport Rule agents – using Transport Rules, all kinds of actions can be applied to messages according to the Rule's filter or conditions. Rules can be applied to internal messages, external messages or both.
- Journaling agents – using the journaling agent, it is possible to save a copy of every message sent or received by a particular recipient.

Since a Mailbox Server does not deliver any messages, every Mailbox Server in an Active Directory site requires a Hub Transport Server in that site. The Hub Transport Server also needs a fast connection to a Global Catalog server for querying Active Directory. This Global Catalog server should be in the same Active Directory site as the Hub Transport Server.

When a message has an external destination, i.e. a recipient on the Internet, the message is sent from the Hub Transport Server to the "outside world." This may be via an Exchange Server 2010 Edge Transport Server in the DMZ, but the Hub Transport Server can also deliver messages directly to the Internet.

Optionally, the Hub Transport Server can be configured to deal with anti-spam and anti-virus functions. The anti-spam services are not enabled on a Hub Transport Server by default, since this service is intended to be run on an Edge Transport Service in the DMZ. Microsoft has supplied a script on every Hub Transport Server that can be used to enable their anti-spam services if necessary.

Anti-virus services can be achieved by installing the Microsoft Forefront for Exchange software. The anti-virus software on the Hub Transport Server will scan inbound and outbound SMTP traffic, whereas anti-virus software on the Mailbox Server will scan the contents of a Mailbox Database, providing a double layer of security.

Edge Server role

The Edge Server role was introduced with Exchange Server 2007, and provides an extra layer of message hygiene. The Edge Transport Server role is typically installed as an SMTP gateway in the network's "Demilitarized Zone" or DMZ. Messages from the Internet are delivered to the Edge Transport Server role and, after anti-spam and anti-virus services, the messages are forwarded to a Hub Transport Server on the internal network.

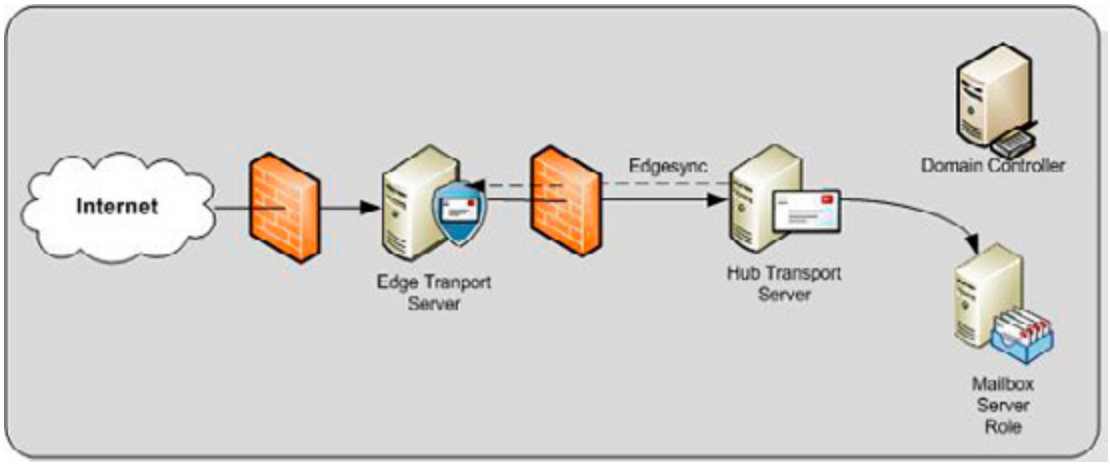

Figure 9. The Edge Transport Server is installed between the Internet and the Hub Transport Server.

The Edge Transport Server can also provide the following services:

- Edge Transport Rules – like the Transport Rules on the Hub Transport Server, these rules can also control the flow of messages that are sent to, or received from the Internet when they meet a certain condition.
- Address rewriting – with address rewriting, the SMTP address of messages sent to or received from the Internet can be changed. This can be useful for hiding internal domains, for example after a merger of two companies, but before one Active Directory and Exchange organization is created.

The Edge Transport Server is installed in the DMZ and cannot be a member of the company's internal Active Directory and Exchange Server 2010 organization. The Edge Transport Server uses the Active Directory Lightweight Directory Services (AD LDS) to store all information. In previous versions of Windows this service was called Active Directory Application Mode (ADAM). Basic information regarding the Exchange infrastructure is stored in the AD LDS, like the recipients and the Hub Transport Server which the Edge Transport Server is sending its messages to.

To keep the AD LDS database up to date, a synchronization feature called EdgeSync is used, which pushes information from the Hub Transport Server to the Edge Transport Server at regular intervals.

Unified Messaging Server role

The Exchange Server 2010 Unified Messaging Server role combines the mailbox database and both voice messages and e-mail messages into one store. Using the Unified Messaging Server role it is possible to access all messages in the mailbox using either a telephone or a computer.

The phone system can be an IP based system or a "classical" analog PBX system, although in the latter case, a special Unified Messaging IP Gateway is needed to connect the two.

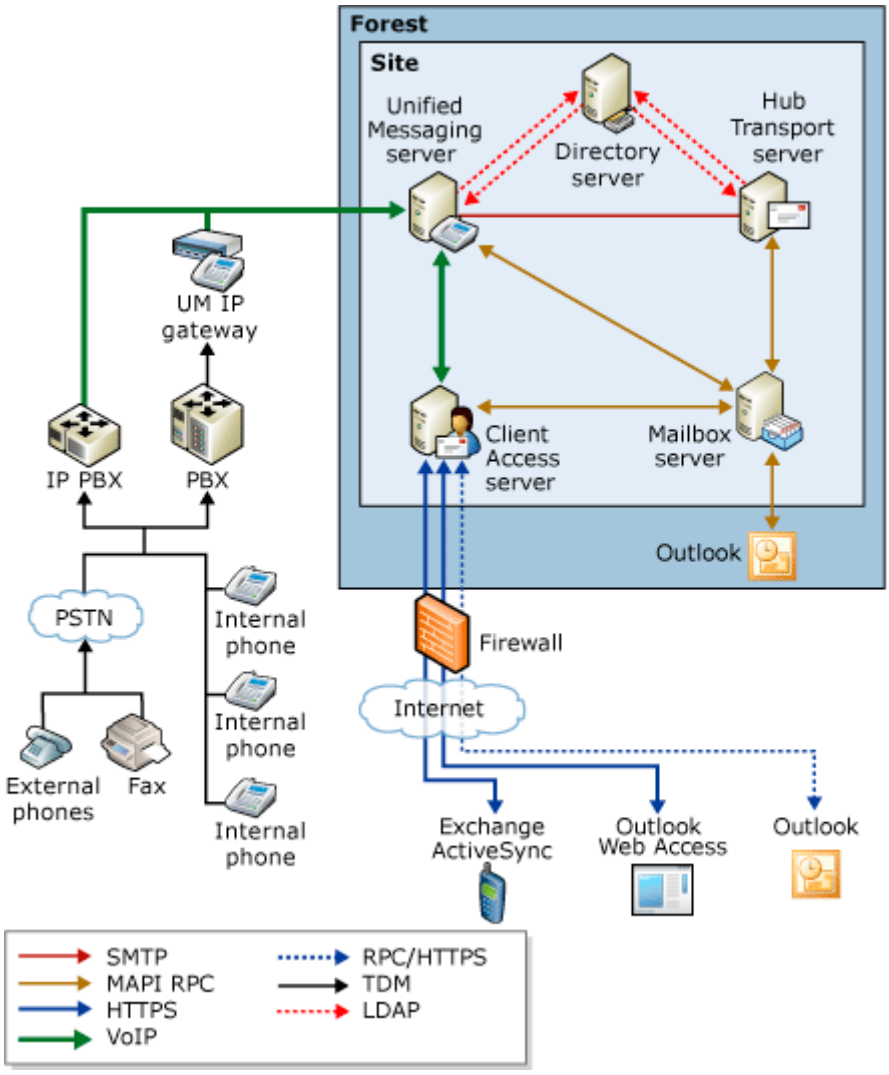

Figure 10. Overview of the Unified Messaging Infrastructure.

The Unified Messaging Server role provides users with the following features:

- Call Answering – this feature acts as an answering machine. When somebody cannot answer the phone, a personal message can be played after which a caller can leave a message. The message will be recorded and sent to the recipient's mailbox as an .mp3 file.
- Subscriber Access – sometimes referred to as "Outlook Voice Access." Using Subscriber Access. users can access their mailbox using a normal phone line and listen to their voicemail messages. It is also possible to access regular mailbox items like messages and calendar items, and even reschedule appointments in the calendar.

- Auto Attendant – using the Auto Attendant, it is possible to create a custom menu in the Unified Messaging system using voice prompts. A caller can use either the telephone keypad or his or her voice to navigate through the menu.

The Unified Messaging service installed on the Unified Messaging Server role works closely with the Microsoft Exchange Speech Engine Service. This Speech Engine Service provides the following services:

- Dual Tone Multi Frequency (DTMF) also referred to as the touchtone (the beeps you hear when dialing a phone number or accessing a menu).
- Automatic Speech Recognition.
- Text-to-Speech service that's responsible for reading mailbox items and reading the voice menus.

The Unified Messaging Server role should be installed in an Active Directory site together with a Hub Transport Server, since this latter server is responsible for routing messaging to the Mailbox Servers. It also should have a fast connection to a Global Catalog server. If possible, the Mailbox Server role should be installed as close as possible to the Unified Messaging Server role, preferably in the same site and with a decent network connection.

Summary

Exchange Server 2010 is the new Messaging and Collaboration platform from Microsoft, and it has a lot of new, compelling features. The new High Availability, management and compliancy features make Exchange Server 2010 a very interesting product for the Exchange administrator. In fact, the new features in Exchange Server 2010 will generally result in *less* complexity, which is always a good thing!

The whole eBook can be downloaded HERE.

The Ego and the System Administrator

22 October 2009

by Matt Simmons

> ***Because the Sys Admins and DBAs who are efficient end up wielding a great deal of power, there is always a danger that it can 'go to their head'. If ever you start imposing your decisions on the organisation for no better reason than the fact that they're yours, then it is time to read Matt's simple guide on how to prevent your actions being controlled by your ego.***

Tim, the Sys Admin, took a moment from flitting back and forth between email, IM, and the ticket queue, in order to assess the state of his network using a screen full of graphs and a highly tuned monitoring system. No aberrant event could occur in the infrastructure without his knowing. Almost as a premonition, he flipped to the virtual terminal that held his email, just as his iPhone chirped an alert. It was a security warning. One of the users had attempted to log into a machine to which they didn't have access.

Tim checked the username, and wasn't surprised. Bill was continually being a user who caused trouble. He installed software on his PC, even after Tim had locked it down. He was one of those users who tried to solve their own problems and, in Tim's opinion, caused more trouble than they were worth.

So why would Bill have been trying to log into the authentication server itself? He should know better than that. Tim reached for the phone to find out. Bill apologized, and tried to explain himself by pointing to an email that he had sent to Tim last week, talking about his login time limits. Bill said that he was only trying to alter the settings on his account so that he could work on the weekend, but Tim suspected that he was trying to get administrative rights on his computer again. Immediately, Tim suspended Bill's account. Now, Bill only had local access, and Tim laid out, in no uncertain terms, what would happen if he caught Bill trying his boundaries again. Bill remembered back to the time he had been suspended from remote access for six months. He apologized again, and then hung up the phone.

Like Tim, we SysAdmins tend to get to the stage where we defend our networks aggressively against the slightest provocation or threat. This isn't a good thing to do, for a number of reasons. It isn't healthy on a personal level; it can have a counter reaction from those we offend that can cause collateral damage to the infrastructure and to other staff members; We lose our reputation as helpful people. We can, and must, respond in a more adaptive way.

The underlying problem is that system administrators have egos. Usually, we have big egos. We play God with the infrastructures that we run, and we are used to being treated as such. When we create a mandate, it goes. When we make a change, it is unquestioned. We say jump, the servers say "how high." We're egotistical because we're the final say in almost every way that counts.

Sysadmins aren't the only profession that is stricken with this particular issue. Surgeons, lawyers, politicians, and online columnists are all frequently seen as egotistical, and although the latter's reputation may be blamed on their editors, there's a kernel of truth to the accusation. Whenever a person is placed in a position of power and given the final say on a matter, the opportunity arises for ego to assert itself.

For the past few weeks, I've been working on an upgrade for my company's telephone infrastructure. We've got a couple of offices, and I'd like to unite them using Voice over IP. The upgrade will involve improving the wiring infrastructure of each office. I was pretty excited to get the chance to improve the physical infrastructure, since it had been necessary for a while. Although it was going to be expensive, it was a step forward.

While I was reviewing my options, one of my users had the nerve to bring a perfectly acceptable option to me that would cut our costs by nearly an order of magnitude. Yes, their solution would work, but it was *their* solution, not mine. Mine was better...because it was mine. What did I say? Well, for a few minutes, I didn't say anything. I wanted to upgrade the network. I wanted to run new lines. I didn't want to use their solution, even if it was better. *Especially* if it was better. My ego got the best of me; At least for a while.

After a few minutes, I was able to step back from the situation and look at their idea logically. It was acceptable, both from a technical and a financial point of view. Although it wouldn't result in the physical infrastructure being upgraded, I couldn't ignore the massive savings and the fact that it would solve the problem at hand. From a purely logical point of view, their solution was better. So what did I do? I told them so. Though it still hurt to do it, I told them that it was better, and I meant it. The only way that this was possible was that I realized that the well being of the company was more important than my ego being satisfied by solving the problem.

It's hard to avoid being controlled by one's ego, and it takes practice. It is worth it because well-being of the company is your real goal, and it can best be served by doing whatever is right. Remember that whoever is presenting an opposing idea isn't your enemy. They want the same thing that you do, or at least, you should treat their request as if that's the case.

Implementing Windows Server 2008 File System Quotas

19 November 2009

by Ben Lye

> ***File system Quotas are used to restrict the amount of space users can consume or to report on the space consumed by them. They are useful for reporting on those users or folders that are consuming large amounts of disk space on a file server. Ben Lye shows that File system quotas are quick and easy to set up, with three different methods available for configuring them.***

Disk quotas have been available in Windows since Windows 2000 was released, and could be used by administrators to limit the amount of space users could use on an NTFS volume. Disk quotas are based on file ownership rather than folder structure and because of this they are not particularly useful in all situations. For example, if your server had a single storage volume and you need to apply quotas to different folders on the volume then disk quotas will not help.

File system quotas, which were first introduced in Windows Server 2003 R2, and are a part of the File Server role in Windows Server 2008 (and Windows Server 2008 R2), offer many benefits over disk quotas. With file system quotas we can set quotas for specific folders on the volume, we can use templates to ensure consistent application of quotas, and we can set quotas which are automatically applied to all sub-folders of a folder.

Additionally, file system quotas are useful not just for limiting the amount of space users can consume, but also for reporting on space used – quotas can be set with so-called "soft" limits which are used for monitoring rather than enforcing limits. This functionality can be extremely useful for quickly determining which users or folders are consuming large amounts of disk space on a file server.

Quota thresholds can be configured so that users or administrators receive notifications when quotas have been reached or are about to be reached. Multiple thresholds can be configured for individual quotas, and actions can include sending e-mail messages, logging to the Windows event log, running commands or scripts, or generating storage reports.

In Windows Server 2008 file system quotas are managed with the File Server Resource Manager (FSRM) console (which is installed as a role service in the File Services role), the command line utility dirquota, or with Windows PowerShell using a COM API.

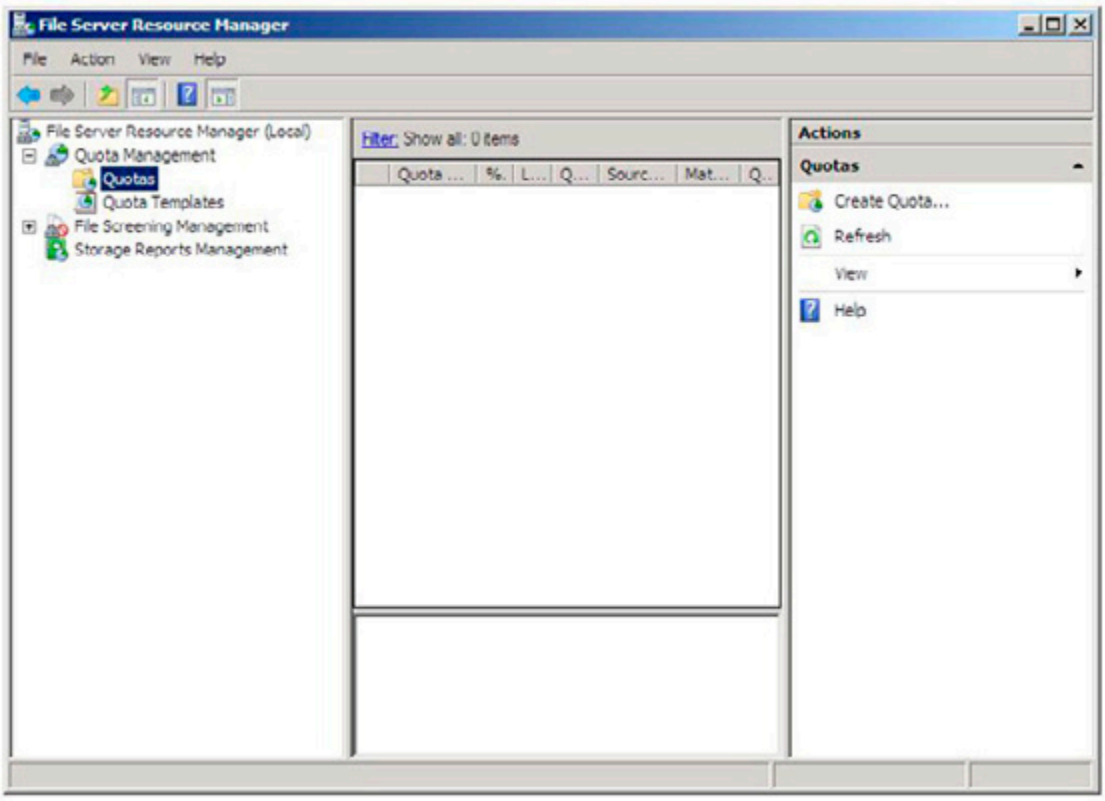

Figure 1: Windows Server 2008 File Server Resource Manager.

There are two kinds of quota available - hard quotas which set a limit and enforce it, and soft quotas which set a limit but only report on it. Soft quotas are useful for monitoring disk space use. Quotas are commonly applied using quota templates, which are a mechanism for easily applying the same quota settings to one or more folders. Quota Templates are the recommended way to configure quotas and FSRM includes some example templates which cover a range of scenarios, including using both hard and soft quota types.

Before we start to configure quotas which will generate e-mail messages, the quota File Server Resource Manager needs to be configured with an SMTP server, and optionally, the default administrator recipients, and the default "from" address.

Like all aspects of quota management, the FSRM settings can be applied using three different tools and you can choose the method appropriate to your needs.

To configure FSRM using the FSRM console:

- Launch the File Server Resource Manager.
- Select the root node, labelled "File Server Resource Manager."
- In the Action Pane click "Configure Options..."
- Enter an SMTP server and if desired configure the other settings.

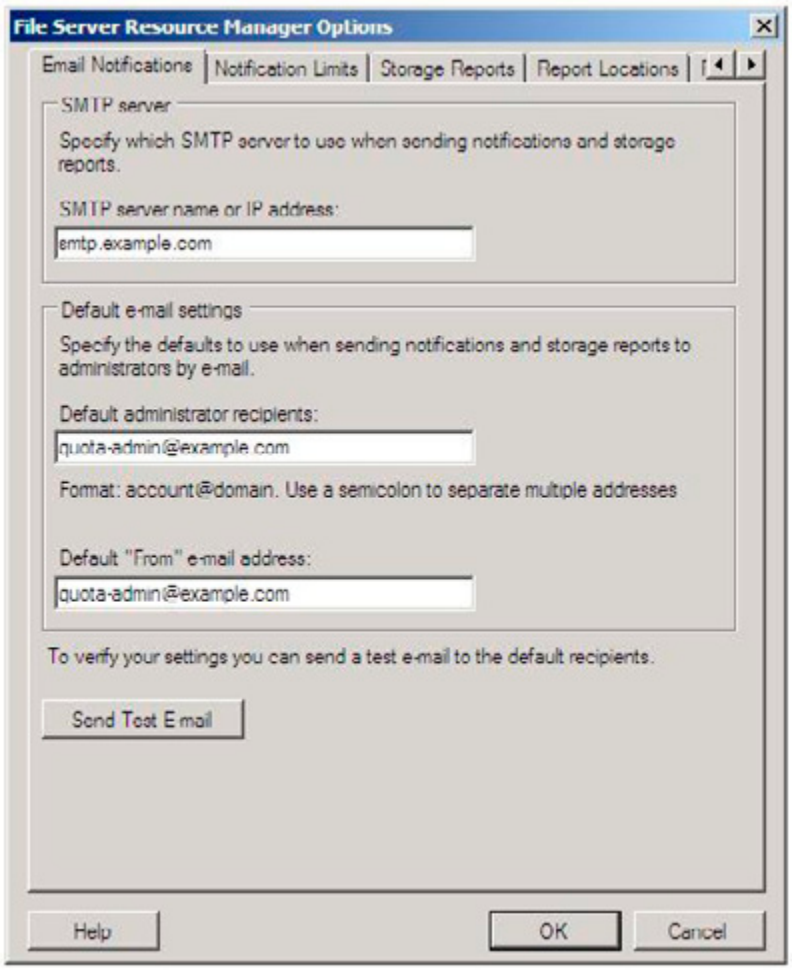

- Click the "OK" button.

To configure FSRM using the command line:

- Open an elevated command prompt window.
- Enter the command "dirquota admin options /From:quota-admin@example.com /AdminEmails:quota-admin@example.com /SMTP:smtp.example.com."

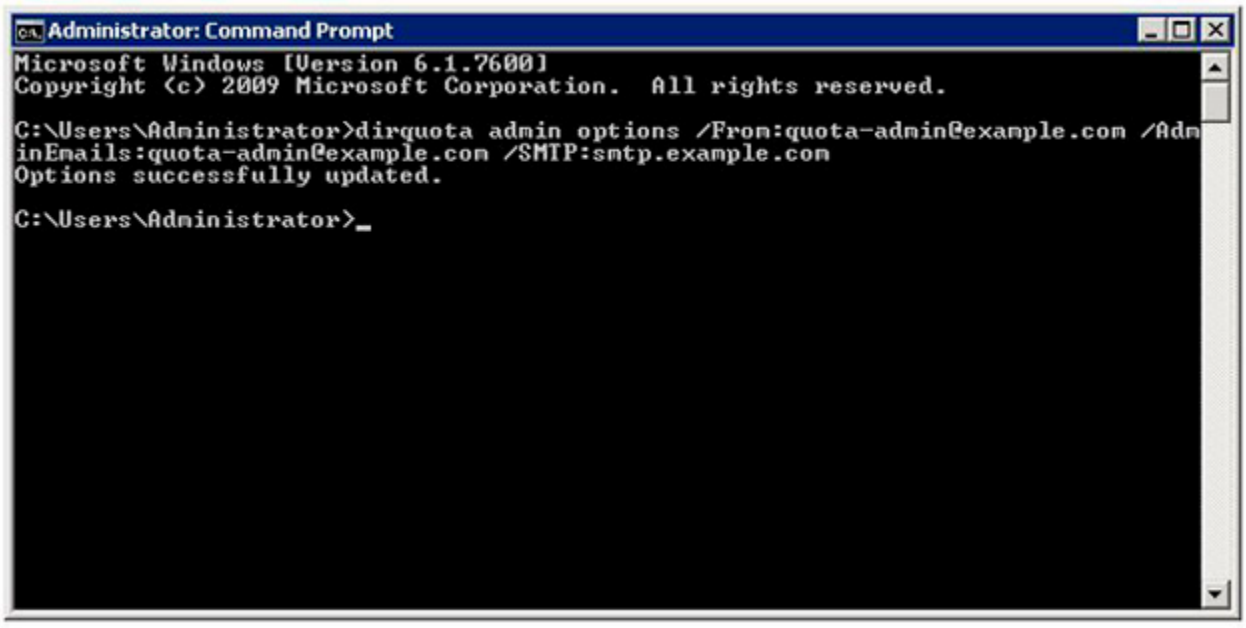

To configure FSRM using Windows PowerShell:

- Open Windows PowerShell.
- Enter these commands (or save them as a script and run it):

```
# Create a new COM object to access the FSRM settings
$fsrm = New-Object -com Fsrm.FsrmSetting

# Set the default administrators e-mail address
$fsrm.AdminEmail = "quota-admin@example.com"

# Set the from address
$fsrm.MailFrom = "quota-admin@example.com"

# Set the SMTP server
$fsrm.SmtpServer = "smtp.example.com"
```

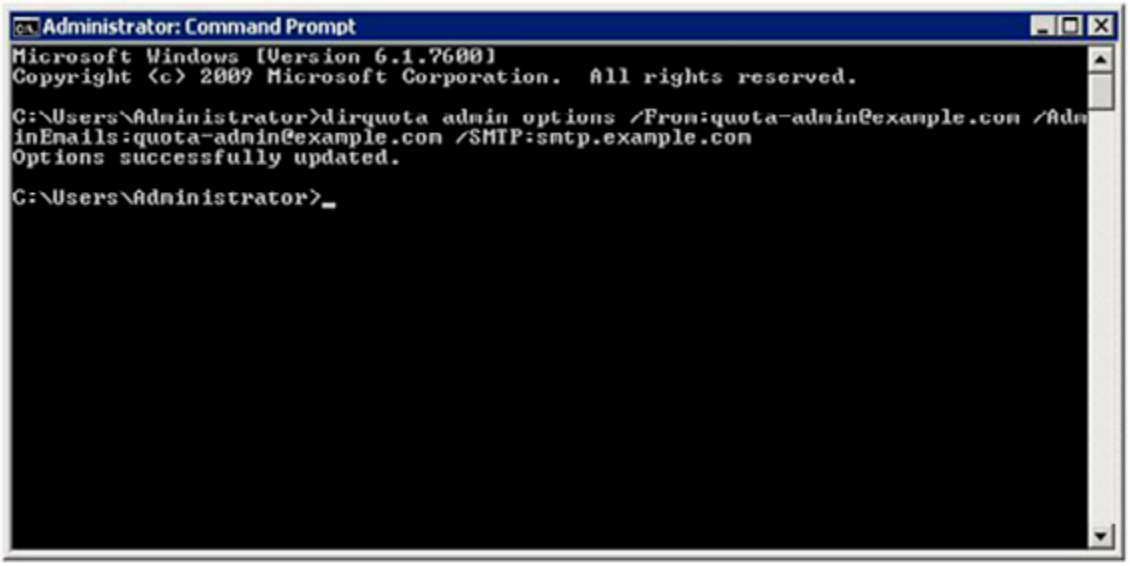

Quota Example – Home directories with a 5GB limit

A common use of file system quotas is to put limits on the size of user's personal storage space (folders which are often referred to as home directories) on a file server. The requirements of this scenario are to limit the space each user can use to 5GB, alert administrators when 90% of the quota has been reached, and automatically apply quotas to new home directories. The solution requires the implementation of new quota template and an auto apply quota.

Step 1: Create the new quota template

The first step is to create a new template, which we will use later to apply the quota to the file system. Using a template means we can easily make changes to all folders where we have applied the template quota settings. The template can be created using the FSRM, the **dirquota** command line tool, or PowerShell, meaning you can choose the tool with which you feel comfortable with and that fits most of your scenarios.

To create the new quota template using the FSRM:

- Launch the File Server Resource Manager.
- Expand "Quota Management" ▸ "Quota Templates".
- In the Action Pane click "Create Quota Template".
- Enter the template name and set the space limit.

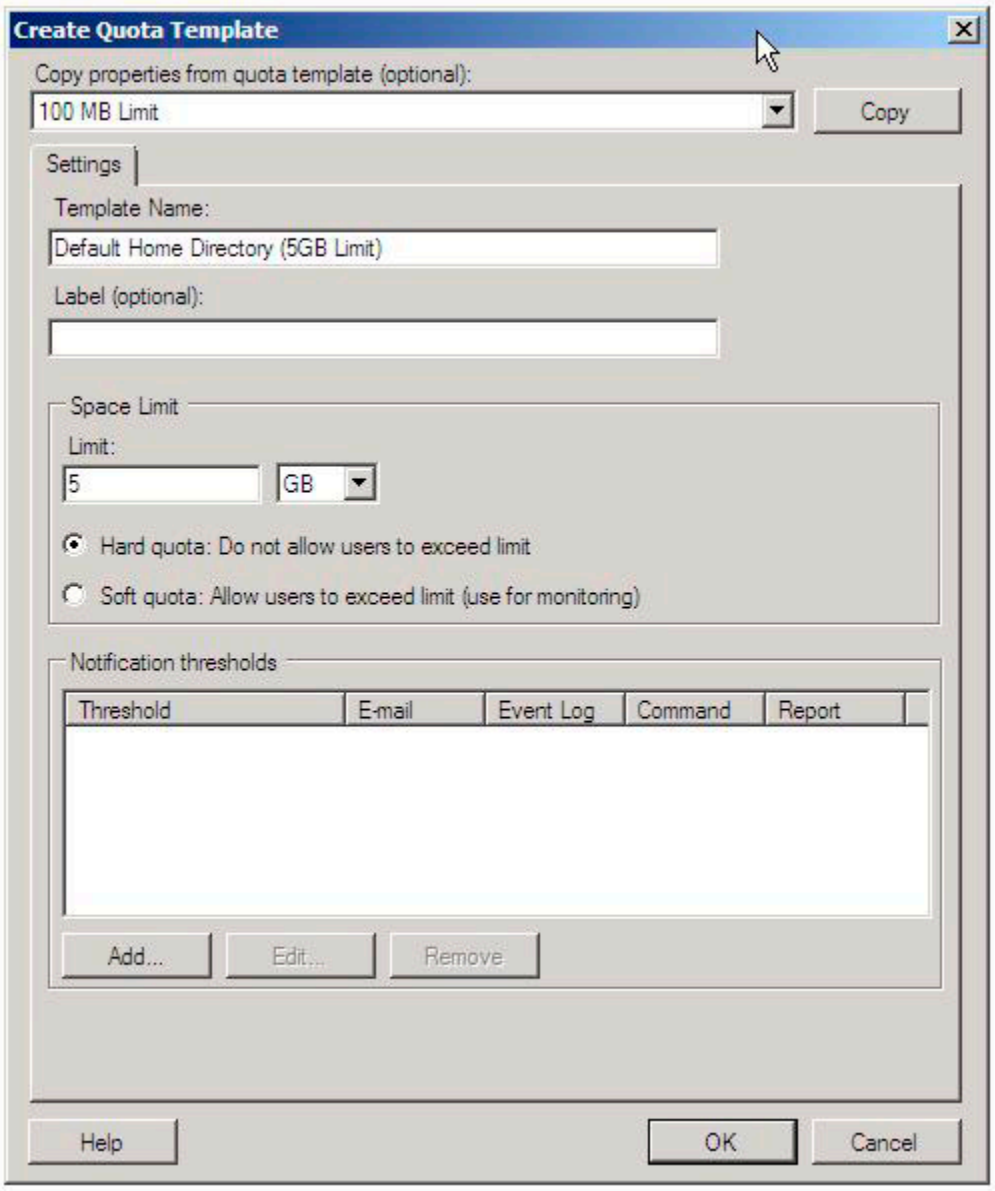

Note

To set a soft quota (for monitoring only) check the "Soft Quota" radio button.

- Click the "Add" button to add a notification threshold.
- Set the notification percentage to 90, check the "Send e-mail to the following administrators," and enter an appropriate destination e-mail address. You can also customise the message text.

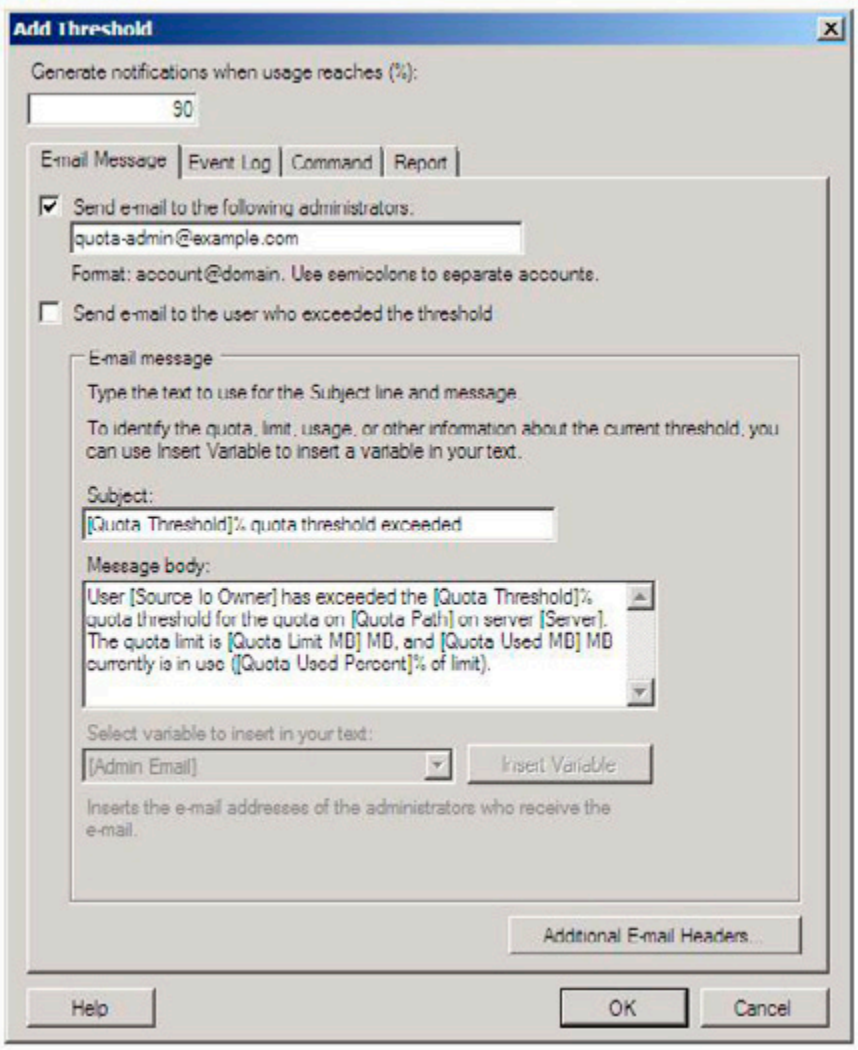

- Click the "OK" button twice.

To create the new quota template using the **dirquota** command line utility:

- Open an elevated command prompt (or Windows PowerShell) window.
- Create a text file called notification.txt containing the text of the notification message (an example of this text message can be downloaded from this article).
- Enter this command:

```
dirquota Template Add /Template:"Default Home Directory (5GB Limit)" /Limit:5GB /Type:Hard /Add-Threshold:90 /Add-Notification:90,M,notification.txt
```

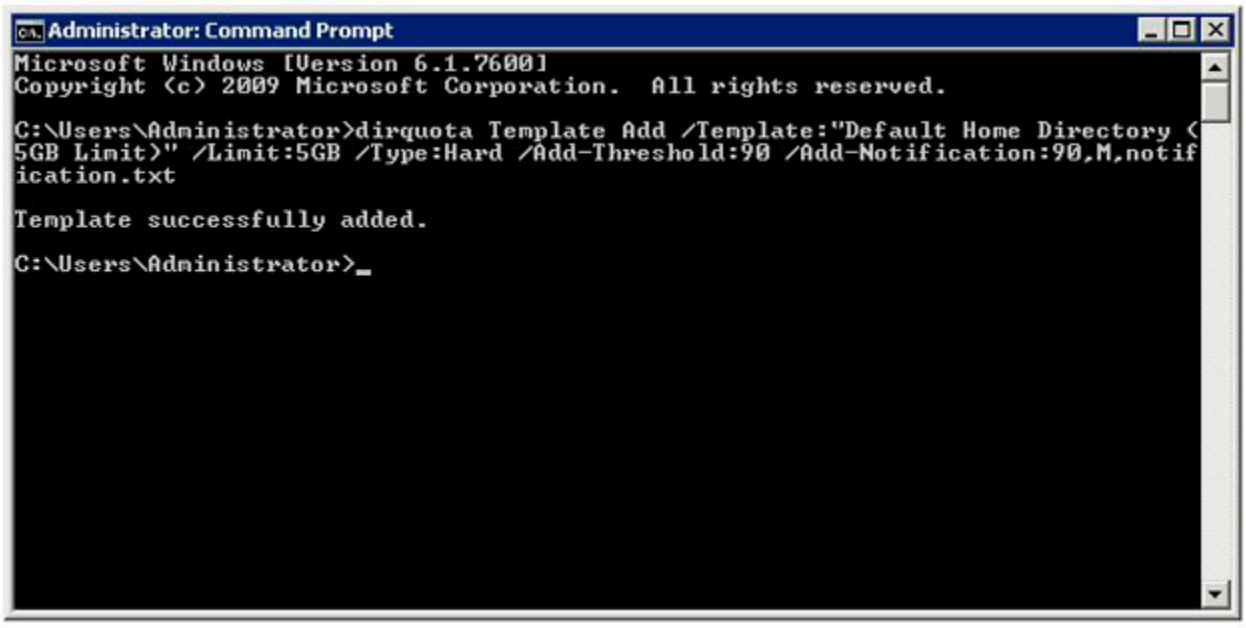

Note: To set a soft quota (for monitoring only) change "/Type:Hard" to "/Type:Hard"

To create the new quota using Windows PowerShell:

- Open Windows PowerShell.
- Enter these commands (or save them as a script and run it):

```
# Create a new COM object to access quota templates
$fqtm = New-Object -com Fsrm.FsrmQuotaTemplateManager

# Create a new template object
$template = $fqtm.CreateTemplate()

# Set the template's name
$template.Name = "Default Home Directory (5GB Limit)"

# Set the quota limit
$template.QuotaLimit = 5GB

# Set the quota type to hard limit (the flag for a hard limit is 0x100)
$template.QuotaFlags = $template.QuotaFlags -bor 0x100

# Add a quota threshold
$template.AddThreshold(90)

# Add a threshold e-mail action
$action = $template.CreateThresholdAction(90, 2)

# Set the e-mail message recipient
$action.MailTo = "[Admin Email]"

# Set the e-mail message subject
```

```
$action.MailSubject = "[Quota Threshold]% quota threshold exceeded"

# Set the e-mail message text
$action.MessageText = "User [Source Io Owner] has exceeded the [Quota Threshold]% " + `
"quota threshold for the quota on [Quota Path] on server [Server]. The quota limit " + `
"is [Quota Limit MB] MB, and " + ` "[Quota Used MB] MB currently is in use ([Quota " + `
"Used Percent]% of limit)."
```

Note

To set a soft quota (for monitoring only) change "$template.QuotaFlags = $template.QuotaFlags -bor 0x100" to "$template.QuotaFlags = $template.QuotaFlags -bxor 0x100" to disable the hard limit flag.

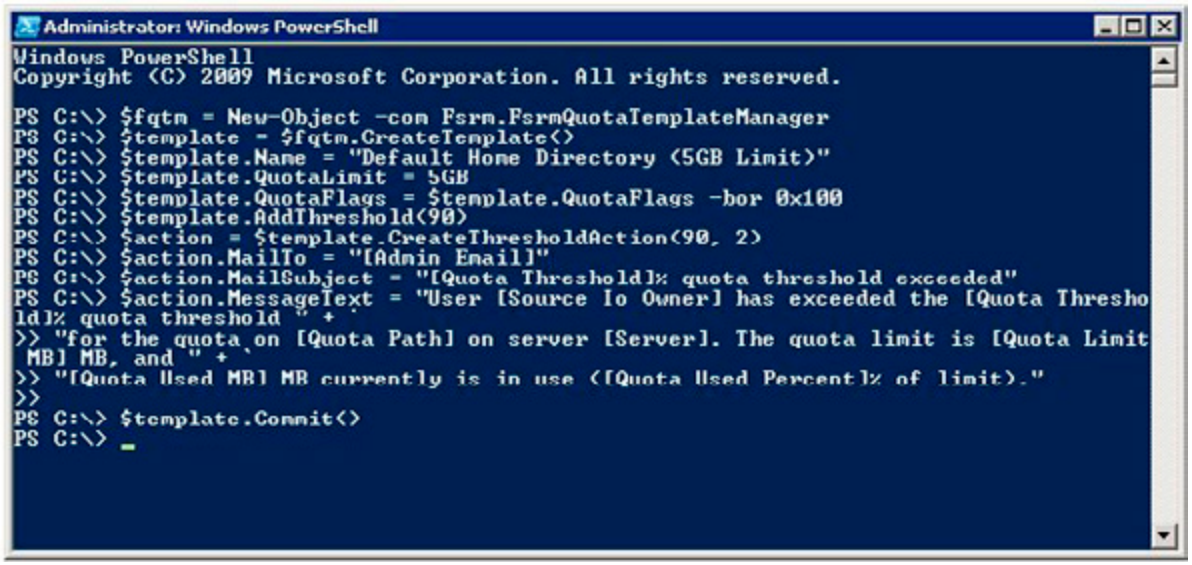

Step 2: Create the Quota

The next step is to use the new quota template to apply the quota to the file system.

In this example we'll say that the home directories are all subfolders of C:\Home. Because we want any new home folders to automatically have the quota applied we need to create an Auto apply Quota. Auto apply quotas are applied to all existing subfolders and any future folders.

To create the quota using the FSRM:

- Launch the File Server Resource Manager.
- Expand "Quota Management" ▸ "Quotas".
- In the Action Pane click "Create Quota".
- Enter the quota path and choose the appropriate template.

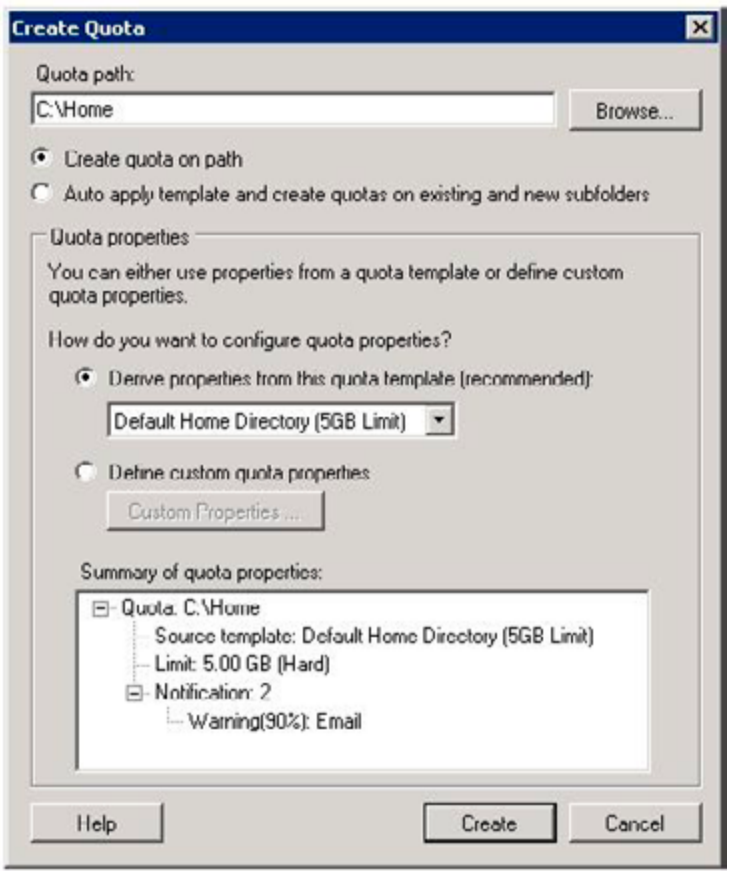

- Click the "Create" button.

To create the quota using the **dirquota** command line tool:

- Open an elevated command prompt window.
- Create a text file called notification.txt containing the text of the notification message (an example of this text message can be downloaded from this article).
- Enter the command "dirquota autoquota add /Path:C:\Home /SourceTemplate:"Default Home Directory (5GB Limit)".

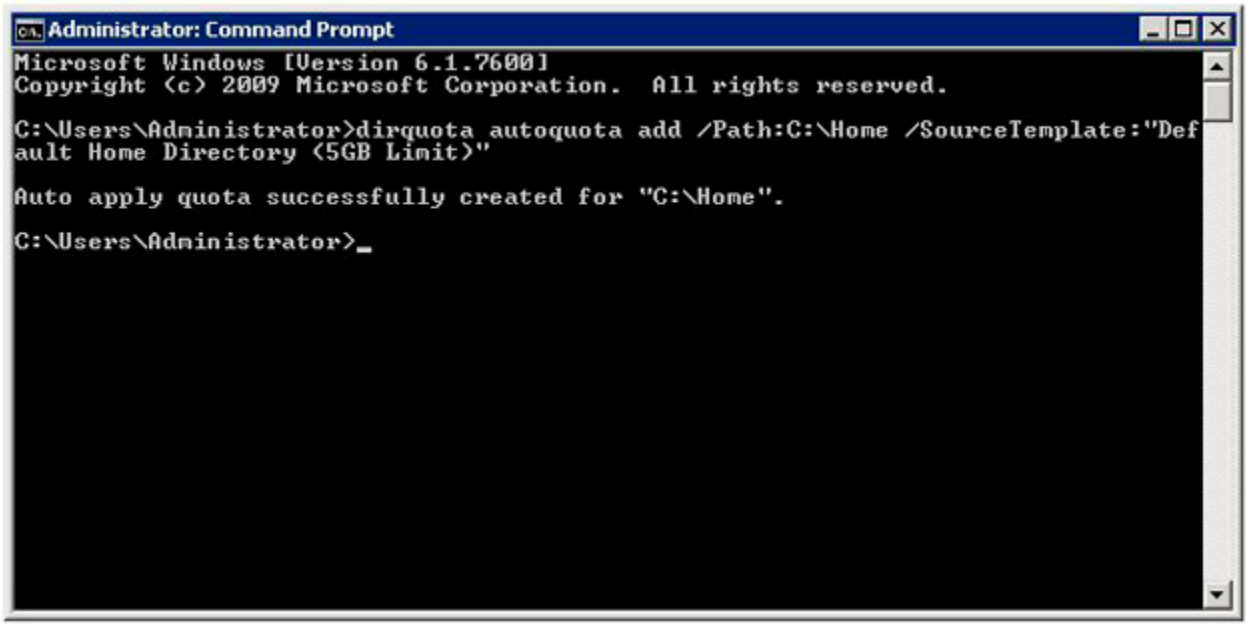

To create the quota using Windows PowerShell:

- Open Windows PowerShell
- Enter these commands (or save them as a script and run it):

```
# Create a new COM object to access quotas
$fqtm = New-Object -com Fsrm.FsrmQuotaManager

# Create the new quota
$quota = $fqtm.CreateAutoApplyQuota("Default Home Directory (5GB Limit)," "C:\Home")

# Save the new quota
$quota.Commit()
```

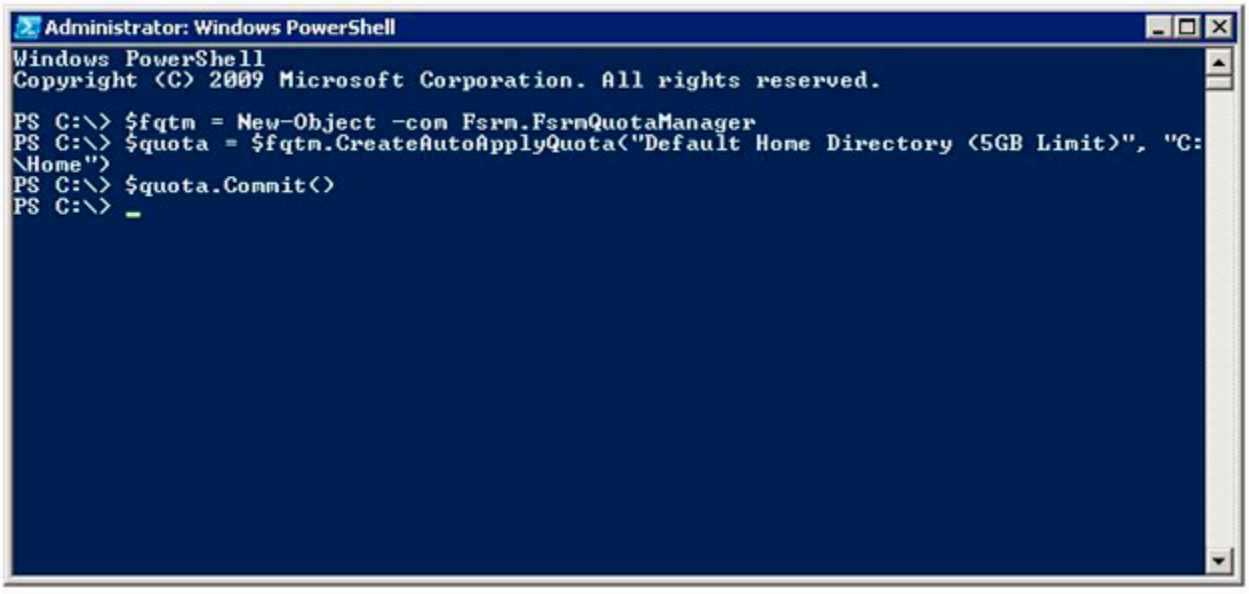

Quota Exceptions / Folder-Specific Quotas

Naturally there will be occasions when a folder needs to be excluded from a template or auto apply quota. In these situations you can easily add a specific quota for that folder to either increase the limit or to disable the quota entirely.

To create the quota exception using the FSRM:

- Launch the File Server Resource Manager.
- Expand "Quota Management" ▸ "Quotas".
- Select the folder you wish to make the exception for.
- In the Action Pane click "Edit Quota Properties...".
- Enter new limit for the quota.

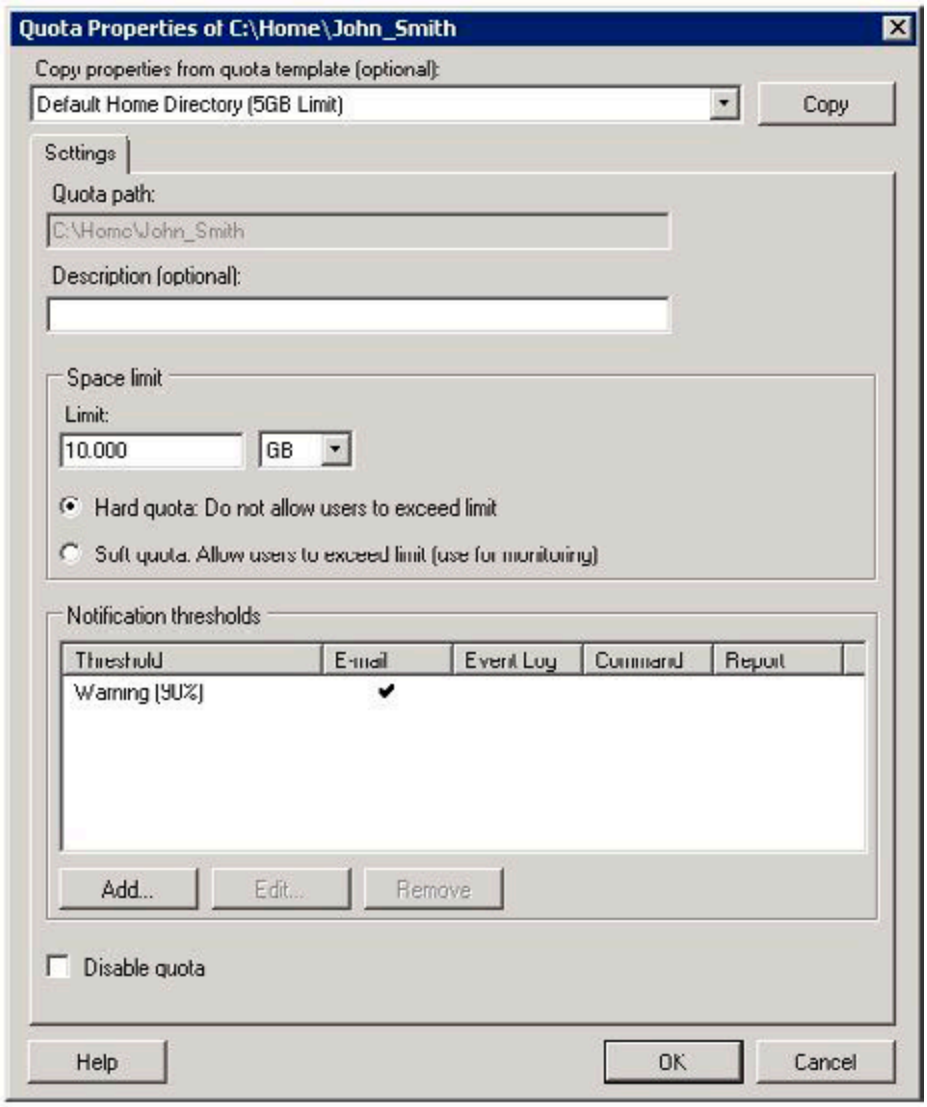

- Click "OK."

Note

Check the "Disable quota" box to disable the quota.

To create the quota exception using the **dirquota** command line tool:

- Open an elevated command prompt window.

Enter the command:

```
"dirquota quota add /Path:C:\Home\John_Smith /SourceTemplate:"Default Home Directory (5GB Limit)"
/Limit:10GB /Overwrite"
```

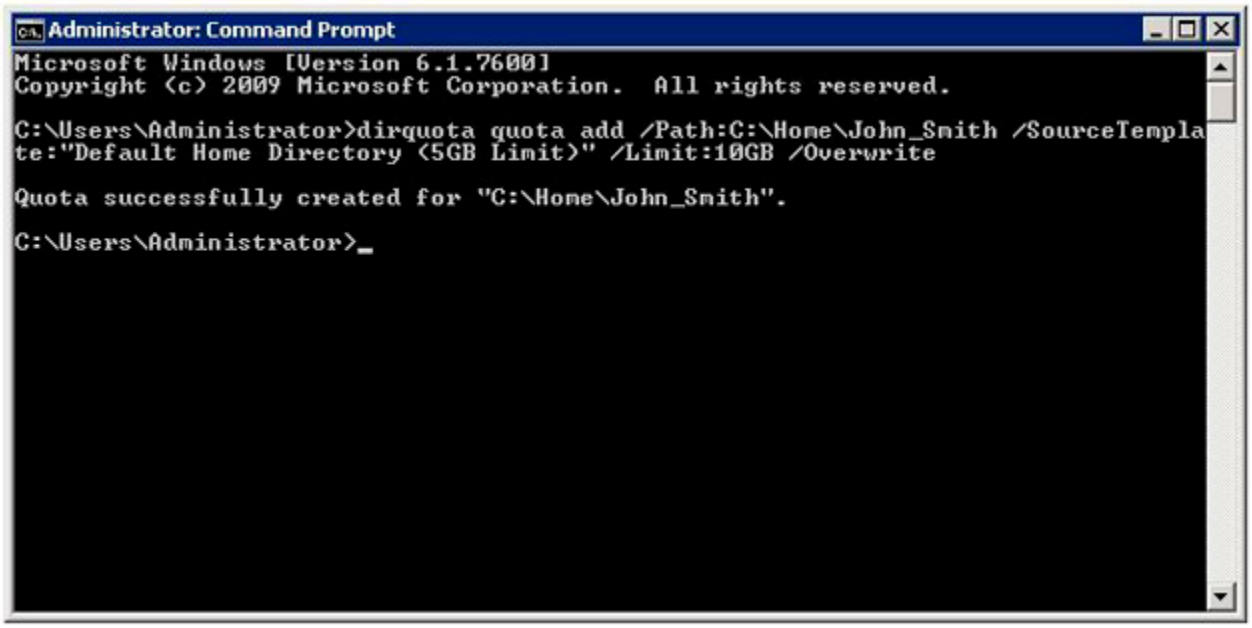

Note

To disable the quota append the command with "/status:disabled."

To create the quota exception using Windows PowerShell:

- Open Windows PowerShell.
- Enter these commands (or save them as a script and run it):

```
# Create a new COM object to access quotas
$fqtm = New-Object -com Fsrm.FsrmQuotaManager

# Get the existing quota
$quota = $fqtm.GetQuota("C:\Home\John_Smith")

# Set the new quota limit
$quota.QuotaLimit = 10GB

# Save the quota
$quota.Commit()
```

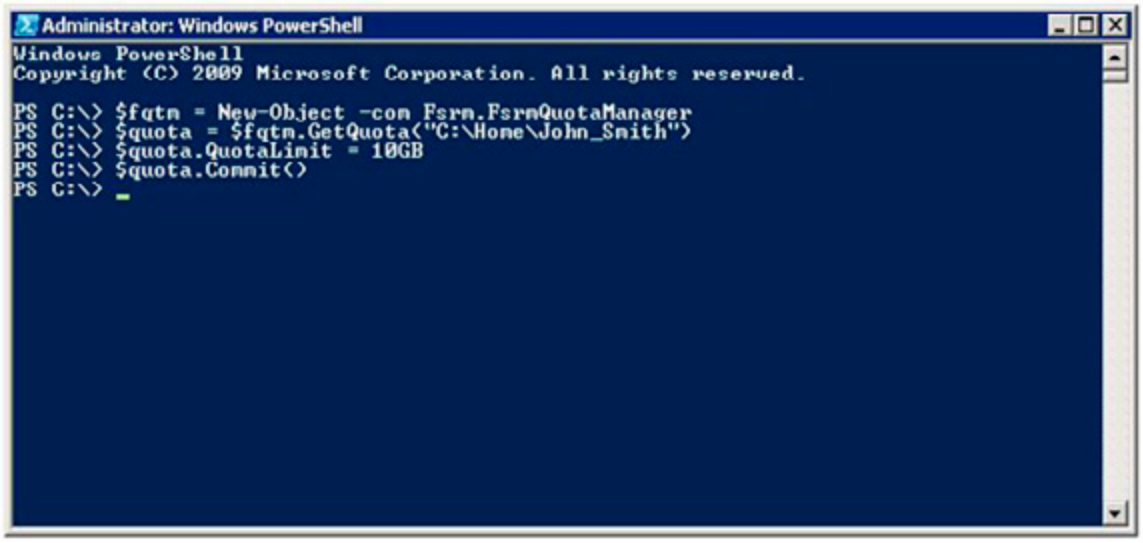

Note

To disable the quota, insert the line "$quota.QuotaFlags = $quota.QuotaFlags -bor 0x200" before saving the quota.

How Quotas Affect Clients

When a client maps a network drive to a folder which has a hard quota applied, the size of the volume and the amount of available disk space shown is equal to the quota settings.

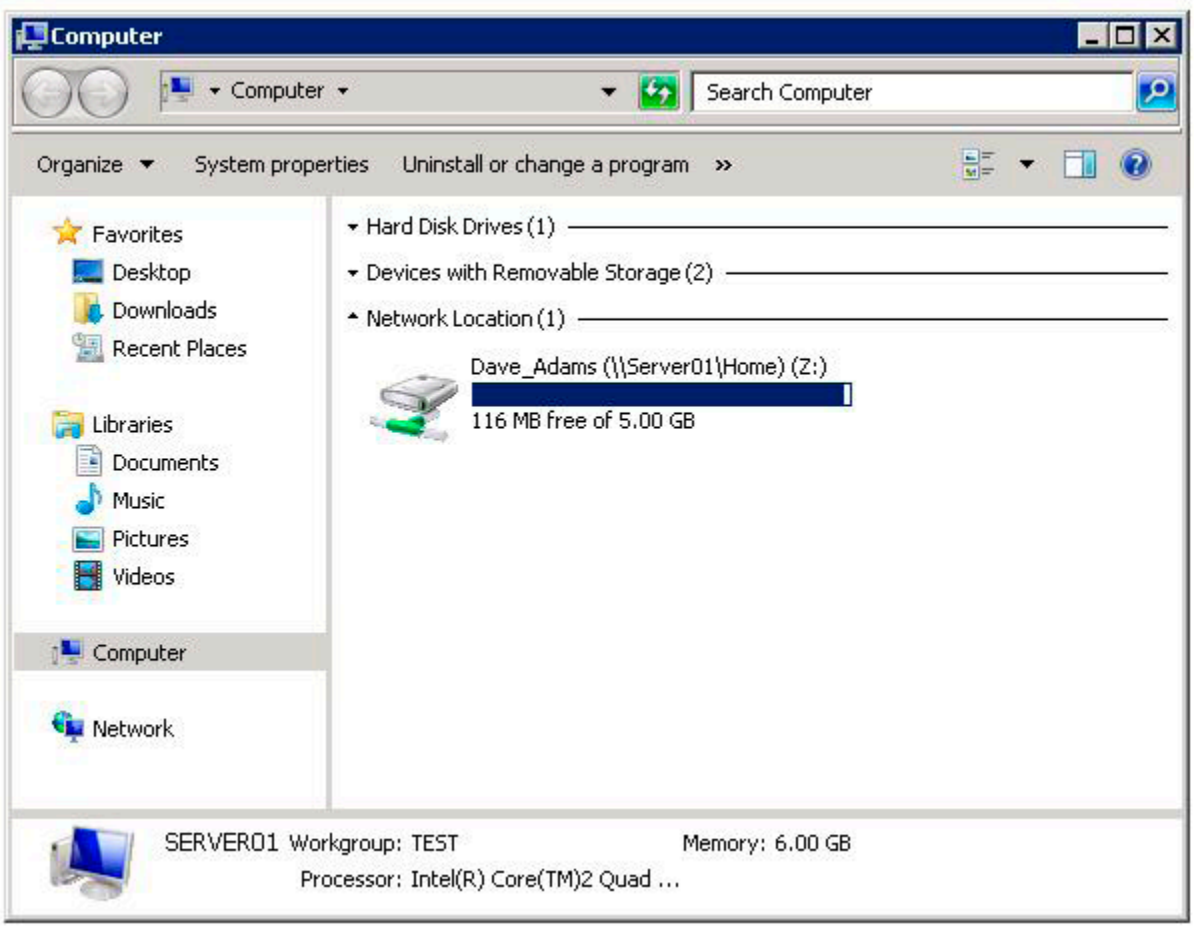

Figure: Mapped drive showing hard quota limit as volume size.

When a hard quota is met or exceeded clients will receive a message telling them that the volume is out of disk space.

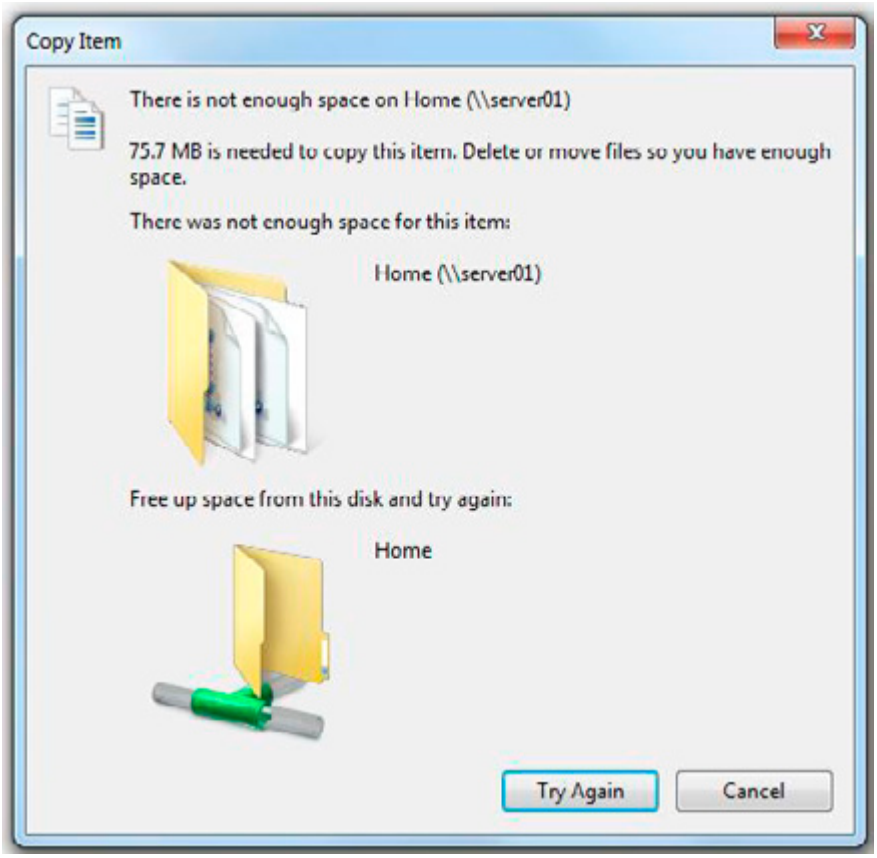

Figure: Insufficient disk space message.

Viewing Quotas

Administrators can view hard and soft quotas using FSRM, and viewing quotas this way can be a quick method for finding large folders or large consumers of space.

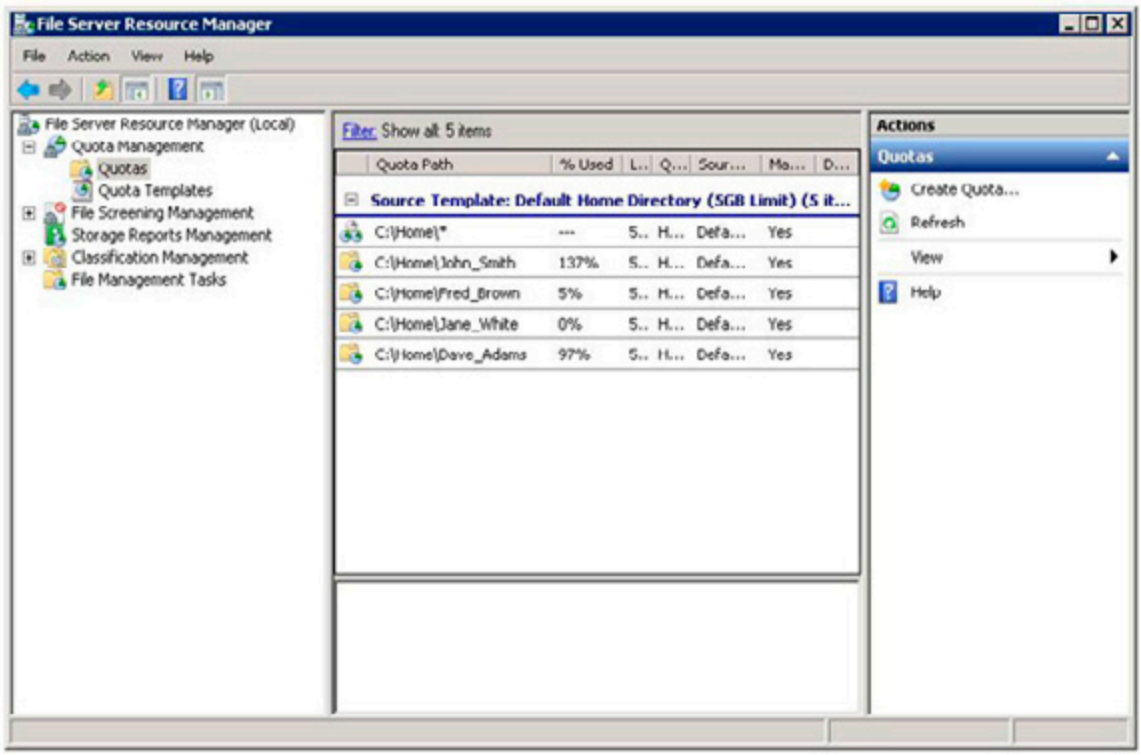

Figure: Quotas in File Server Resource Manager.

File system quotas are quick and easy to set up, with three different and flexible methods available for configuring them. Properly applied they can be a good tool to help ensure efficient use of storage resources, a convenient countermeasure against storage waste, and a useful tool for reporting on storage utilisation. There is an IO performance penalty for using quotas, but the benefits will probably outweigh the small performance cost.

More information about File Server Resource Manager is available in Microsoft TechNet:
http://technet.microsoft.com/en-us/library/cc771092%28WS.10%29.aspx.

More information about the **dirquota** command line tool is also available in TechNet:
http://technet.microsoft.com/en-us/library/cc754836%28WS.10%29.aspx.

More information about the COM API for working with FSRM is available in MSDN:
http://msdn.microsoft.com/en-us/library/bb972746%28VS.85%29.aspx.

Implementing Cluster Continuous Replication, Part 2

19 November 2009

by Brien Posey

Cluster continuous replication (CCR) helps to provide a more resilient email system with faster recovery. It was introduced in Microsoft Exchange Server 2007 and uses log shipping and failover. configuring Cluster Continuous Replication on a Windows Server 2008 requires different techniques to Windows Server 2003. Brien Posey explains all.

In the first part of this article series, I showed you how to configure Windows to form a Majority Node Set Cluster. Although the procedure that I gave you works well, it is intended for use with Windows Server 2003. You can also host a clustered mailbox server on Windows Server 2008, but the procedure for creating the cluster is quite a bit different from what you saw in the first article. In this article, I want to show you how to create the required cluster on Windows Server 2008.

In the previous article, I spent a long time talking about the hardware requirements for building a cluster, and about the cluster planning process. This same information roughly applies to planning for a Windows Server 2008 cluster. You are still going to need two similar servers with two NICs installed. The requirements for the number of names and IP addresses that I covered in Part 1 also remain the same. Once again, I will be referring to the cluster node that is initially active as Node1, and the node that initially acts as the passive node as Node 2.

Deploying the Failover Cluster Feature

Before you can configure Windows 2008 Servers to act as clusters, you must install the Failover Cluster feature. To do so, begin by opening the Server Manager, and selecting the Features container. Click the Add Features link, and Windows will display a list of the various features that you can install. Select the Failover Clustering check box, as shown in Figure A, and click Next. At this point, you will see a warning message telling you that a server restart will be required. You don't have to do anything to acknowledge the message though. Just click the Install button to install the feature. When the installation process completes, click Close and reboot your server if necessary. Now, perform this procedure on your second cluster node.

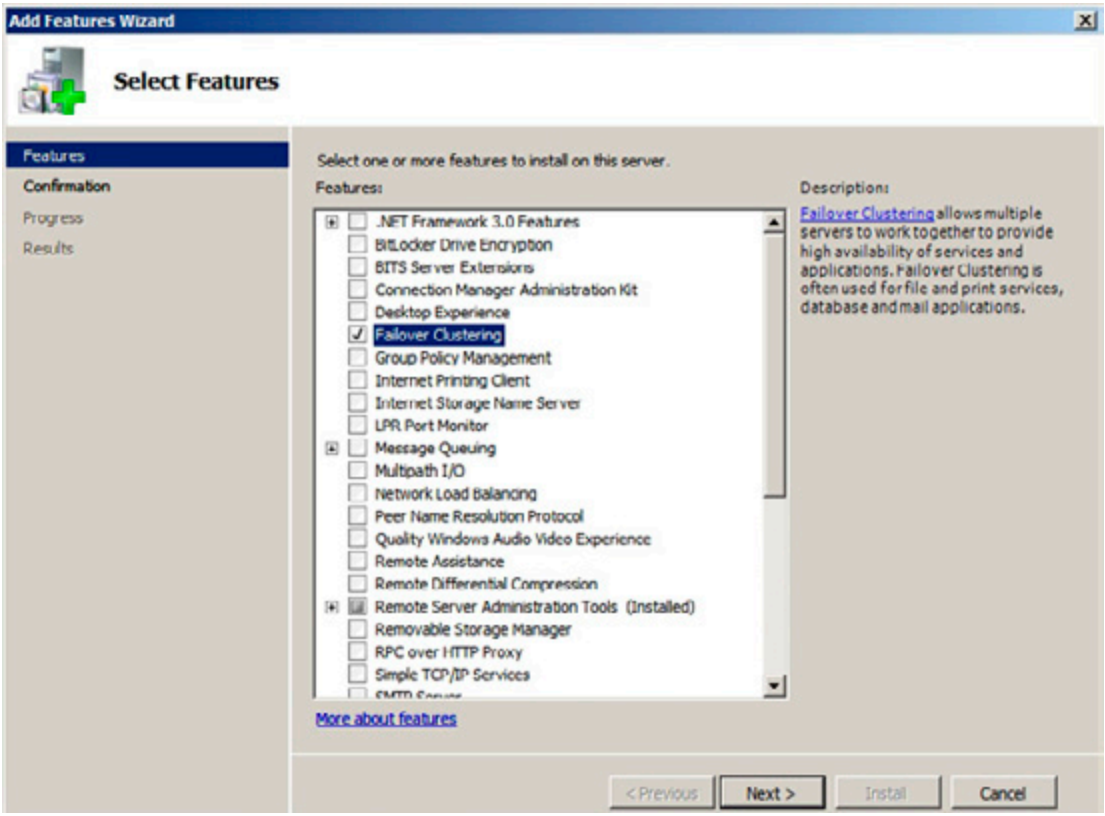

Figure A.

Select the Failover Cluster option, and click Next.

The Failover Cluster Management Console

When we set up clustering in Windows Server 2003, we used a wizard for the configuration process, but we had to call the wizard from the command line via the Cluster command. Although the Cluster command still exists in Windows Server 2008, the individual commands that I showed you in Part 1 for invoking the cluster wizard no longer work. The good news is that you no longer have to launch the wizard from the command line. In Windows Server 2008, Microsoft gives us an administrative tool that you can use to perform various clustering functions. You can access this tool by selecting the Failover Cluster Management command from the server's Administrative Tools menu.

Creating the Cluster

You are now ready to create our cluster. As before, we are going to be creating an active and a passive node, which I will be referring to as Node 1 and Node 2. We can get started by opening the Failover Cluster Management Console on Node 1. When the console opens, click the Create Cluster link, found in the console's Management section. When you do, Windows will launch the Create Cluster Wizard.

Click Next to clear the wizard's Welcome dialog, and the wizard will prompt you to enter the names of the cluster nodes. Enter the name of Node 1, and click the Add button. Once Windows validates the server's name and adds it to the list, repeat the process by entering the name of Node 2 and clicking the Add button. When both cluster nodes have been added to the list, click Next.

You should now see a dialog warning you that Microsoft does not support the use of clusters unless the cluster configuration can be validated. You are now presented with the option of running a validation wizard. Choose the option labeled Yes, When I Click Next, Run Configuration Validation Tests, and Then Return to the Process of Creating the Cluster

Click Next, and the wizard will take you to an introductory dialog, as shown in Figure B,that describes the validation process. Go ahead and click Next to move on. You are now asked if you want to run all of the validation tests, or if you want to select specific tests to run. Choose the option to run all of the tests and click Next. You should now see a screen that confirms the tests that you are about to perform. Click Next to begin the validation tests. When the validation tests complete, you may receive some warnings (particularly in regard to storage), but you should not receive any errors. Click Finish.

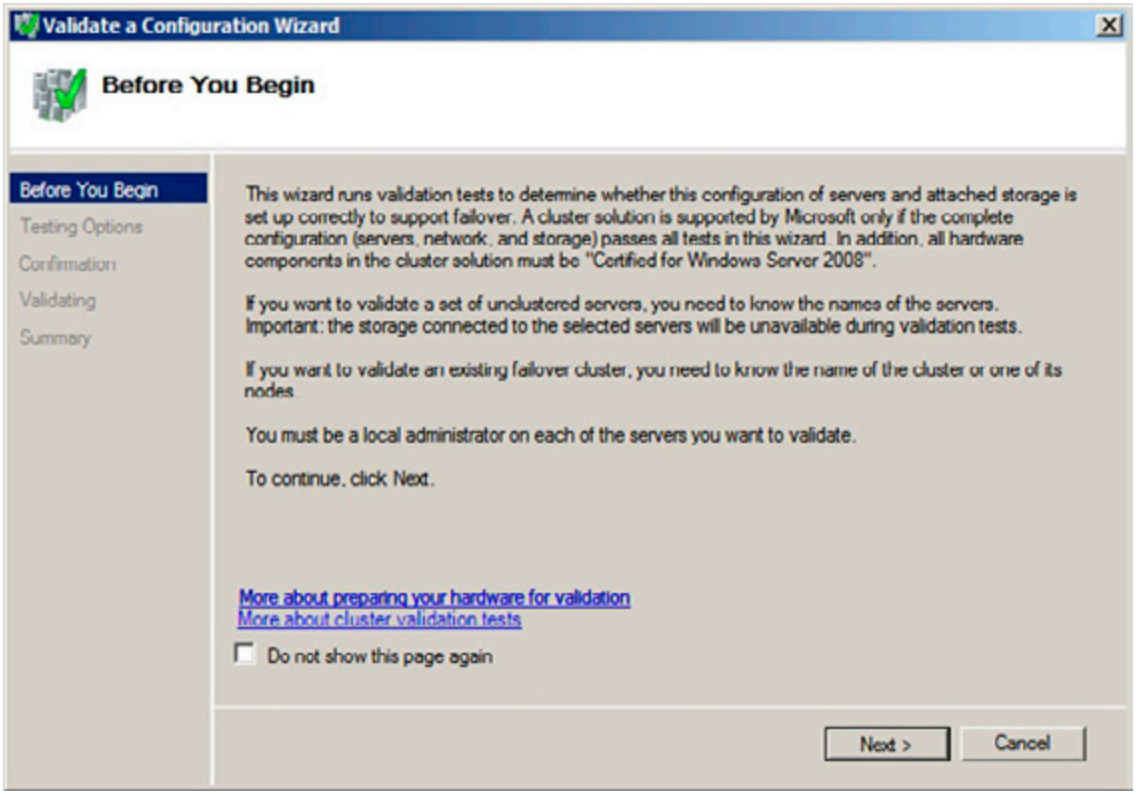

Figure B.

This dialog introduces you to the validation tests.

Now that the validation is complete, Windows returns us to the Create Cluster Wizard. At this point, you must enter the name and IP address that you wish to assign to the cluster, as shown in Figure C. Be sure to make note of which names and addresses you use, because you will need to reference them again later on.

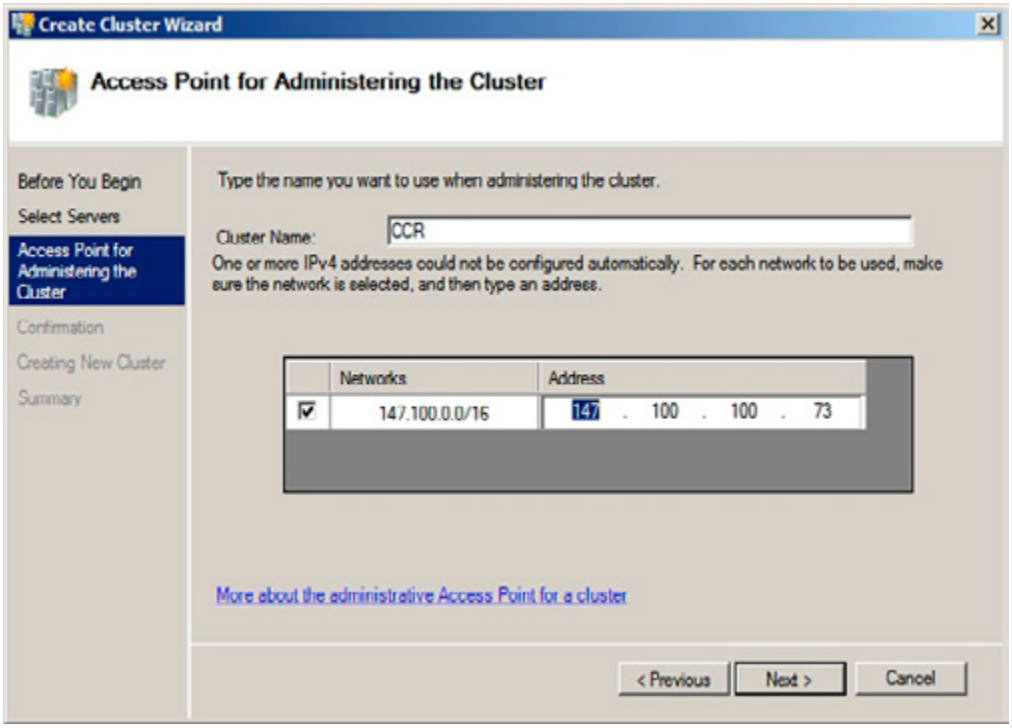

Figure C.

Enter the name and IP address that you want to assign to the cluster.

After entering the cluster's name and IP address, click Next, and you should see a summary screen outlining the details of the cluster that you are creating. Take a moment to read the summary and make sure that everything is correct. Assuming that all is well, click Next, and Windows will create your cluster. When the process completes, click Finish.

Configuring Networking for the Cluster

Now that you have finished creating the cluster, you must make sure that the network adapters are configured properly for the cluster. Windows Server 2008 normally does a pretty good job of automatically configuring the network interfaces for you, but it is still important to double check the configuration.

Navigate through the Failover Cluster Management Console to Failover Cluster Management | <your cluster name> | Networks. You should now see listings for both of your network segments, as shown in Figure D. They should be listed as Cluster Network 1 and Cluster Network 2 respectively.

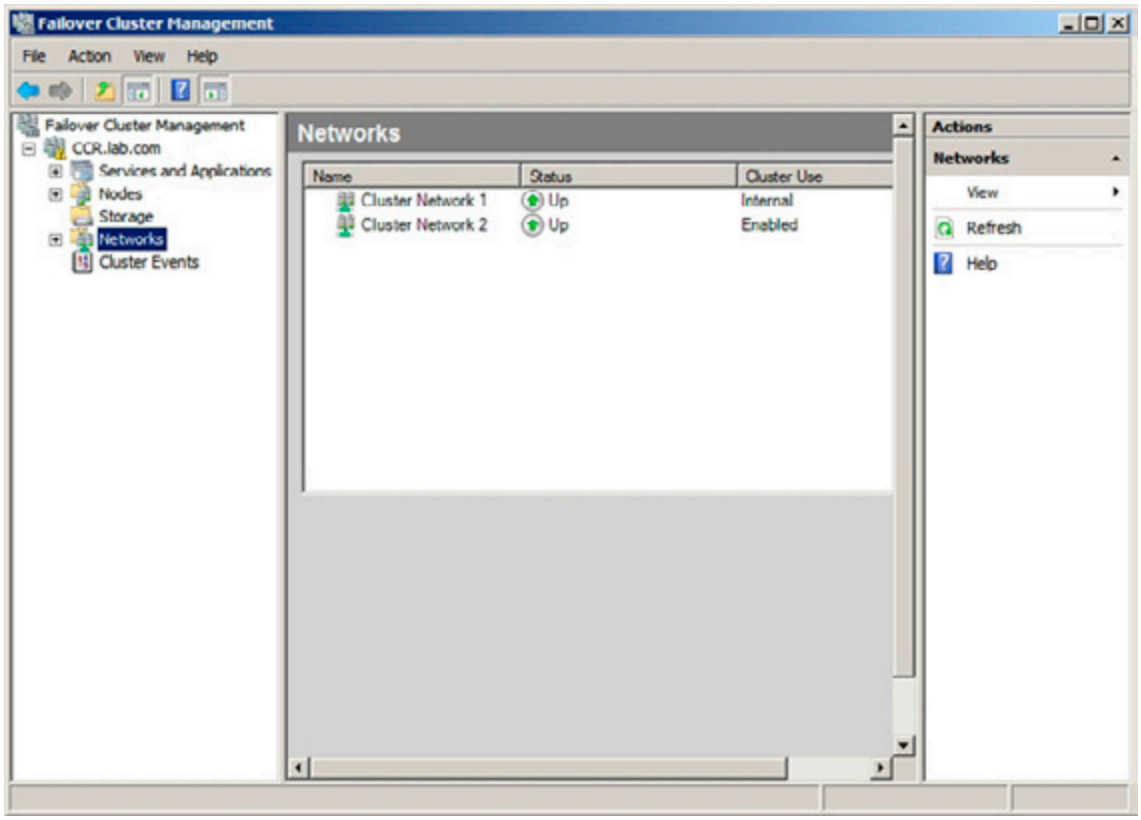

Figure D.

Both networks should be displayed in the Failover Cluster Management Console.

Right click on Cluster Network 1, and choose the Properties command from the shortcut menu. When you do, you will see a properties sheet for the connection. Make sure that this is the connection isn't the private segment that you are reserving for communication between cluster nodes. Now, verify that the Allow Cluster to Use this Connection option is selected You should also make sure that the Allow Clients to Connect Through This Network check box is selected. Click OK to close the properties sheet.

Now, right click on Cluster Network 2, and select the Properties command from the resulting shortcut menu. This network connection should be the private segment that services the cluster nodes. Once again, you want to make sure that the Allow Cluster to Use This Network Connection option is selected. However, the Allow Clients to Connect Through This Network option should not be selected. Click OK to close the properties sheet and return to the Failover Cluster Management Console.

Configure the Node and File Share Majority Quorum

In the previous article, I explained that a majority node set cluster has to have more than two nodes, because otherwise there is no node majority in a failover situation. Since CCR only works on two node clusters though, we need to create a file share witness on our hub transport server. This allows the file share that we create to take the place of the third cluster node.

The first thing that we need to do is to create a share on a hub transport server that can be used as the file share witness. To do so, go to your hub transport server, and open a Command Prompt window. You must now enter the following commands:

```
C:
CD \
MD FSM_DIR_CCR
Net Share FSM_DIR_CCR=C:\FSM_DIR_CCR /Grant:<cluster name>$, FULL
CACLS C:\FSM_DIR_CCR /G BUILTIN\Administrators:F <cluster name>$:F
```

In the code above, we are granting the cluster access to the FSM_DIR_CCR folder on our hub transport server. You will notice that every time I reference the cluster name, I am putting a dollar sign after the cluster name. The dollar sign tells Windows that we are granting access to a computer account rather than to a user account.

Now that we have prepared our hub transport server, we need to configure the cluster quorum settings. To do so, go back to one of your cluster nodes and open the Failover Cluster Management Console. Once the console opens, right click on the listing for your cluster, and select the More Actions | Configure Cluster Quorum Settings commands from the resulting shortcut menus.

At this point, Windows will launch the Configure Cluster Quorum Wizard. Click Next to clear the wizard's Welcome screen, and you will be taken to a screen that asks you to select the quorum configuration that you want to use. Choose the Node and File Share Majority (For Clusters With Special Configurations) option, as shown in Figure E, and click Next.

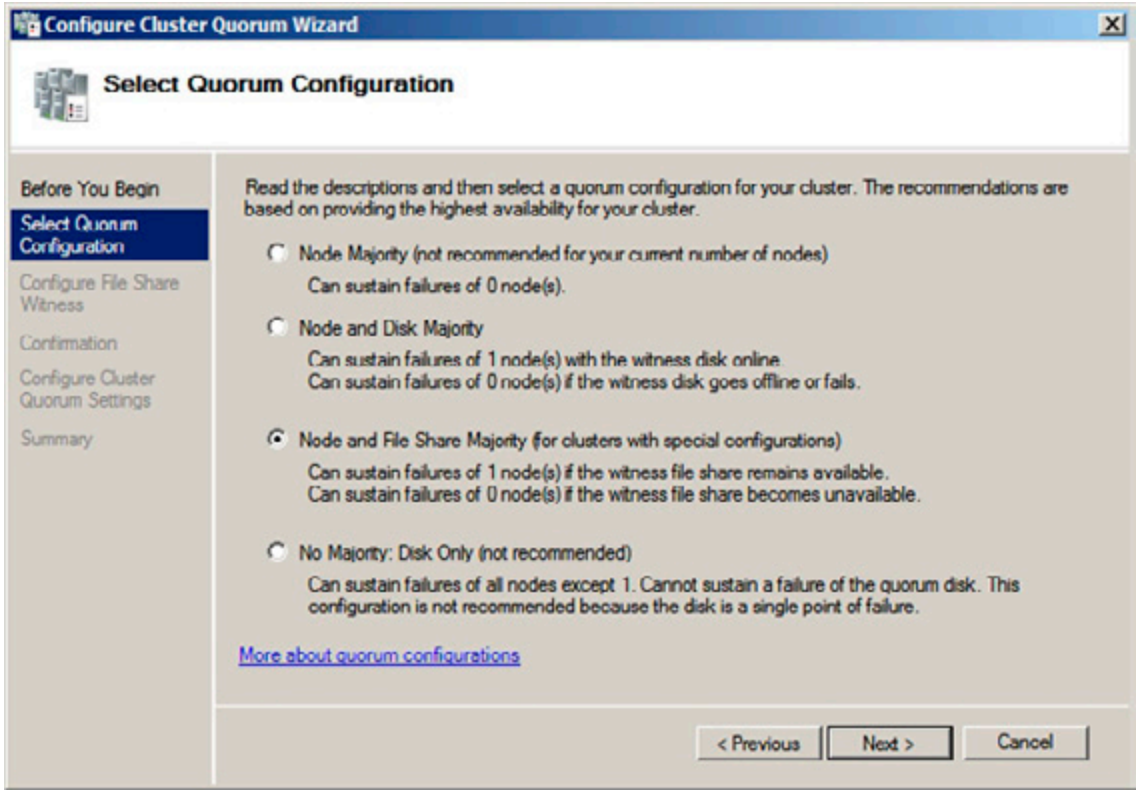

Figure E.

Choose the Node and File Share Majority option

The next screen that you will encounter asks you for the path to the file share that you want the cluster to use. Enter the path to the share that you created earlier in UNC format (\\<your hub transport server's name)\FSM_DIR_CCR). You can see an example of this in Figure F. Click Next, and the wizard will take you to a confirmation screen. Take a moment to make sure that the information that is presented on this screen appears correct, and then click Next. Windows will now configure the quorum settings. When the process completes, click Finish.

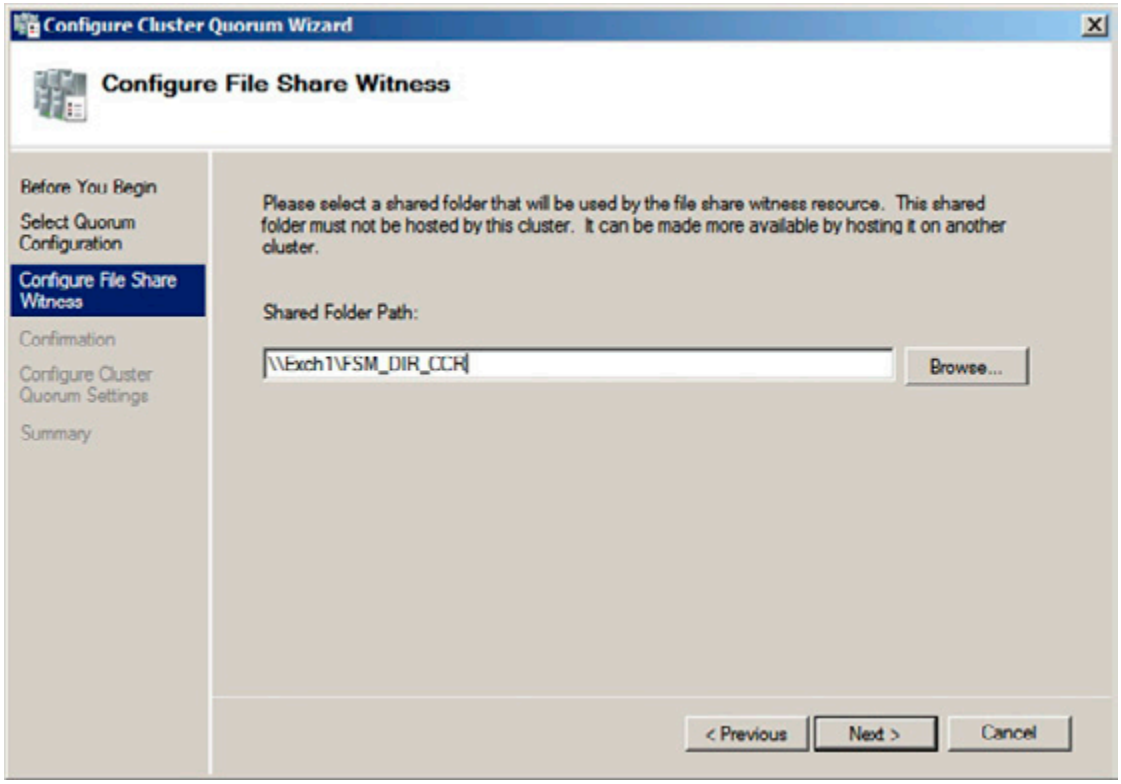

Figure F.

Enter the path to the file share that you created.

Installing Exchange Server

Now that we have created our cluster, it is time to install Exchange Server 2007. There are a few things that you need to take into account prior to beginning the installation process.

For starters, I am assuming that both of your cluster nodes have the Exchange Server prerequisites installedYou should also keep in mind is that the Clustered Mailbox Server role cannot coexist on a server with any other Exchange Server roles. This means that at a minimum you will already need to have an Exchange 2007 hub transport server and client access server in place prior to deploying your clustered mailbox server.

One last thing that you need to keep in mind is that if you are going to be installing Exchange onto a Windows Server 2008 based cluster, then you must install Exchange 2007 SP1 or higher. You can't install the RTM release and then upgrade to SP1 later on. You must start with SP1 or higher. If you are going to be installing Exchange onto Windows Server 2003, then you can install Exchange 2007 with or without the service pack, although it is always preferable to use the latest service pack.

Configuring the Active Node

The process for setting up the active node isn't all that different from configuring a typical mailbox server. Since we must begin by setting up the active node, you should perform these steps on Node 1.

Begin the installation process by double clicking on Setup.exe. When the Exchange Server 2007 splash screen appears, click on Step 4: Install Microsoft Exchange Server 2007 SP1. This will cause Setup to launch the installation wizard.

The first screen that you will see is really nothing more than a welcome screen. You can just click Nextto skip it. You will now be prompted to accept the server's license agreement. After doing so, click Next.

The next screen that you will encounter asks you whether or not you want to enable error reporting. Whether or not you want to use error reporting is really up to you. After you make your decision, click Next.

The following screen will ask you if you want to perform a typical Exchange Server installation or a custom installation. You can only configure a clustered mailbox server by performing a custom installation. Therefore, choose the Custom Exchange Server Installation option, and click Next.

You will now be prompted to select the Exchange Server roles that you want to install. Select the Active Clustered Mailbox Server Role check box, as shown in Figure G. When you do, all of the other options will be grayed out, because the clustered mailbox server role cannot coexist with any other Exchange Server role.

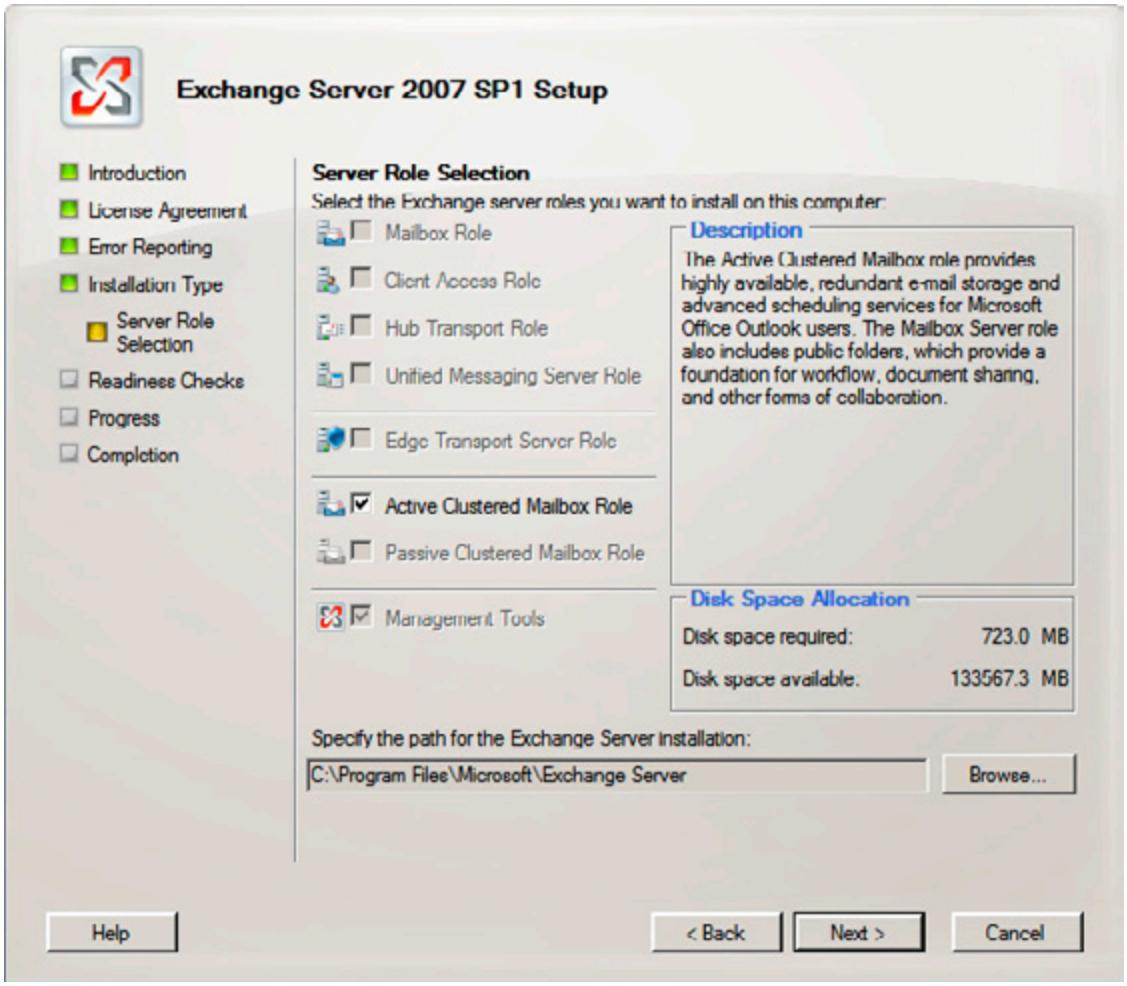

Figure G.

Select the Active Clustered Mailbox Role check box.

Click Next, and Setup will ask you what type of clustered mailbox server you want to create. Choose the Cluster Continuous Replication option. Before continuing, you must enter the name that you want to assign to your clustered mailbox server, as shown in Figure H. This is the name that will be unique to Exchange Server.

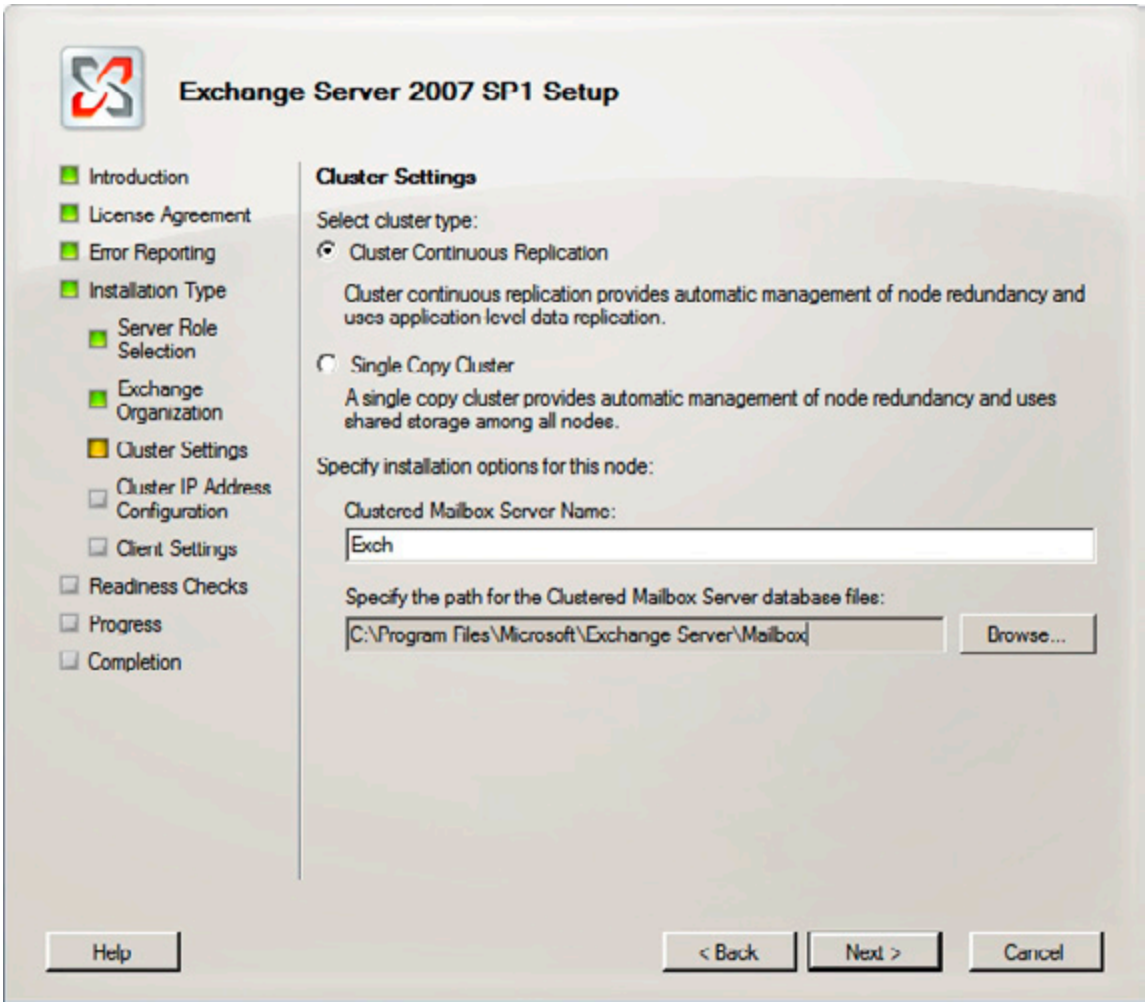

Figure H.

Choose the Cluster Continuous Replication option and then enter the name that you want to use for your clustered mailbox server.

Click Next, and you will be prompted to enter the server's IP address. You should enter the IP address that will be unique to Exchange Server.

Click Next, and Setup will perform a prerequisite check to make sure that all of the necessary components are in place. Assuming that the check succeeds, click the Install button to begin the installation process.

When the installation process completes, click Finish. You may see a message telling you that a restart is required before placing the server into production. If you receive such a message, just click OK to acknowledge it. Rather than restarting the server immediately, we need to stop the clustered mailbox server service, and move it to our passive node. As you will recall though, we haven't actually installed Exchange onto the passive node yet. This makes things a little bit messy, but it isn't going to cause us any problems. My personal preference is to perform this operation from the command line, but you can use the Failover Clustering Console.

To stop the Clustered Mailbox Service, open the Exchange Management Shell, and enter the following command:

```
Stop-ClusteredMailboxServer <your clustered mailbox server name> -StopReason Setup –Confirm:$False
```

Stopping the Clustered Mailbox Service allows our passive node to take ownership of the service. You can now safely restart Node 1. When the node restarts though, Node 2 still has ownership of the clustered mailbox server. To fix this, open a command prompt window (not EMS) and enter the following command:

```
Cluster Group <your clustered mailbox server name> /Move:<the name of node 1>
```

Once the move has completed, we just need to start the Clustered Mailbox Service. To do so, open the Exchange Management Shell and enter the following command:

```
Start-ClusteredMailboxServer <your clustered mailbox server name>
```

Installing the Passive Node

Now that our active node is operational, we can install our passive node. Begin the installation process by double clicking on Setup.exe. When the Exchange Server 2007 splash screen appears, click on Step 4: Install Microsoft Exchange Server 2007 SP1. This will cause Setup to launch the installation wizard.

The first screen that you will see is really nothing more than a welcome screen. You can just click Next to skip it. You will now be prompted to accept the server's license agreement. After doing so, click Next.

The next screen that you will encounter asks you whether or not you want to enable error reporting. Whether or not you want to use error reporting is really up to you. After you make your decision, click Next.

The following screen will ask you if you want to perform a typical Exchange Server installation or a custom installation. You can only configure a clustered mailbox server by performing a custom installation. Therefore, choose the Custom Exchange Server Installation option, and click Next.

You will now be prompted to select the Exchange Server roles that you want to install. Select the Passive Clustered Mailbox Role. Once again, the other roles will be grayed out to prevent you from selecting them, but the management tools will be installed automatically. You also have the option on this screen of specifying a database path. If you do enter a non default path, it must match the path used by the active node.

Click Next, and Setup will perform a readiness check to make sure that the server has been properly prepared. When the readiness check completes, click the Install button. When the installation completes, click Finish.

Once again, you will receive a message telling you that your server needs to be restarted. This time though, you can restart the server without having to do anything special because this is a passive node.

Conclusion

In this article I have explained how to deploy Cluster Continuous Replication in a Windows Server 2008 environment. In Part 3, I will conclude the series by showing you various techniques for managing your cluster.

Active Directory Management with PowerShell in Windows Server 2008 R2

19 November 2009

by Jonathan Medd

> ***One of the first things you notice with Windows Server 2008 R2 is that PowerShell 2.0 has become central to the admin function There is a powerful Active Directory module for Powershell that contains a provider and cmdlets that are designed to allow you to manage Active Directory from the command line. Now, you can also use versions for previous versions of Windows Server.***

Windows Server 2008 R2 and Windows 7 both ship with PowerShell 2.0 installed by default, a fact which begins to demonstrate how important a consistent command line interface is becoming to Microsoft products. Which is why you as a sysadmin should be aware that, in order to excel as a Windows administrator in the 21st Century you will need to get to grips with PowerShell. Starting with products like Exchange Server 2007 and System Center Virtual Machine Manager (SCVMM), moving into the core OS with Windows 2008 R2 and Windows 7, other product groups are now providing PowerShell support with their latest releases.

Active Directory Domain Services in Windows Server 2008 R2 ships with PowerShell support via cmdlets and a provider; in addition the new Active Directory Administrative Center is built on top of Windows PowerShell technology. In this article we will look at how you can use these new tools to more effectively manage your environment.

Active Directory Scripting

When responsible for an Active Directory environment the larger it becomes the more likely it is you are going to want to use some kind of automation tool to manage it effectively rather than be constantly clicking through a GUI interface to complete the same repetitive tasks. Even in environments of 100 users if you were tasked to provision 10 new users last thing on a Friday night ready for a Monday morning would you really want to click through the New User Wizard 10 times and could you guarantee you wouldn't make any typing mistakes? Prior to the release of 2008 R2 some of the typical options for automating Active Directory with scripting were:

- VBScript. There are hundreds of examples on the Internet, one of the best resources being the Microsoft Script Center.
- PowerShell using ADSI, similar to how the VBScript examples are put together.
- PowerShell using .NET Directory Services Classes.
- PowerShell using Quest's Active Directory cmdlets.

Each approach had its good and bad points, but going forward using the PowerShell cmdlets which ship as part of Active Directory in Windows Server 2008 R2 will be the route to take given that will be Microsoft's focus for administrators compared to the above examples.

Active Directory Web Services

Introduced natively as part of Active Directory in Windows Server 2008 R2 is a new service, Active Directory Web Services. This technology permits remote management of any local directory service instance using web service protocols, which by default uses TCP port 9389. This service is included as part of an Active Directory Domain Services (or Active Directory Lightweight Directory Services) installation and is configured to run as an automatic service alongside the main ADDS service.

The Active Directory PowerShell cmdlets and provider which ship with Windows Server 2008 R2 are included as part of a module (modules are essentially the evolution of snapins from version 1 of PowerShell) and use this Web Service to manage Active Directory.

The good news for organisations who will not be immediately upgrading their Active Directory environments to Windows Server 2008 R2 and are currently running either Windows Server 2003 SP2 (or R2 SP2) or Windows Server 2008, with or without SP2, is that this Web Service has been made available as a Release To Web update. This means that you can have the same automation experience today with your current environment without going through the upgrade process for Active Directory which typically for most organisations involves significant planning and quite possibly cost.

There are some system requirements for this downlevel release which you should be aware of before you go diving straight in:

- .NET Framework 3.5 with SP1 must be installed on the Windows Server 2003 or 2008 Domain Controller.
- PowerShell 2.0 for Windows Server 2003 or 2008 must be installed, it's available from Knowledge Base (KB) article 968929.
- For Windows Server 2003 and 2008 based Domain Controllers, you must download and install the hotfix that is described in KB article 969166.
- For Windows Server 2003& based Domain Controllers, you must download and install the hotfix that is described in KB article 969429.
- For Windows Server 2008 based Domain Controllers without SP2, you must download and install the hotfix that is described in KB article 967574.

Note that it is not possible to install the Active Directory module on these downlevel servers and you will instead require a Windows Server 2008 R2 or Windows 7 instance to remotely manage these systems with the Active Directory module. Whilst not a requirement that you install the Web Service on every single downlevel Domain Controller, you may run into issues if you do not give enough coverage across your Domain Controllers. The majority of the Active Directory cmdlets though do include a **Server** parameter which allows you to specify a particular Domain Controller to connect to.

Getting Started

Once you have the Active Directory Web Service installed, either as part of a native Windows Server 2008 R2 installation or using one of the downlevel versions as described above, you can start to manage Active Directory with the PowerShell module. To use this module you can either Remote Desktop to connect to a Domain Controller in your environment or more typically, and also better practise, use tools on your management workstation.

To do this from a Windows 7 based workstation you obviously already have PowerShell 2.0 installed by default, you should also install the Microsoft Remote Server Administration Tools for Windows 7 (RSAT). Amongst the many tools available in this package for remotely managing Windows Server is the Active Directory PowerShell module. After installation of RSAT the feature can be turned on by navigating to **Control Panel, Programs and Features, Turn Windows Features On** or **Off**.

The next step is to open a PowerShell session and use the Import-Module cmdlet to enable the use of the ActiveDirectory module. If this is something you would regularly use on your management workstation then it would be well worth your while adding this line to your PowerShell profile so that this module is available to you in every session without having to type this command.

Import-Module ActiveDirectory

Active Directory Provider

Windows PowerShell 1.0 shipped with a number of providers which gave you access to navigate and update data stores, such the file system or the registry in the Windows OS. Some other products also lend themselves to the concept of a provider and in Windows Server 2008 R2 the Active Directory module ships with a provider. This provider allows you to traverse and update Active Directory just like you were using the old style command prompt to navigate the Windows file system.

After importing the ActiveDirectory module, running the below cmdlet:

Get-PSProvider

will display the list of providers available to you on the system.

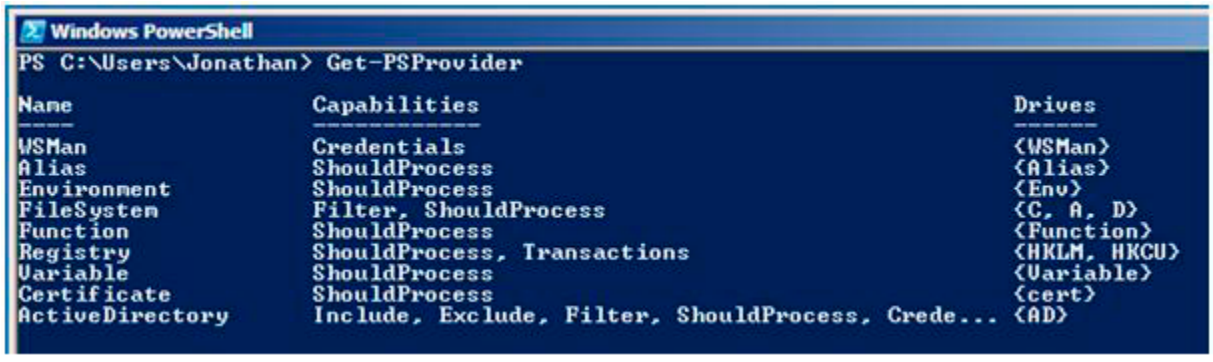

To begin you can use the familiar **cd** (which is an alias in PowerShell for **Set-Location**) to change your location to that of the Active Directory hierarchical navigation system.

You can use other familiar commands to navigate your way around and make changes. We can use the well-known **dir** to show us what is available. Then **cd 'DC=test,DC=local'** to start exploring the domain and then **dir** again will give a well-known view to that seen when first opening the Active Directory Users and Computers GUI tool.

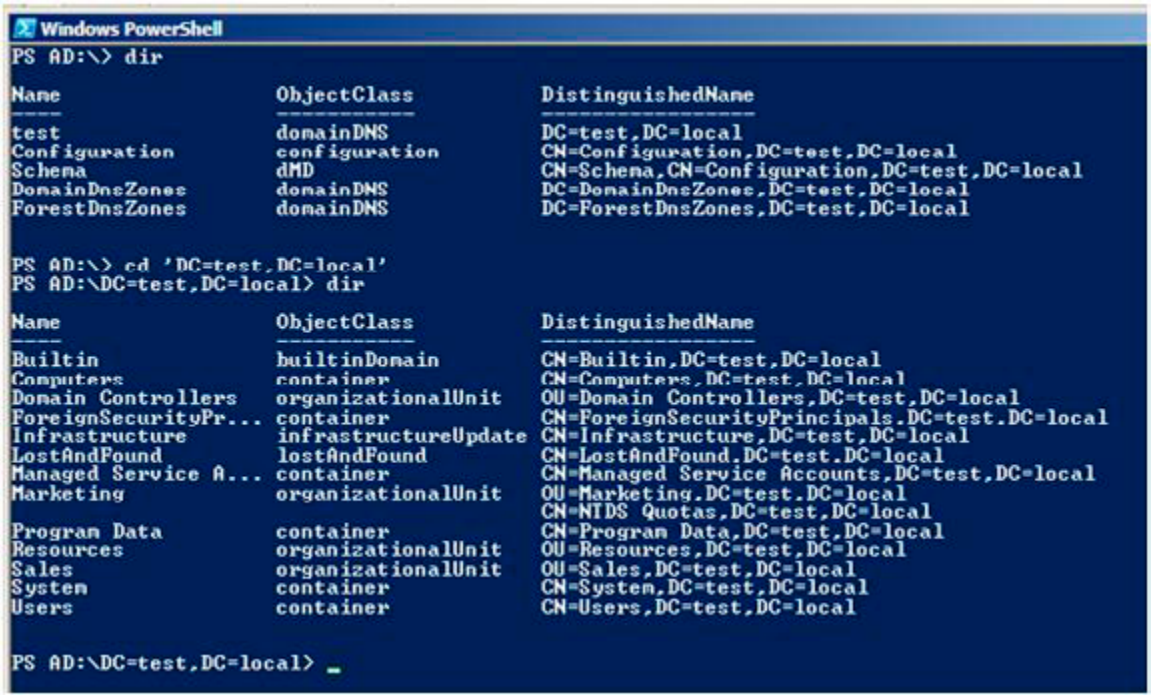

If we then drill down further into the **Users** OU below the **Sales** OU we can observe some user accounts again by using **dir**.

This time we would like to update the Description field for the Joe Bloggs user account to read "Marketing Manager'. To do that we can use the **Set-ItemProperty** cmdlet with the path to the object we wish to change and the values we wish to set:

Set-ItemProperty -Path "CN=Joe Bloggs" -Name Description -value "Marketing Manager"

We can then read this property back with the Get-ItemProperty cmdlet:

Get-ItemProperty -Path "CN=Joe Bloggs" -Name Description

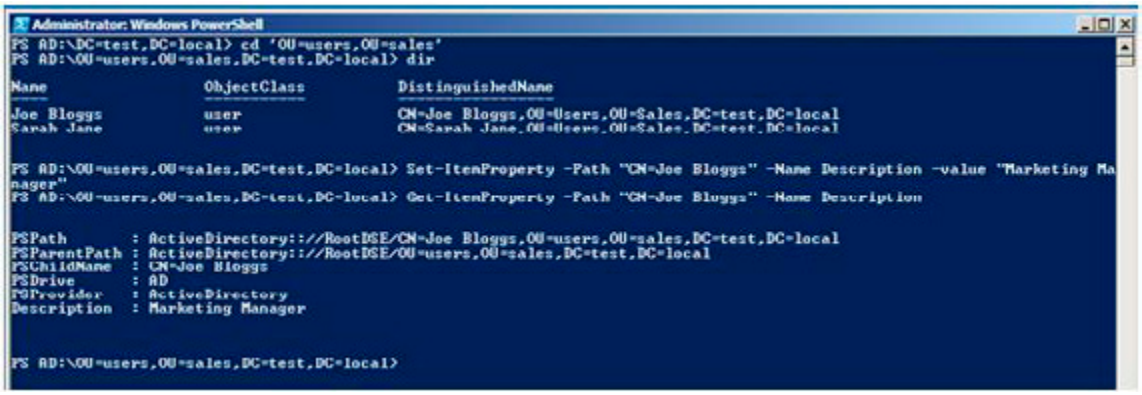

Or in a GUI view:

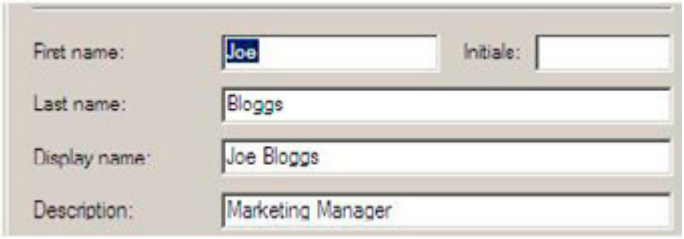

Active Directory Cmdlets

There are 76 cmdlets which make up the Active Directory module in Windows Server 2008 R2. You can view details of them via the Get-Command cmdlet and specifying the module of interest:

Get-Command –Module ActiveDirectory

The best way to learn about what each one does is to check them out in your test environment and also use the built-in PowerShell help documentation via the **Get-Help** cmdlet. For example:

Get-Help Get-ADUser -Full

This will give you detailed information about the cmdlet and how you should go about using it. **Tip:** I often find using the **–Examples** parameter of the Get-Help cmdlet to be a great kick start to learning how a cmdlet works; rather than wade through the text help you can instantly see how you might use it via some examples.

So let's check out some of these cmdlets.

Creating User Accounts

Earlier in the article we talked about how when creating even a small number of users, repeatedly working through an individual wizard for each one could be a dull task and prone to consistency mistakes. Let's say the HR department supplies you with a CSV file containing information about 10 new users and needs you to create these accounts ASAP before Monday.

	A	B	C	D	E	F	G
1	Name	SamAccountName	Description	Department	EmployeeID	Path	Enabled
2	John Smith	john.smith	Sales Manager	Sales	45896	ou=users,ou=sales,dc=test,dc=local	$true
3	Jane Bloggs	jane.bloggs	Sales Director	Sales	45897	ou=users,ou=sales,dc=test,dc=local	$true
4	Freddie Montana	freddie.montana	Sales Assistant	Sales	45898	ou=users,ou=sales,dc=test,dc=local	$true
5	Jo Clark	jo.clark	Administrator	Sales	45899	ou=users,ou=sales,dc=test,dc=local	$true
6	Mark Brown	mark.brown	Account Manager	Sales	45900	ou=users,ou=sales,dc=test,dc=local	$true
7	Anne Wilson	anne.wilson	Marketing Manager	Marketing	45901	ou=users,ou=marketing,dc=test,dc=local	$true
8	George Maker	george.maker	Marketing Director	Marketing	45902	ou=users,ou=marketing,dc=test,dc=local	$true
9	Lisa Taylor	lisa.taylor	Marketing Assistant	Marketing	45903	ou=users,ou=marketing,dc=test,dc=local	$true
10	Joanne Martin	joanne.martin	Administrator	Marketing	45904	ou=users,ou=marketing,dc=test,dc=local	$true
11	Clive Mason	clive.mason	Administrator	Marketing	45905	ou=users,ou=marketing,dc=test,dc=local	$true

For simplicity's sake we'll keep it to a few basic properties like **Name** and **Description**, of course you could and normally would have significantly more. Let's also say to make your life easier you've added a few extra columns of you own like the **SamAccountName** and the **Path** (OU) where you would like the account to be created.

We can take of advantage of a standard PowerShell cmdlet **Import-CSV** which will read in a standard CSV file and create a set of objects based on the data inside the file. We can then send the results of this cmdlet down the pipeline to the **New-ADUser** cmdlet from the ActiveDirectory module and create the 10 accounts in a matter of seconds.

Import-CSV C:\scripts\users.csv | New-ADUser

It really is as simple as that and now you also can leave early on that Friday night like everyone else! Some of the results of those commands are shown below, nicely created in the OU specified in the CSV file.

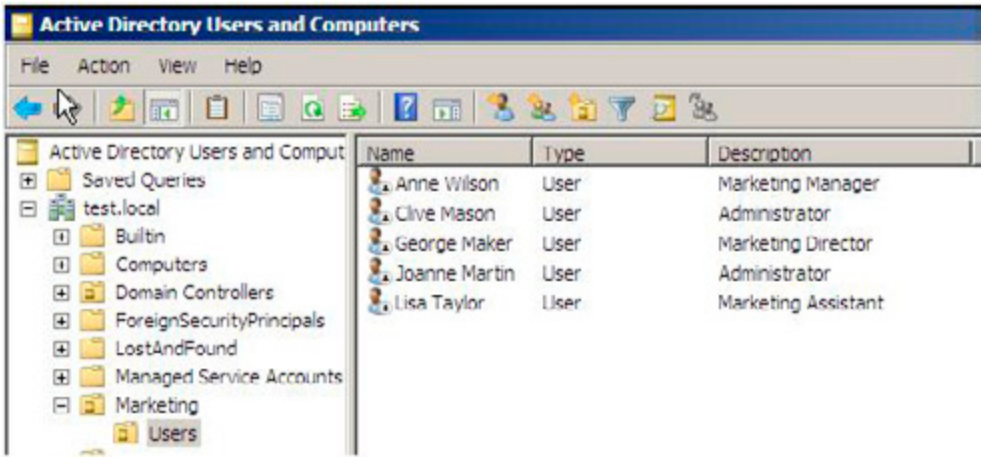

You might be wondering how we made the leap from the information specified in the CSV to the **New-ADUser** cmdlet without having to specify any parameters with this cmdlet. In fact you could create a new user account with something like this:

New-ADUser –Name "John Smith" –SamAccountName "john.smith" –Description "Sales Manager" –Department "Sales" –EmployeeID "45896" –Path "ou=users,ou=sales,dc=test,dc=local" –Enabled $true

We have actually taken advantage of something known as "Pipeline Parameter Binding'. Since we gave the columns in the CSV file the same names as the parameters of the **New-ADUser** cmdlet and these parameters "Accept Pipeline Input ByPropertyName" PowerShell is clever enough to match them up and use them. It's almost like someone has thought to try and make your administrator life just that little bit easier.

Administering Groups

Another frequent maintenance task for the Active Directory administrator is the upkeep of groups and in particular membership of them. Let's take the scenario of the newly created Marketing user accounts and a request to add all of them to the "Marketing Application" group which is used to provide access to one of their main tools.

Again you could do that through Active Directory Users and Computers, edit the membership of the "Marketing Application" group and manually enter all of their names to add them to the group. Or you could be smart, use the **Get-ADUser** cmdlet with a filter to select those users, pipe the results to the standard PowerShell cmdlet **ForEach-Object** and then use the **Add-ADGroupMember** cmdlet to populate the group. OK it might not make that big a difference for only five users, but imagine if that was five thousand.

Get-ADUser -filter * -SearchBase "ou=users,ou=marketing,dc=test,dc=local" | ForEach-Object {Add-ADGroupMember -Identity 'Marketing Application' -Members $_}

Note

The **$_** *at the end of that command refers to the 'current object in the pipeline', i.e. imagine it cycling through each of those users in the filter and substituting that user account for* **$_** *each time.*

The results of those cmdlets displayed in the GUI view of the group.

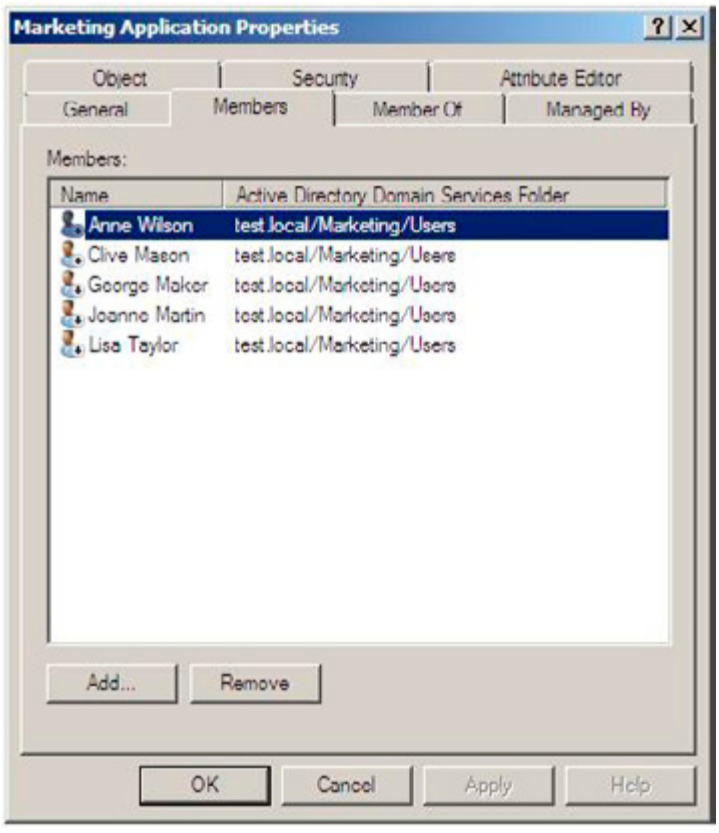

Fine-Grained Password Policies

You may be aware that in Windows Server 2008 a new feature was introduced whereby you were able to add multiple password policies into a single domain. In previous versions of Windows Server there was a restriction of a single password policy in a domain and consequently many organisations ended up deploying multiple domains just for that reason.

So a great new feature, but unfortunately the management tool for configuring these new password policies was wait for it.....ADSI Edit! Whilst there were a number of freeware tools available around the time 2008 shipped which made this easier for the administrator, there are now out of the box PowerShell cmdlets with Windows Server 2008 R2 to help you manage these policies.

The first cmdlet deals with the Default Domain Password Policy **Get-ADDefaultDomainPasswordPolicy**, this will retrieve information about the policy which will apply to all users if there are no Fine-Grained polices.

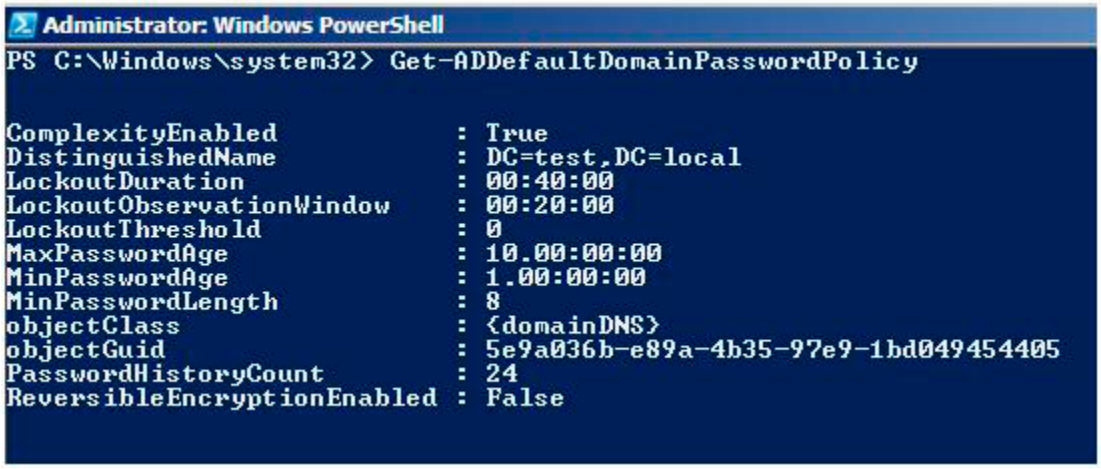

There's an equivalent **Set-ADDefaultDomainPasswordPolicy** which will allow you to make changes to any of the above settings.

You could though have done all of this through Group Policy Editor. The real gains come from being able to create and manage a Fine-Grained Password Policy. Use:

New-ADFineGrainedPasswordPolicy -Name "Standard Users PSO" -Precedence 500 -ComplexityEnabled $true -Description "Standard Users Password Policy" -DisplayName "Standard Users PSO" -LockoutDuration "0.00:12:00" -LockoutObservationWindow "0.00:15:00" -LockoutThreshold 10

to create a new Fine-Grained Password Policy called "Standard Users PSO" and policy settings defined as per the parameter values.

In the above example the Lockout parameters should be specified in the format Day.Hour:Minutes:Seconds, i.e. **LockoutDuration** has been set to 12 minutes.

Fine-Grained Password Polices are applied to Users or Groups, so to apply the **Standard Users PSO** policy to the **Marketing Users** group the **Add-ADFineGrainedPasswordPolicySubject** cmdlet is available to you. Use:

Add-ADFineGrainedPasswordPolicySubject 'Standard Users PSO' -Subjects 'Marketing Users'

to make this change.

Of course you may now have more than one Fine-Grained Password policy. To view details of them all use:

Get-ADFineGrainedPasswordPolicy -Filter {name -like "*"}

and you will see results similar to the below so that you can compare them.

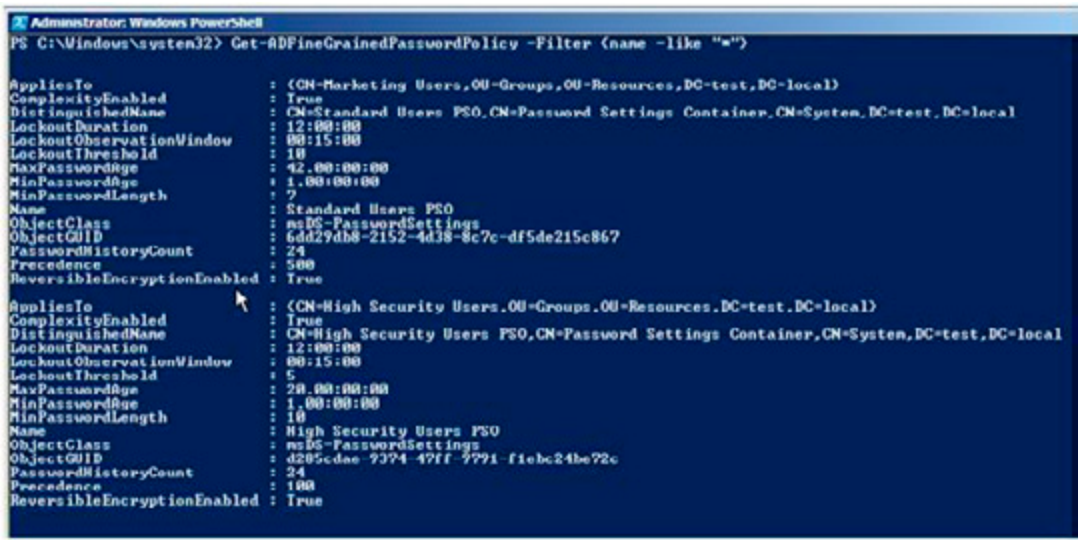

Consequently some users may have more than one policy applied to them, so how do you tell which policy will be effective on them?

- If a user is directly linked to a particular policy then that policy wins. (Generally it is bad practise to link a policy directly to a user; groups should be used for more effective management.)
- If a user is a member of different groups which each have different policies applied to them then the policy with the lowest **Precedence** value will win.
- If there are no Fine-Grained polices created then the Default Domain Policy will apply.

The best way to determine what policy will be applied to a user is to use the **Get-ADUserResultantPasswordPolicy** cmdlet.

For example, George Maker is in the **Marketing Users** and **High Security Users** groups, both of which have different Fine-Grained Password Policies applied to them. By running the command:

Get-ADUserResultantPasswordPolicy -Identity 'George.Maker'

we will be given the resulting answer of which password policy will be applied to him. In this case it's the **"High Security Users PSO"** which is applied because it had the lowest precedence value.

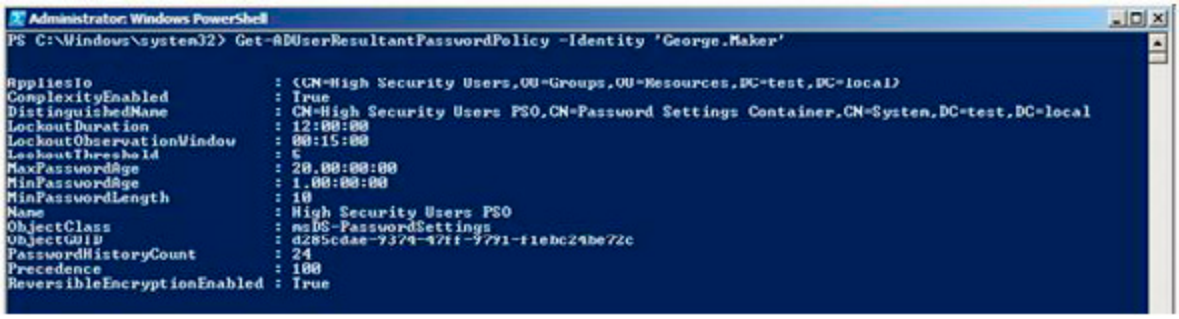

FSMO Roles

In an Active Directory environment all Domain Controllers are equal, although some are more equal than others. There are five roles known as Flexible Single Master Operation roles which typically will live on different domain controllers.

- Domain Naming Master.
- Schema Master.
- Infrastructure Master.
- PDCEmulator.
- RID Master.

You can use the **Get-ADForest** and **Get-ADDomain** cmdlets to determine which Domain Controllers are holding these roles. (In my test environment there is only one DC, but you get the idea)

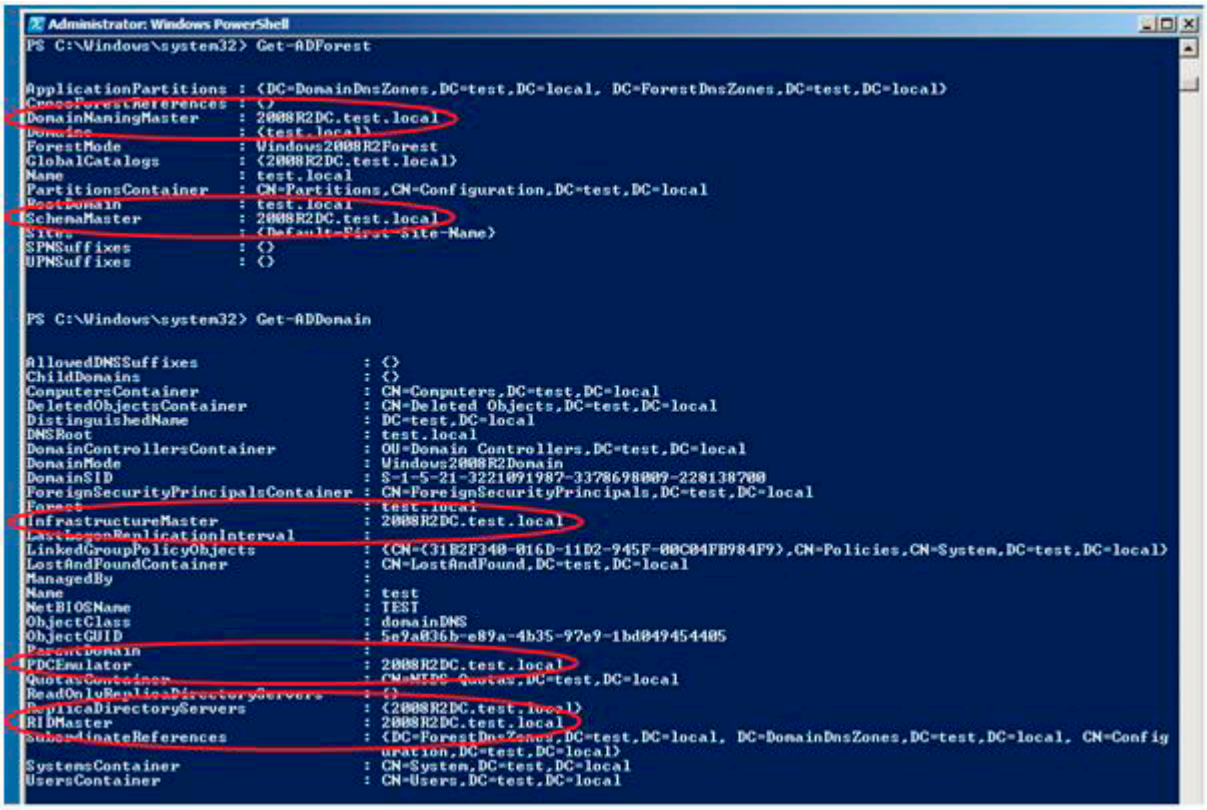

Whilst you could use the Netdom.exe command line tool to obtain the same information, you would need to use an alternative command line tool Ntdsutil.exe, with a different style and syntax, to transfer a FSMO role between Domain Controllers. Consequently by using the **Move-ADDirectoryServerOperationMasterRole** cmdlet to transfer FSMO roles between Domain Controllers, the administrator using the PowerShell module benefits from a consistent command line experience.

```
Move-ADDirectoryServerOperationMasterRole -Identity "2008R2DC2" -OperationMasterRole
PDCEmulator,RIDMaster
```

Further Information

If you want to know more, the best place to look is the Active Directory PowerShell Blog maintained by the team who put the module together. Also, I have made an Active Directory PowerShell Quick Reference Guide which you might find useful for getting started or a handy reference to keep by your desk.

Summary

Windows Server 2008 R2 ships with both PowerShell 2.0 and an Active Directory module containing a provider and cmdlets to enable you to manage Active Directory from the command line. This module has also now been made available for downlevel versions of Windows Server making it more readily accessible to those who might not be in a position to upgrade their Active Directory environments just yet.

It is highly worth spending some time with this module, even if it just getting to grips with the basics; it will help make you a better and more effective administrator of Active Directory.

Upgrade Exchange 2003 to Exchange 2010

11 December 2009

by Jaap Wesselius

In this article, the first of two in which Jaap describes how to move from Exchange Server 2003 straight to Exchange Server 2010, he shows what is required before moving mailboxes from Exchange Server 2003 to Exchange Server 2010. He shows how to upgrade Active Directory, install both Exchange Server 2010 and certificates, and set up the Public Folder replication.

Microsoft released Exchange Server 2010 in October 2009, and this new version of Exchange Server contains a lot of compelling new features such as the new High Availability, the facility to store your Exchange databases on JBOD (Just a Bunch of Disks), the archiving option and the new Outlook Web App. Oh, and do not forget the new Windows Mobile 6.5 and its new mail client.

If you have an Exchange Server 2003 environment you may want to skip Exchange Server 2007 and move directly to Exchange Server 2010. The easiest way to achieve this is to integrate Exchange Server 2010 into the existing Exchange Server 2003 environment, a so called intra-organizational migration. This is also known as transitioning from Exchange Server 2003 to Exchange Server 2010. But what does it take and what issues might arise? This is part 1 of a series of two about moving from Exchange Server 2003 to Exchange Server 2010 and in this document I'll show you what's needed before you start moving mailboxes from Exchange Server 2003 to Exchange Server 2010.

Exchange Server 2003

Suppose we have a fictitious company called Inframan, which is a consulting company specializing in bridges, tunnels, buildings etc. Inframan has approximately 500 employees, 50 employees are working in the office, 450 employees are working "in the field'. Employees within the office have their own desktop which connects to an Exchange 2003 Mailbox Server using Outlook 2003 and Outlook 2007. Employees outside the office connect to the office using their company laptop with Outlook 2007 and Outlook Anywhere and with Windows Mobile devices. When needed they can use their PC at home to use Outlook Web Access to access their mailbox. Typical usage profile is "light," approximately 25 messages are received per day and 10 messages are sent per day, per user that is. Behind the firewall is an ISA Server 2006 acting as a reverse proxy to publish all Exchange Services to the Internet. Inframan's environment will look something like this:

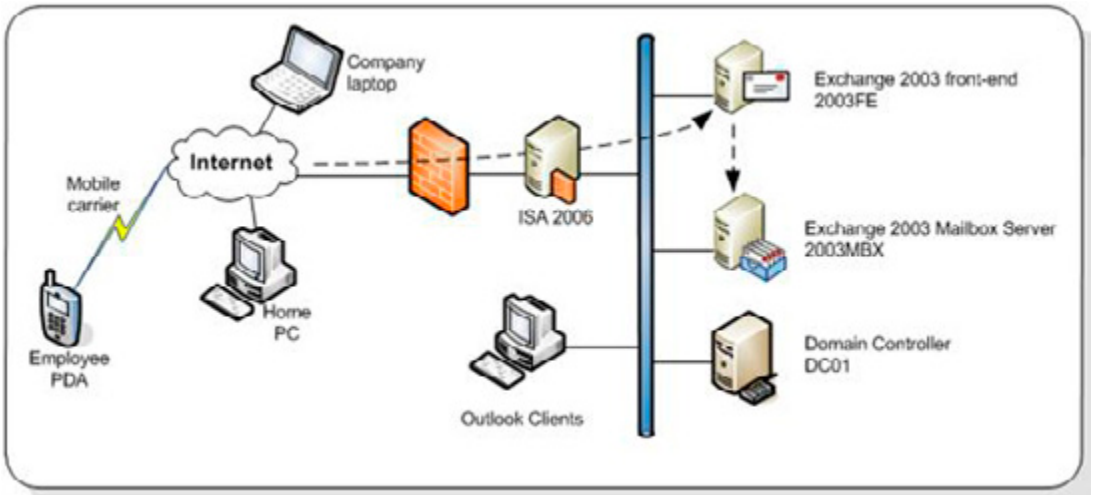

Inframan is using only one namespace for accessing all services from the Internet: webmail.inframan.nl. This is used for Outlook Web Access, Outlook Anywhere and Windows Mobile devices.

Recently Inframan has been thinking about upgrading to Exchange Server 2007, but decided to move directly to Exchange Server 2010.

Coexistence with Exchange Server 2010

Exchange Server 2010 can easily coexist in a Exchange Server 2003 organization as long as the Exchange Server 2010 prerequisites are met:

- The Active Directory forest needs to be in Windows Server 2003 forest functionality mode.
- All domains that contain Exchange recipients need to be in Windows Server 2003 domain native mode.
- The Global Catalog Servers and the Active Directory Schema Master need to be at a minimum level of Windows Server 2003 SP1 (which equals to Windows Server 2003 R2).
- The Exchange 2003 organization needs to be running in "native mode."
- Link State updates on all Exchange Server 2003 servers need to be disabled according to Microsoft knowledge base article KB 123456.

Be careful when upgrading your Active Directory Domain Controllers since not all versions are supported to run with Exchange Server 2003. For a complete overview check the Microsoft Technet Site: http://technet.microsoft.com/en-us/library/ee338574.aspx.

Inframan will build two new Exchange Server 2010 servers, one combined Hub Transport Server / Client Access Server and one dedicated Mailbox Server. These Servers will be installed in the same Windows Server 2003 Active Directory domain as the Exchange Server 2003 organization. This will greatly improve the ease of moving mailbox from Exchange Server 2003 to Exchange Server 2010.

Moving from Exchange Server 2003 to Exchange Server 2010 in the same Active Directory forest is called transitioning. Building a new Active Directory forest with a new Exchange Server 2010 organization and moving mailboxes from the old Active Directory to the new Active Directory is called migrating.

The interim messaging environment, where both Exchange Server 2003 and Exchange Server 2010 coexist in the same Active Directory domain will look like this:

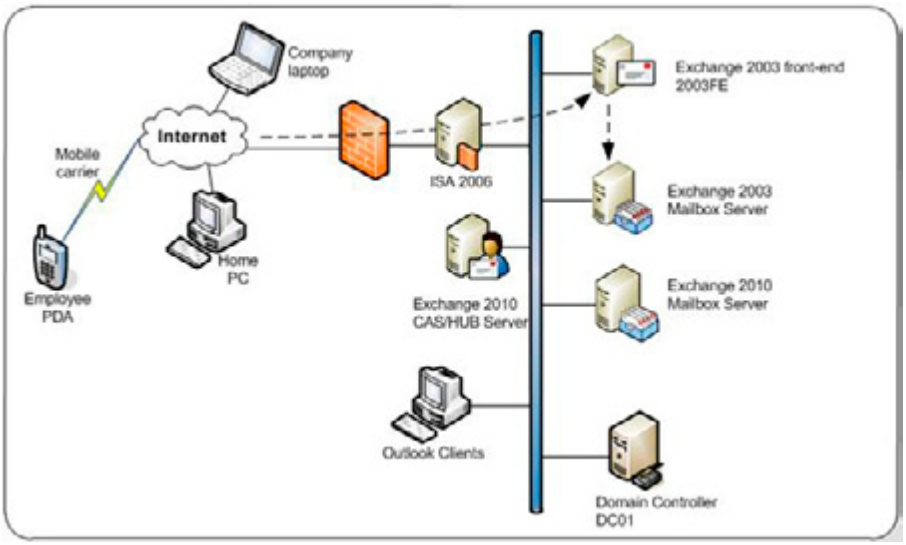

In Exchange Server 2007 Internet clients could connect to the Exchange Server 2007 Client Access Server while the mailbox was still on Exchange Server 2003. The Client Access Server retrieves the data out of the mailbox and sends it back to the Internet client. In Exchange Server 2010 this has changed. When a client connects to Exchange Server 2010, it actually connects to the Exchange Server 2010 Client Access Server and if the mailbox is still on the Exchange Server 2003 Mailbox Server then the client is redirected to the Exchange Server 2003 front-end server. This front-end server then handles the connection request. This automatically means the namespaces of the Exchange environment will change. For Inframan this means that the following namespaces are used:

https://webmail.inframan.nl – This is used by all Internet clients that connect to the Exchange environment. This name is not different than in the Exchange Server 2003 namespace, but it will now point to the Exchange Server 2010 Client Access Server.

https://autodiscover.inframan.nl – This is used by Outlook 2007 and (Outlook 2010) clients for autodiscover purposes.

https://legacy.inframan.nl – This will be the new namespace for the Exchange Server 2003 front-end server. This automatically means that the namespace for the Exchange Server 2003 front-end server is going to change!

The servers that will hold the Exchange Server 2010 server roles have the following prerequisites:

- The servers need to be running on Windows Server 2008 or Windows Server 2008 R2.
- .Net framework 3.5 with SP1 needs to be installed.
- PowerShell 2.0 needs to be installed.
- Office 2007 Filter packs needs to be installed for the Hub Transport Server role and the Mailbox Server role.

Make sure that after installing Windows on the servers that they are up-to-date with the latest hotfixes and service packs.

The first step for Exchange Server 2010 Server is to upgrade the Active Directory schema to contain the Exchange Server 2010 extensions. This is achieved by using the Exchange Server 2010 setup application followed by a number of parameter:

Setup.com /PrepareLegecyExchangePermissions – Exchange Server 2003 uses the Recipient Update Service to stamp the user with the appropriate Exchange attributes during provisioning. This is replaced in Exchange Server 2010 by E-Mail Address Policies. The /PrepareLegacyExchangePermissions parameter changes security settings so that both the Recipient Update Service and E-mail Address Policies can coexist in the same Active Directory;

Setup.com /PrepareSchema – This command upgrades the Active Directory schema to include the Exchange Server 2010 extensions. This can be checked by using ADSIEDit and checking the value of the UpperRange parameter of the CN=ms-Exch-Schema-Version-Pt object in the Schema. This should have one of the following values:

Value	Corresponding Exchange version
6870	Exchange Server 2003 RTM
6936	Exchange Server 2003 service pack 2
10628	Exchange Server 2007 RTM
11116	Exchange Server 2007 service pack 1
14622	Exchange Server 2007 service pack 2
14622	Exchange Server 2010 RTM

Note that the value is the same in Exchange Server 2007 service pack 2 and in Exchange Server 2010 RTM – this is because Exchange Server 2007 service pack 2 will install the Exchange Server 2010 schema extensions.

Setup.com /PrepareAD – This command upgrades the Exchange organization, which is stored in the configuration partition in Active Directory to support Exchange Server 2010. In Exchange Server 2003 information is stored in the "First Administrative Group" or perhaps more if you created additional Administrative Groups. The Exchange Server 2010 setup application will create a new Administrative Group called "Exchange Administrative Group (FYDIBOHF23SPDLT)" where all Exchange Server 2010 configuration information is stored. This will be visible in the Exchange Server 2003 System Manager:

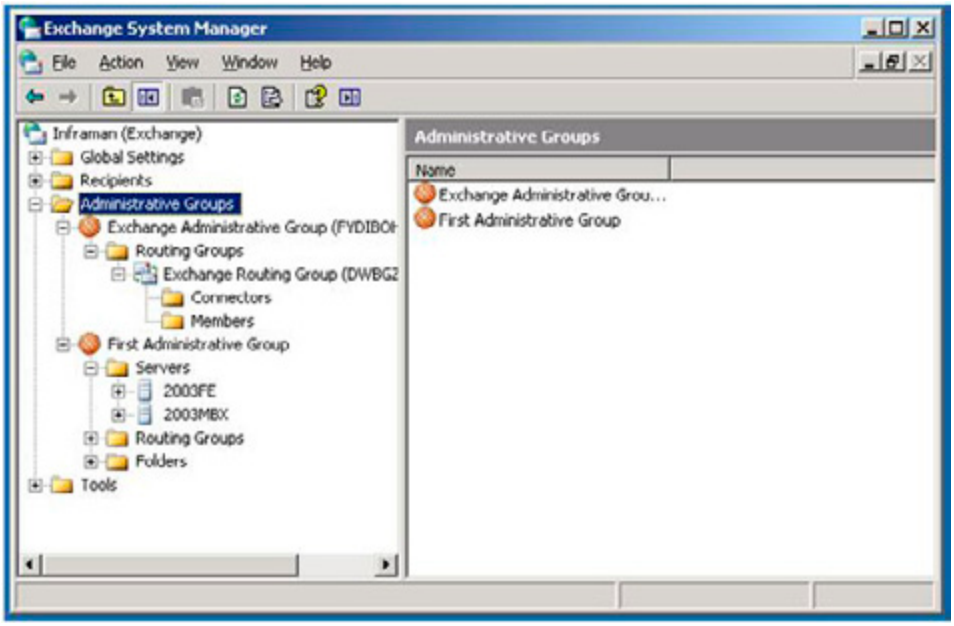

Setup.com /PrepareDomain – This is the last step in preparing the Active Directory and will create all necessary groups in the domain being prepared.

When Active Directory is fully prepared we can continue with installing the first Exchange Server 2010 server in the environment. For our example, this has to be the combined Hub Transport and Client Access Server. Start the graphical setup program (setup.exe) and download the Language File bundle if needed. If you select "install only languages from the DVD" only the language setting of your DVD (for example English or French) will be available. This is used not only for the language of the Exchange Server, but also the available language settings for the clients being used.

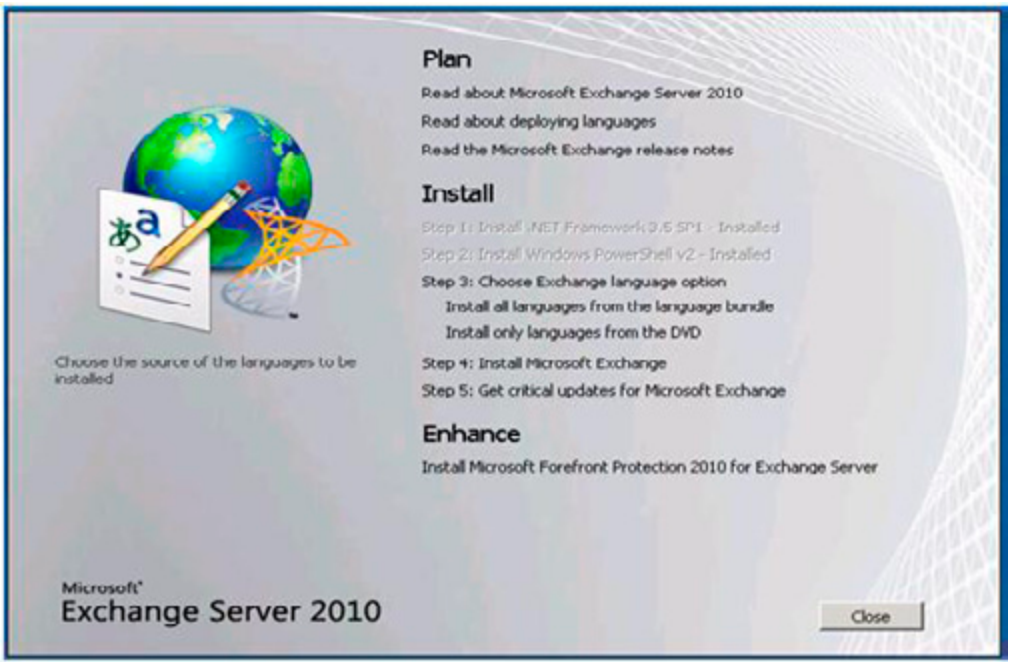

During the installation of the combined Hub Transport and Client Access Server a so called "custom setup" will be used. This means we can select which server roles will be installed. In the Inframan example the following needs to be selected during setup:

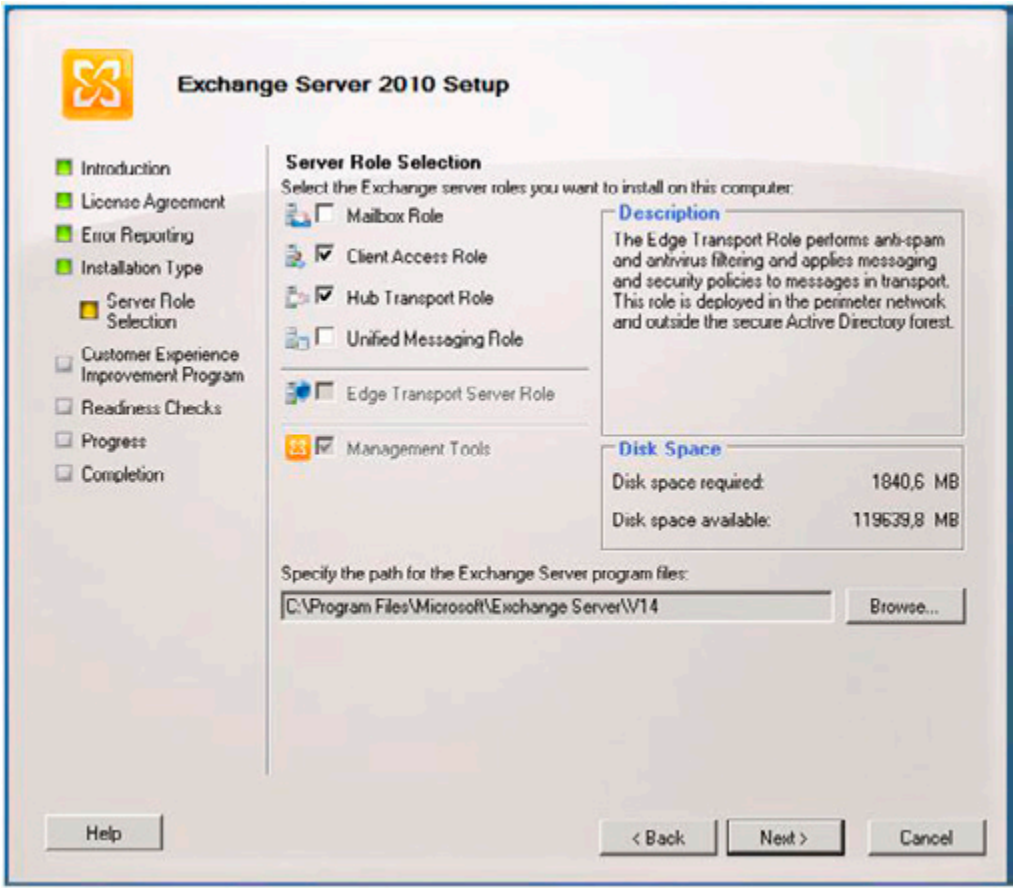

When continuing the setup application a window will be shown asking if this Client Access Server is Internet facing and if so, what the external domain will be. This is an important step because it configures the Client Access Server automatically with the appropriate settings. Check the "The Client Access server will be Internet-facing" option and enter the external domain name. This is "webmail.inframan.nl" in our example.

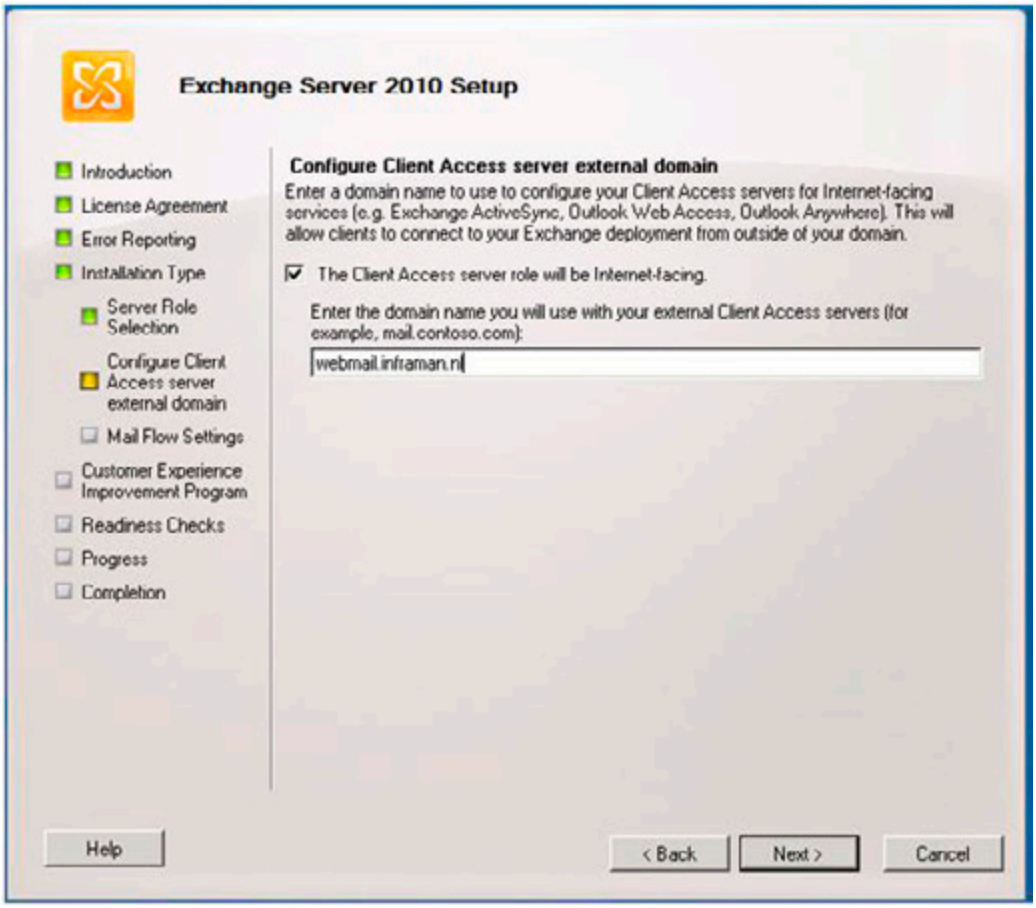

Exchange Server 2003 uses Routing Groups to determine the proper way to route messages while Exchange Server 2010 uses Active Directory sites for routing. These are not compatible with each other so a legacy Routing Group Connector will be created within Exchange Server 2010. This legacy connector connects Exchange Server 2010 with Exchange Server 2003 so messages can be sent between the two Exchange versions. During setup of the first Hub Transport Server an Exchange Server 2003 Hub Server needs to be selected. This is the server the legacy Routing Group Connector will connect to:

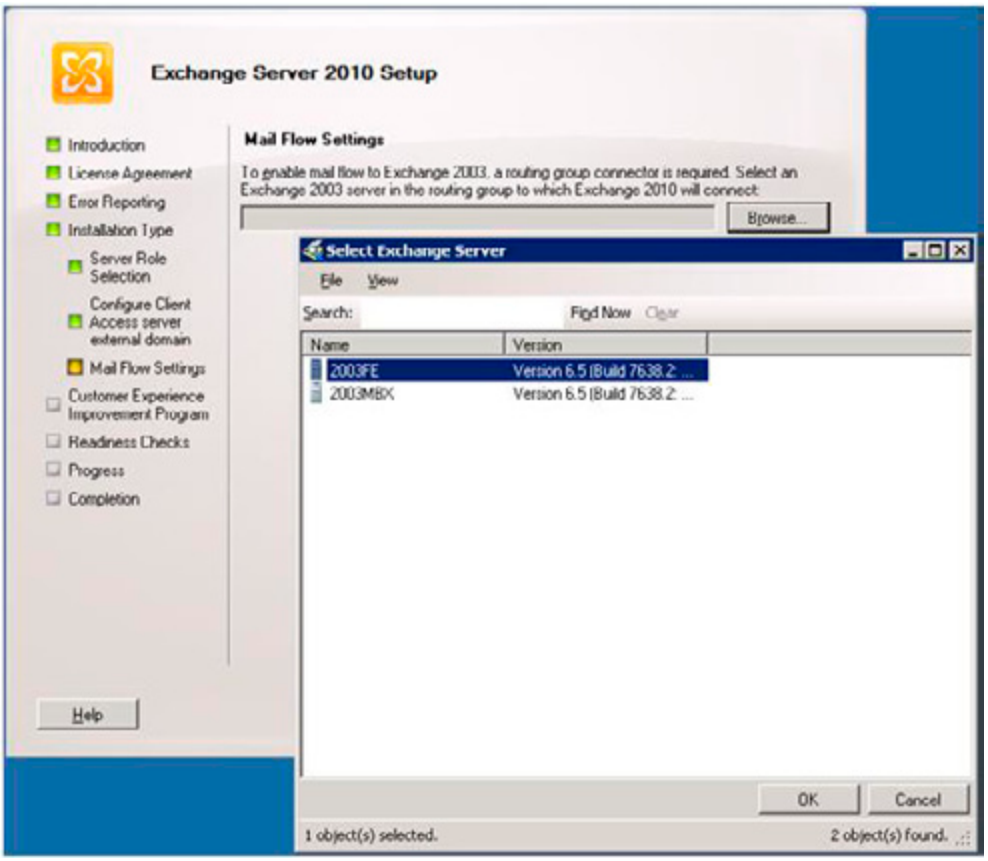

Note that this choice can be changed and/or added to after setup is complete.

Now finish the setup wizard and install the Client Access and Hub Transport Server roles on this server.

It is also possible to use the command line setup application to setup the above mentioned configuration. Open a command prompt, navigate to the installation media and enter the following command:

```
Setup.com /mode:install /roles:ht,ca,mt /ExternalCASServerDomain:
webmail.inframan.nl /LegacyRoutingServer:2003FE.inframan.local
```

Mailbox Storage Design

Before installing the Exchange Server 2010 Mailbox Server role a proper storage design has to be made. Microsoft has recently released the new storage calculator, which is now called the "Exchange 2010 Mailbox Server Role Requirements Calculator" and can be downloaded here:

http://msexchangeteam.com/archive/2009/11/09/453117.aspx.

The Requirements Calculator needs to be used for a proper storage design. The following variables are used in the Requirements Calculator for our example:

Variable	Value
Number of mailbox servers	1
Number of mailboxes	500
Usage profile	Light
Average message size	75 KB
Personal Archive	0 MB
Mailbox Size	1024 MB
Deleted items retention	14 days

The Requirements Calculator will show the following results:

Variable	Value
Number of Databases/server	4
Number of mailboxes/database	125
Log files generated/day/mailbox	20
Mailbox server internal memory	8 GB
Database size plus overhead	177 GB
Total Log file size plus overhead/database	11 GB
Total database size	707 GB
Total log file size	42 GB
Total LUN size databases	972 GB
Total database required IOPS	72
Total LUN size log files	53 GB
Total log required IOPS	14

An interesting part of Exchange Server 2010 is the database technology. Microsoft has made significant changes to the database structure to lower the disk performance requirements. It should be sufficient to run the Mailbox databases and its accompanying log files from SATA disks.

In the Requirements Calculator there's the possibility to enter the disk configuration. For the new Inframan Mailbox server 7.200 RPM SATA disks with a capacity of 500 GB will be used for storing the databases and 7.200 RPM SATA disks with a capacity of 250GB will be used for storing the log files. This disk configuration is not exactly a high end configuration, but it is by far the most cost effective solution.

The Requirements Calculator contains a tab called "Storage Design." When using the above mentioned values the Calculator recommends a RAID1/0 configuration with 6 SATA disks for storing the Mailbox Databases and a RAID1/0 configuration with 2 SATA disks for storing the Log Files.

Installing the Mailbox Server role

When the storage solution has been properly designed and implemented the Exchange Server 2010 Mailbox Server role can be installed. As with the Client Access and Hub Transport Server roles make sure you download the Language Pack during setup. Select a "custom setup" and select only the Mailbox Server role when you get to the "Server Role selection" window as shown in Figure 5. Finish the setup wizard and install the Mailbox Server role. After installation of the 2nd server the organization is ready to be configured and we can prepare for start moving mailboxes from Exchange Server 2003 to Exchange Server 2010.

Configuring the Exchange Server 2010 servers

When both Exchange servers are installed it is time to configure the Exchange environment properly before Exchange Server 2010 can be used and mailboxes can be moved. The following needs to be configured:

- Relocate the Mailbox Databases on the new storage solution.
- Unified Communications certificate on the Client Access Server.
- New server certificate on the Exchange 2003 front-end server.
- OWA 2010 needs to be configured for use with Exchange Server 2003.
- Public Folder replication.
- A send and receive connector also have to be configured, but I will describe this in the next article when the mail flow will be changed from Exchange Server 2003 to Exchange Server 2010.

Relocate the Mailbox Databases

On the new Mailbox Server there are two drives, from a hardware perspective configured as outlined before. These drives are F:\ for the Mailbox Databases and the Public Folder database and drive G:\ for the Log Files.

To change the location of the Mailbox Database open the Exchange Management Console and navigate to the Database Management, which can be found in the Organization Configuration. Right click the database and select "Move Database Path." Change the Database file path to a directory on drive F:\ and change the Log folder path to a directory on drive G:\. Repeat this step for the Public Folder database.

If needed create new databases and locate the new database file on drive F:\ and the accompanying log files on driver G:\

Unified Communications Certificate

On the Exchange Server 2010 Client Access Server a new 3rd-party Unified Communications certificate needs to be installed. According to Microsoft knowledge base article 929395 (http://support.microsoft.com/kb/929395) the following Certificate Authorities are supported for use with Unified Communications certificates:

Entrust – http://www.entrust.net/microsoft/

Digicert – http://www.digicert.com/unified-communications-ssl-tls.htm

Comodo – http://www.comodo.com/msexchange

However, most SSL Certificate Authorities can generate UC/SAN certificates that will work just fine. New in Exchange Server 2010 is the possibility to request certificates using the Exchange Management Console. Open the Exchange Management Console and select the Server Configuration in the navigation pane. Select the Exchange Server 2010 Client Access Server and create a new certificate request. For our environment we have to use the following domain names in our certificate:

- Webmail.inframan.nl
- Autodiscover.inframan.nl
- Legacy.inframan.nl

During the coexistence phase Internet clients will connect to the Exchange Server 2010 Client Access Server while their mailbox is still on Exchange Server 2003. The client request will then be redirected to the old Exchange Server 2003 front-end server. This server will therefore get a new FQDN (Fully Qualified Domain Name) and thus need a new certificate. This new FQDN will be legacy.inframan.nl.

OWA Configuration

During installation of the Exchange Server 2010 Client Access Server all settings have been configured for use on the Internet. The only thing that needs to be configured is the coexistence information for Outlook Web App. The Client Access Server needs to be configured in case a mailbox is still on Exchange Server 2003 and the client needs to be redirected to the Exchange Server 2003 front-end server.

On an Exchange Server 2010 server enter the following Management Shell Command:

```
Set-OWAVirtualDirectory <CASHUB01>\OWA "
-ExternalURL https://webmail.inframan.nl/OWA "
-Exchange2003URL https://legacy.inframan.nl/exchange
```

This will make sure that when a user connects to Exchange Server 2010 Client Access Server for Outlook Web Access and the mailbox is still on Exchange 2003 the client will be redirected to the old Exchange Server 2003 front-end server.

Public Folder Replication

During installation of the Mailbox Server a new Exchange Server 2010 Mailbox Database will be automatically created. After installation you have to make sure that this database is moved from the default location to an alternate location for recovery and performance reasons.

A new Public Folder database will also be automatically created on the new Mailbox Server. The hierarchy, which is the structure of all Public Folders will be automatically replicated between all Public Folder Databases in the entire organization. The content replication of the Public Folders will have to be configured manually though.

To replicate the Offline Address Book and Free/Busy folders from Exchange Server 2003 to Exchange Server 2010 open the Exchange System Manager on the Exchange Server 2003 server and navigate to the System Folders in the "Folders" folder in the First Administrative Group Navigate to the first Offline Address Book folder, right click it and select "All Tasks…." The next is to select "Manage Settings."

If you want to toggle between the System Folders and the normal Public Folders, navigate to the Public Folders, right click the Public Folders and select "View System Folders" or "View Public Folders."

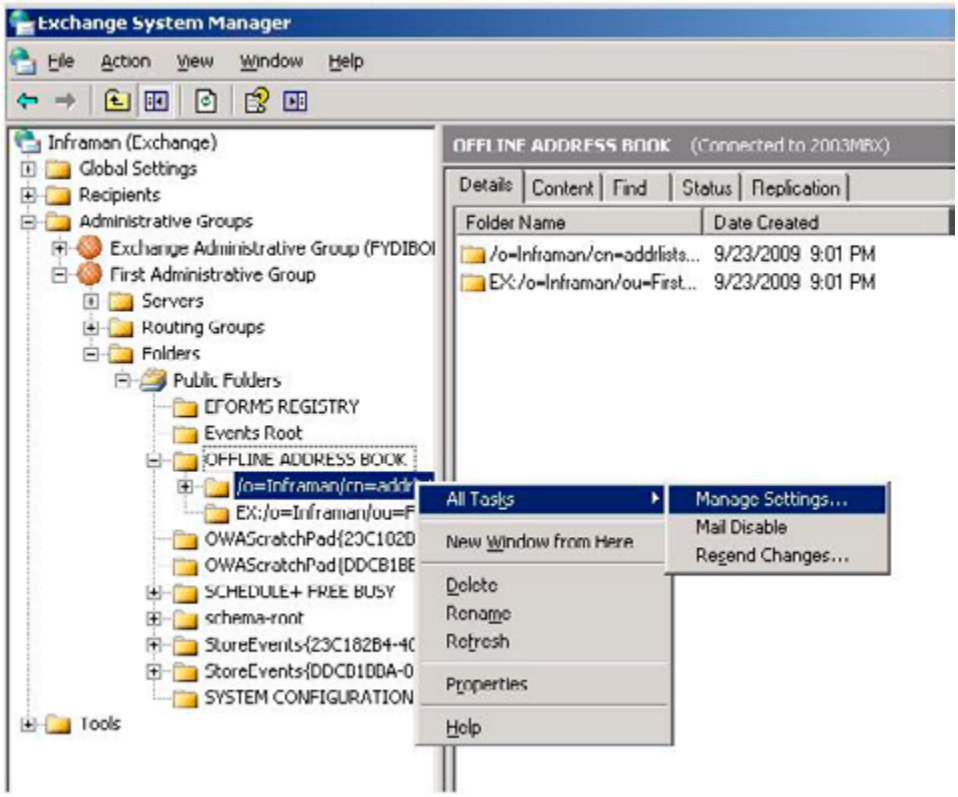

The "Manage Public Folder Settings wizard" will appear. Click Next on the Welcome page and select the "Modify lists of replica servers." Follow the wizard and add the Exchange Server 2010 Mailbox Server role as a new replica. When finished, the folder and all its subfolders will be replicated to the Exchange Server 2010 Public Folder database. Repeat this step for the second Offline Address Book folder and the Schedule+ Free Busy folder.

Notes

When the "Manage Settings" option is not available you can select "Properties" and select the replication tab to add the Exchange Server 2010 Public Folder Database.

Replication of public folders can take quite some time.

The (default) Public Folder that are located on the Exchange Server 2010 Mailbox Server should be replicated to the Exchange Server 2003 Mailbox Server. To accomplish this logon to the Exchange Server 2010 Mailbox Server, open the Exchange Management Console and navigate to the Tools node. Under the Tools node open the Public Folder Management Console.

Right-click the Offline Address Book in the results pane, select Properties and click the Replication tab.

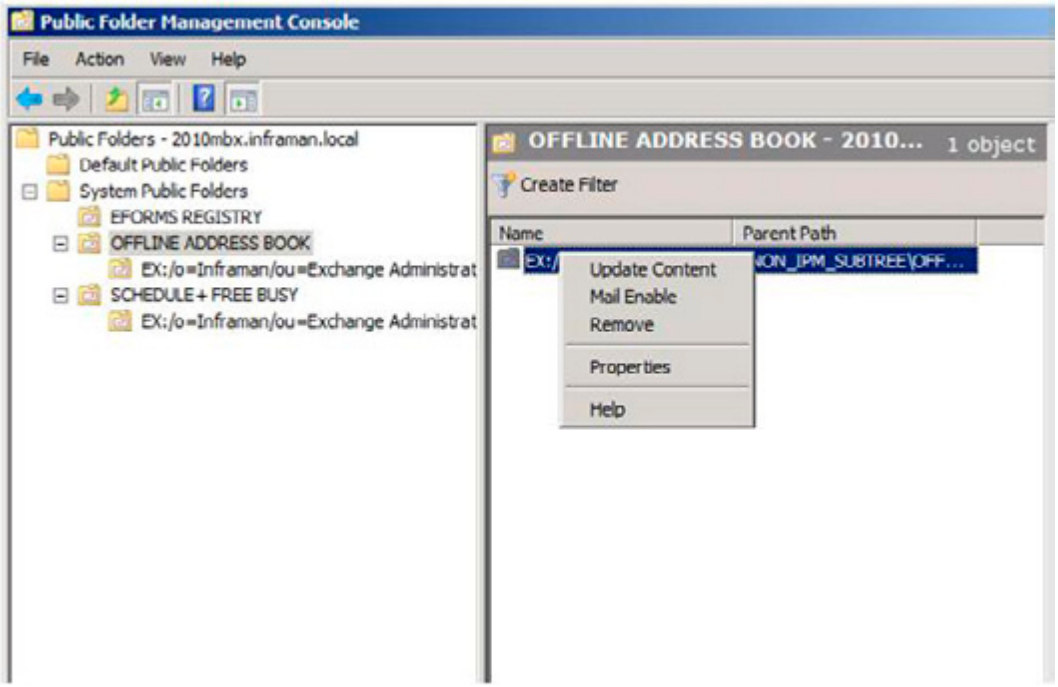

Add the Exchange Server 2003 Mailbox Server to the replica list, the contents will now be replicated to the Exchange Server 2003 Mailbox Server. Be aware that Public Folder replication is a low priority mechanism, so it takes some time before both Public Folder databases are in sync.

Repeat these steps for the Schedule+ Free/Busy folder.

Summary Part I

In this first article out of a series of two, I explained what steps are needed before you can start moving mailboxes from Exchange Server 2003 to Exchange Server 2010. In this article the Active Directory was upgraded, two servers with Exchange Server 2010 were installed, certificates were installed and the Public Folder replication was setup.

In the next article I will cover the actual movement of the mailboxes and the steps that are needed to decommission the Exchange 2003 servers like moving the Offline Address Book generation server and conversion of Recipient Policies and Address Books. Stay tuned!

Customizing the Outlook Address Book

17 December 2009

by Ben Lye

It is possible to change the fields in the Outlook address book to make them a better fit for your organisation. Exchange provides templates for Users, Contacts, Groups, Public Folders, and Mailbox Agents that are downloaded to Outlook. Any changes will be visible in the address book. As usual with Exchange, careful planning before you make changes pays dividends, Ben Lye explains.

I was recently asked if it was possible to change the information fields which are displayed in the Outlook Address Book for users – the person making the request wanted to add an additional telephone number field on the General and Phone/Notes property pages to display an internal extension.

This kind of customisation is probably something that many Exchange organisations can benefit from, and the changes to the Outlook Address Book can be implemented easily within Exchange using the Exchange Details Template Editor.

The Details Template Editor is an MMC snap-in which provides a GUI for editing the object properties which are displayed when an object opened from the address book. Details templates can be modified for Users, Contacts, Groups, Public Folders, and Mailbox Agents. The Advanced Find search dialogue box can also be edited. Each of the six template types can be modified in 50 different languages. The Details Template Editor is installed along with the Exchange Management Tools.

To start the Details Template Editor in Exchange 2007 RTM:

- On the taskbar, click Start, and then click Run.
- Type "mmc" in the Open field.
- On the Console menu bar, click File, and then click Add/Remove Snap-in.
- In Add/Remove Snap-in, on the Standalone tab, click Add.
- In Add Standalone Snap-in, select Details Templates Editor from the list of available stand-alone snap-ins, and then click Add.
- Click Close to close the Available Snap-ins dialog box, and then click OK on the Add/Remove Snap-in dialog box.

To start the Details Template Editor in Exchange 2007 SP1 or SP2 or Exchange 2010 either use the method above, or:

- Launch the Exchange Management Console.
- Select Toolbox in the console tree.
- Double-click Details Template Editor in the results pane.

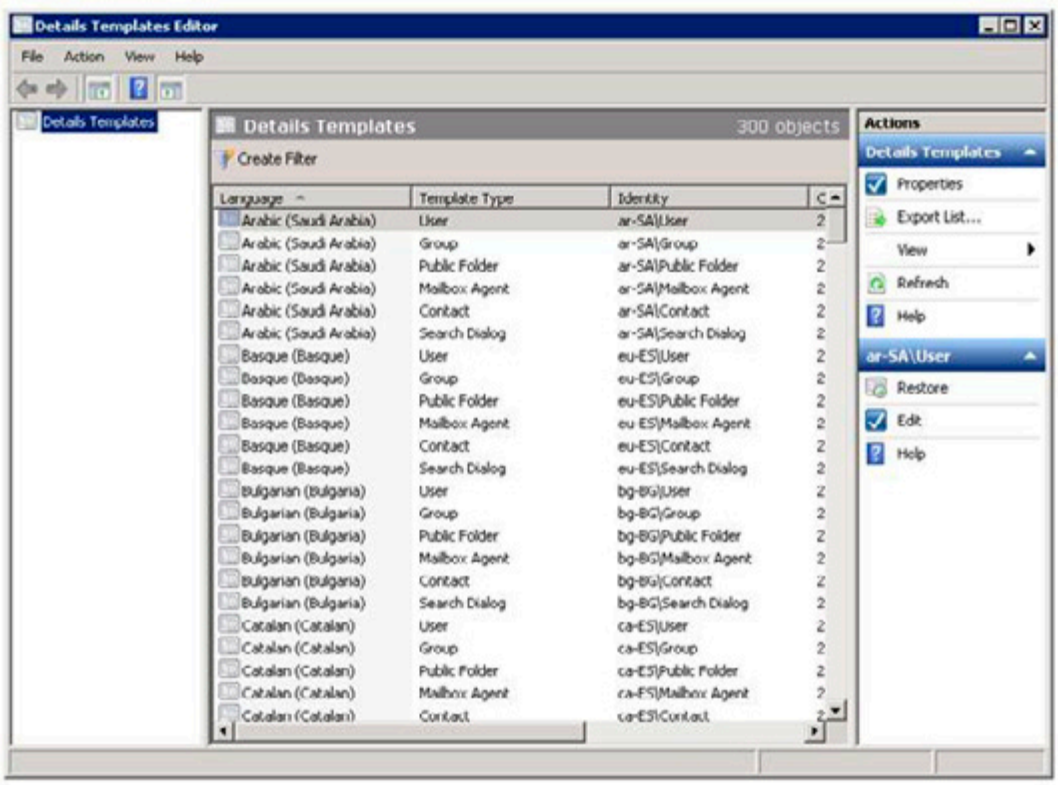

Figure – Exchange Details Template Editor.

Note

In order to use the Details Template Editor you need to be delegated the Exchange Organization Administrator role.

In my case I needed to edit the English User template – to open the template for editing double-click it and the template editor window is shown.

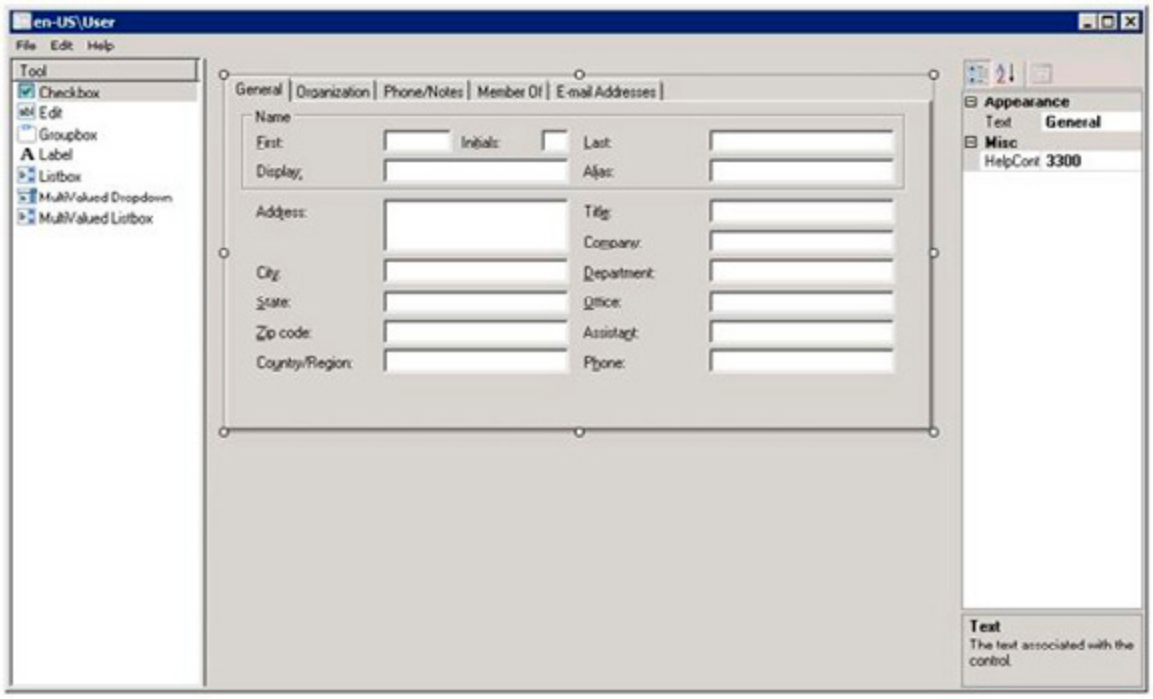

Figure – Editing the English User template.

The template editor is divided into three panes: the toolbox pane on the left, the designer pane in the centre, and the properties pane on the right. New fields can be added to the template by selecting the appropriate element type in the toolbox pane, placing the element in the designer pane, and linking it to an object attribute in the properties pane. The properties pane can also be used for fine-grained control over the size and position of elements as well as tab order and text size.

Note

The template editor does not include any undo functionality – you cannot undo any changes made in the editor, but the template can be reverted back to default settings. Templates are restored by right-clicking the template to be restored and selecting "Restore."

The object attributes which can be displayed on a template are limited to those provided by Microsoft. While it is technically possible to extend the set of attributes by modifying the Active Directory schema, doing so is not supported by Microsoft. If additional attributes which are not in the standard set are required the supported method of displaying the data is to use one of the fifteen Exchange extension attributes.

In my case I wanted to use the Active Directory attribute IpPhone, which had already been populated with the IP telephone numbers for our staff. As this attribute is not one included in the standard set I had to copy the data to another attribute which could be used. To do this I copied the data from the IpPhone attribute on each user record in AD to the Exchange extension attribute extensionAttribute1. The easiest way to do this in the Exchange Management Shell is with a short script.

This script will copy the value of the IpPhone attribute to the Exchange extensionAttribute1 attribute for all enabled user objects:

```
# Script to copy the IpPhone attribute the extensionAttribute1 attribute
# Written by Ben Lye - 07 December 2009

# Find the users in AD using an ADSI search method
$searcher = new-object DirectoryServices.DirectorySearcher([ADSI]"")

# Filter for enabled user accounts with a value in the IpPhone attribute $searcher.filter =
"(&(IpPhone=*)(objectCategory=person)(!(useraccountcontrol:1.2.840.113556.1.4.803:=2)))"

# Return the sorted results
$objects = $searcher.findall() | Sort-Object -Property cn

# Loop through all the objects that the search returned
ForEach ($object in $objects) {

    # Store some attribute values into variables
    $ipphone = $object.properties.ipphone
    $extensionattribute1 = $object.properties.extensionattribute1
    $dn = $object.properties.distinguishedname
    $adspath = $object.properties.adspath

    # If IpPhone is not equal to extensionAttribute1 then process this object
    If ($ipphone -ne $extensionattribute1) {
       # Get the ADSI object
       $adsiobject = [ADSI]"$adspath"

       # Set the attribute
       $adsiobject.extensionattribute1 = $ipphone

       # Commit the changes
       $adsiobject.SetInfo()

       # Output what just changed
       Write-Host $dn ":" $extensionAttribute1 "-->" $IpPhone
    }
}
```

Once the data is in an attribute which can be exposed via the details templates, then modifying the templates is relatively easy. An existing element can be re-labelled and linked to the new data, or a new element can be added. When modifying or adding elements, you must match the element type to the AD field type – single-valued elements (checkbox or edit box) must be used for single-valued AD attributes (such as primary telephone number fields), and multi-valued attributes (listbox, multivalued dropdown and multivalued listbox) should be used for multi-valued AD attributes (such as "other" telephone number fields). Mismatching single and multi-valued elements and AD attributes will result in data not displaying.

I decided to replace "Assistant" on the "General" tab of the user details with the new IP phone number. To complete my changes I opened the English user details template, changed the label text from "Assistant" to "IP phone," and changed the AttributeName property of the element from "ms-Exch-Assistant-Name" to "ms-Exch-Extension-Attribute-1."

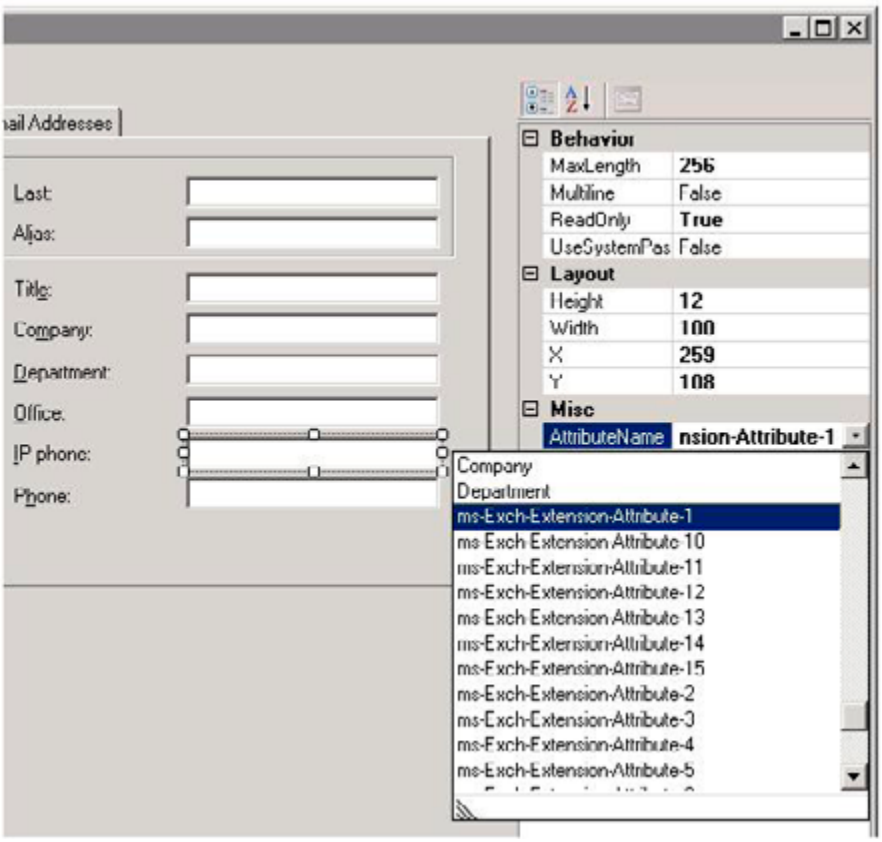

Figure – Editing the elements.

After saving the change to the template the changes are immediately available to Outlook clients which are not in cached-mode. Cached-mode clients will need to wait for the offline address book to be updated and a full download of the offline address book to occur (which by default in Outlook 2007 is only attempted once in a 13-hour period). The server update of the offline address book can be forced to run by running `Get-OfflineAddressBook | Update-OfflineAddressBook` and `Get-ClientAccessServer | Update-FileDistributionService -Type OAB` in the Exchange Management Shell.

Once the server-side update has run and the client has downloaded the new templates the changes will be visible in the address book.

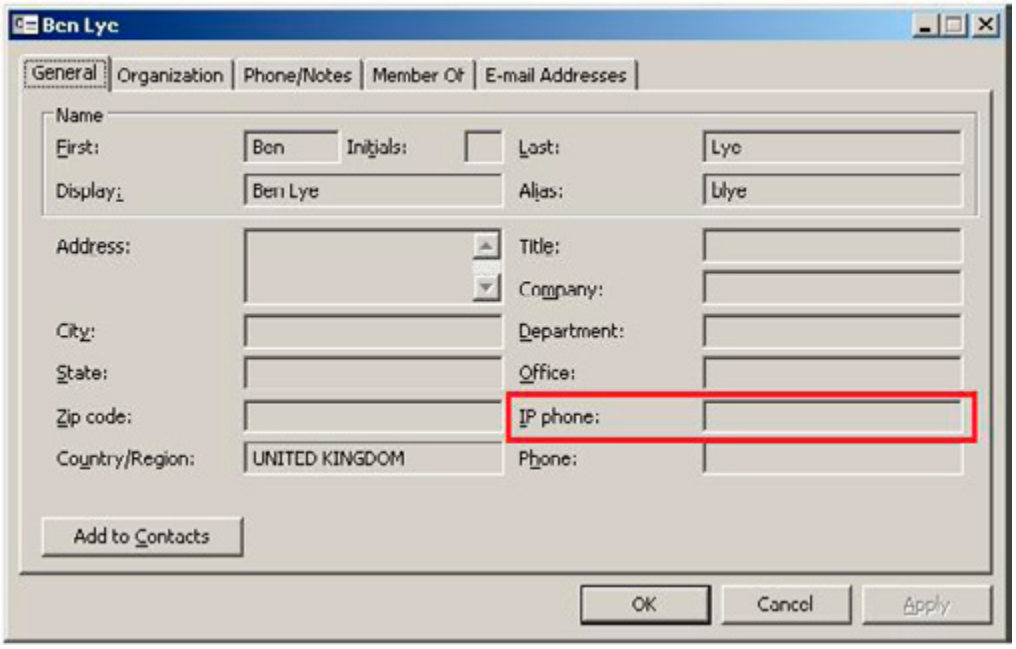

Figure – The new details template in Outlook.

Editing the details templates is relatively easy, but like all things Exchange careful planning will make implementing the changes much easier. When planning for template changes it's important to know what type of data you intend to add to the templates, and understand that while not all data in AD can be exposed directly to the Outlook Address Book there are workarounds available.

More information on customising the Outlook Address Book by editing the details templates is available in Microsoft TechNet:

http://technet.microsoft.com/en-us/library/bb124525.aspx

General Articles

A SysAdmin's Guide to Change Management

06 January 2009

by Matt Simmons

> ***In the first in a series of monthly articles, "Confessions of a Sys Admin," Matt describes the issues involved in Change Management, and gives a simple guide.***

As I sit down in New York's Grand Central station to write this, my first "Confessions of a Sysadmin" editorial for Simple-Talk, with people coming and going all around me, it feels only natural to tackle the topic of **change**, and how to manage it. Each of us has to manage change to some degree, in our working and personal lives, but I'd be willing to bet that whoever said that "change is the only thing that ever stays the same" probably worked in IT.

As a systems administrator, welcome change or shun it, change will happen. It is inevitable, constant, and unyielding. Failing to learn how to use change to your advantage would be foolish and dangerous; your competitors are surely going to. However, in order to benefit from change you need to know how to manage it effectively, and I'd like to review some of the key points and skills that will help you accomplish this goal.

Usually, the spur for change is the recognition that a problem exists in a process, followed by a "grand vision" for a better way to do it. Sounds easy, right? It isn't. I'll show you precisely how difficult it can be as we go through the steps of a successful change.

Early last year, the president of our company came to me with the news that our corporate headquarters was relocating from downtown Manhattan to central New Jersey. We'd also be opening an office in mid-town to keep a presence in New York City, and I was responsible for ensuring that there would be no interruption in IT services before, during, and after the move. Let's see how I handled it and where I could have improved.

Overcoming the Status Quo

Regardless of how great your idea is, you will encounter resistance. People in every organization are attached to the status quo, and will often resist your vision for change. Many are afraid of change in and of itself, others are afraid of losing their control over a system, and some find no profit from changing a system that, in their opinion, "works fine."

You need to create a sense of urgency around the problem, and the need to fix it, but at the same time it is incredibly important that the opinion of each member of the team is heard. If a person is resistant to your ideas, you need to do your best to find the root of their concerns and then assuage their fears. In my experience, sharing feedback and input from the whole team often helps people feel more comfortable with the change. Ideally, every member of the team should subscribe to your vision, although you will inevitably encounter members who consistently disagree with the consensus. Every voice should be heard, but you can't let a vocal minority tie the hands of the entire team. Examine the objection to see if it has merit, and if none exists, move on.

It goes without saying that "buy in" from the people with the power is vitally important to your success. A formal proposal may need to be created and presented to management to begin the process. While it is outside the scope of this article, tutorials on writing business

proposals can be found through your search engine, and once it's written, have others read and revise it for clarity.

In my case, the change was a mandate handed down from above. Rent prices in downtown Manhattan were increasing, and our growing company needed more room. Aside from the interpersonal issues involved with dividing an office, resources were finite. Employees were concerned that their new office wouldn't be equivalent to the previous one, or that they wouldn't be able to perform their tasks as well as they were used to.

I knew that the success of the move would judged in part by how comfortable everyone would be in their new locations, so I took that into account while explaining the various changes to the employees. What I didn't do enough was listen. I talked "to" (and at) people more than "with" them. Had there been more discussion, I believe everyone would have felt more like they were a part of the decision making process, which goes a long way toward accepting the transition.

Research: Mapping your Path

You now have a documented goal, and a team who share your vision. You must now analyze the goal, and carefully map your route from where you are not to your destination. Try to anticipate the roadblocks you will face. Thorough research is vital: you will often find that others have tackled the same problem before you, and you should learn from them. Internet forums, for example, are a great resource of knowledge and experience.

It is vital to have a central repository for the information you gather during the research process, to which every member of the team has access and can contribute. Some form of web based wiki would make an ideal resource, although many network-based project management solutions have this sort of thing built in.

As you start amassing information, ideas will emerge for how to put together a solution which meets the requirements of the goal. Hold regular group meetings and start to carve out a rough sketch of the solution's design.

I planned many aspects of the move, worked with vendors and contractors, managed quotes, and so on, but I personally don't feel like I communicated with management enough during this process. There must have been some amount of trust, as they didn't ask me about much beyond the occasional status update, and I was far too busy to volunteer information. If this were a complex project involving several other people, lack of communication would have sunk me.

Also, because I was the only member of my team, instead of a centralized repository on the network, I used a notebook. At the time I thought it was sufficient, but now I'd like to go back and see what I did, and of course I've misplaced the notebook among a stack of others.

Design and Scheduling

It is very common to go back and forward between the design and research stages, to the point that they may seem to blend together. Be careful not to lose momentum here. Revisions are inevitable, but progress is necessary.

As your ideas evolve and become more detailed, you'll find it useful to segment the goal, so that team members can specialize on specific areas. It makes the whole project more manageable, and having members work exclusively on a specific section increases their efficiency because they begin to "own" that part of the solution, and have an investment in its completion. Make sure that the smaller team reports regularly back to the main group. As the team leader, it becomes your role to ensure that they are making progress.

Certain sections of the task will likely depend on others that need to be completed first. Until you know what you're up against it will be very difficult to generate a schedule. Arrange the tasks in order of necessary completion, and have every team weigh in on the time their particular task is likely to take (all of this should be recorded in your central repository). This forms the basis of your development time schedule. Remember, however, that this is your *ideal* time. It would be wise to remember Hofstadter's Law:

"*It always takes longer than you expect, even when you take into account Hofstadter's Law*."

Documentation of requirements should be produced for each smaller team in order to verify satisfactory completion at a later date.

Always remember to schedule in sufficient time for thorough testing. Many people find that adding one third of the estimate tends to be relatively accurate.

Thought should also be given to the eventual support of the solution which is to be implemented. Does documentation need to be produced, does there need to be a new team created to handle this support? These are things that will need to be examined and decided upon as early in the process as possible.

My design phase went relatively smoothly. The hardest parts for me were keeping up with changes to the spaces we were moving into, and timing the various vendor/contractor interactions. Scheduling circuit installations around wiring-completion dates and phone-vendor delivery dates was difficult, and if I didn't have it all planned out ahead of time, it would have been a disaster.

What I did wrong was to misjudge the amount of network bandwidth necessary for the office to function correctly. I vastly underestimated when I should have relied on graphs. Had I done that, a last minute rush to order dedicated circuits wouldn't have been necessary, and people's initial impressions would have been much better. Rely on facts rather than hunches.

Development

Once the specifics of your design start to emerge, you'll need to begin development in order to keep your schedule. It's easy to become myopic during the development of the solution, particularly when working on a small segment of a much larger design. Through regular meetings, keep the teams aware of each others' status, and keep your individual groups moving towards the original vision and design.

On occasion you will, either through short-sightedness or lack of accurate information, discover that the design which was carefully developed is incompatible with the real world. You may need to fall back into design on a specific task or, if you are particularly unlucky, it may involve several smaller units. Bring together the involved teams and the original design team and work together for a solution. When solutions are updated, remember to update the requirements and user documentation which was produced beforehand.

As your teams approach completion of their specific tasks, care must be taken to ensure that every item is met satisfactorily, according to the agreed specifications. Only in this way can you be sure that the solution will effectively solve the stated problems and create a positive change. As the team leader, sign off when the team completes the requirements placed on them, and reminders may be required if they are neglecting one or more tasks.

It should go without saying, but every major change needs to be built in an environment that is segregated from the "real world," otherwise the usability of whatever you are improving will be at the mercy of your design and development processes. Build a complete solution and use it to replace the status quo, don't destroy the status quo while you build the solution, or you will have a vacuum in the interim and a lot of miserable people.

When you have a solution that satisfies your requirements, you need to create a schedule for replacing the existing system that will not create undue outages or break functionality. My advice is to work with your team to practice putting the change in place, then use that practice to calculate the required time, and multiply that by one third, as discussed previously.

For my migration, IT resources were required to be in place prior to the physical move. This required thorough testing of network ports, telephone access, and the like. As the spaces were independent of our inhabited office, there was sufficient time to test these resources. In a more time-sensitive operation, more manpower would have been needed.

Training and Support

After rolling out the change, you must provide support to the people affected. Providing them with documentation and personal assistance is well worth your time. People resent being left in the cold, and your change will be received poorly if they feel deserted or misused.

A group training session may be necessary and presentations should be delivered by a member of your team who understands the change and its consequences, and they should be prepared to answer pertinent questions. The goal is to make people feel at ease and comfortable with the change, and to reassure them that help will be available if they need it.

As we didn't have the resources to purchase matching phone hardware, normal business telephone lines were installed in one of the locations. This caused some confusion for personnel at that site, because I didn't anticipate the need to train them on these phones. To remedy this, I drafted and made available documents describing how to perform various functions using the available hardware, and reached out personally to the users to help them become accustomed to their new phones. I think that they appreciated the response, and we have had no further problems since then.

Project Evaluation

After a suitable length of time, it is advisable to evaluate the performance of the team and its leaders. Future performance can be improved by examining the mistakes that were made, correcting them and developing better strategy for next time. Review each team member's performance, and offer praise for appropriate conduct, or constructive criticism if it could have been improved. Be fair and honest, and the team members will respect you for it, even if they aren't happy with everything that you have to say.

We've been in the new offices for several months now, and things are going very well. Management follow-up on the move has been positive, and my users seem content in their new locations. Initial problems could have been resolved more quickly or eliminated altogether had I relied on statistics for bandwidth requirements rather than guessing, but overall the transition was smooth. I received kudos from our president for making the migration happen as well as it did, and I've now got several more upcoming changes to manage.

Summary

How can anything so simple be so complex? Change is more than just moving from one state to another. It is an ideology, a force of nature, and a tool for improvement. Change is the ability to alter the world around you. You are not powerless in your life, or in your work, despite what other people may tell you. Put change to work for you and reap the benefits. Form your goals, develop your plan, and implement the changes.

A SysAdmin's Guide to Users

26 February 2009

by Matt Simmons

> ***What level of trust do you afford users by default? What level of support do you offer users who have low technical skills? Are you creating a system which is more difficult to use with little or no payback? Are you adopting administrative policies as a punitive measure? do users believe that you are deriding them for their ignorance? Matt provides some timely advice.***

If management is the art of utilizing resources to accomplish goals, then administration might be thought of as the maintenance of those resources to ensure that they can be properly utilized. Administration of users is, for many, the most challenging aspect of their positions. In addition to the technical and logistical challenges provided by all resources, interpersonal relationships add a complexity that cannot be abstracted away in a script.

To help administrators cope with these issues, this column will explore various administrative tactics, how their effectiveness varies according to the traits displayed by their co-workers and we'll learn from examples of positive and negative administrative decisions.

There are many pressures placed on the IT administrators of the world. Corporate regulations, compliance requirements, and security best-practices compete for our time and resources. After meeting legal demands, there are sometimes few tactics available with which to approach administration. Many times, the tone and phrases used make greater differences in the perceived attitude than actual policies.

For our purposes, we can consider the laws and business requirements immutable. What we have left is really the core of administrative policies. What level of trust do you afford users by default? What level of support do you offer users who have low technical skills? Are the decisions you are making creating a system which is more difficult to use with little or no payback? These questions should be in your mind as you design system policies.

User Trust

In IT security, it is widely deemed to be the best practice to deny access by default. It has also become apparent that the most potent threat to an IT infrastructure comes in the form of insider access. This leads to the solution that users should have no unnecessary trust, be universally logged, and treated with general suspicion.

From the strictest security perspective, this makes sense. In practice, however, mistrust breeds mistrust, and you may soon find that the users who's lives are made difficult may find ways to return the favor to the administrators.

I propose that there may be a happy center ground. There are various occupations with extremely high turnover where extreme caution is called for. The same goes for certain financial, medical, and governmental institutions, but the majority of users and administrators find themselves under somewhat less pressure. In these cases, I believe that it is in the best interest of administrators to give the users leeway.

Many long-term technically oriented users have displayed responsibility in their decision making. Affording them administrative access to the machine that they use day in and day out shows an amount of reciprocity in appreciation. There is the chance that the user could install non-approved software, or through inattention allow a virus onto the machine. Proper end point protection software can all but alleviate the latter, and to be honest, the former probably isn't the end of the world. Obviously the situation is different if the software in question is illegal or opens an attack vector, but barring these complications, perhaps you should think twice before punishing the user for

attempting to improve their environment.

If you find that several users have all installed the same (or similar) unsupported programs, you should examine why it is unapproved. If so many people find it important enough to install, it may fill a niche and should be added to the support list. If not that particular program, then at least one which accomplishes the same task. Conduct non-confrontational interviews with users who installed the software to find out what it gave them. You may be surprised in the ways that your users are accomplishing their tasks.

I know that I have some problems sometimes when designing security policies. I tend to get carried away, and I start to resemble Mordak the IT preventor from Dilbert. Whenever I find this happening, I close my eyes, count to ten, and try to re-evaluate the situation from a new perspective. Sometimes, I have to stop working on the policy and sleep on it. This fresh view is enough to tell me that I'm being particularly retentive, or that the threat is warranted. I'll also ask advice from another admin. Getting the additional input helps a great deal.

User Support

It may help to examine how the typical administrator is viewed by the typical user. Obviously, all users are different, as are all administrators, but as a generalization, the relationship might be thought of as similar to that of a car owner and an automobile mechanic. The owner shows faith that the mechanic will properly maintain or repair the vehicle, and the mechanic makes their best effort to do so. Of course, there are mechanics who will abuse this relationship, as are there IT administrators.

No one wants to work with someone who will deride them for ignorance, and belittling someone is not an administrative tactic, despite how boneheaded some peoples" action might seem. On the other hand, coddling users and excessive handholding will ensure that a user who doesn't know how to do something will never have to learn. Where is the line drawn?

As each user is different, the approach taken must vary as well. Some users show eagerness to acquire new skills. Repeatedly performing a task for these users will do nothing but cause them frustration. Instead, demonstrate the task, and ask if there are any questions. The next time it is to be performed, be on hand to offer support if needed, but otherwise remain in the background. Offer encouragement and praise, or constructive criticism as necessary. This user's goal is to learn how to accomplish the task, and they will appreciate any efforts you make in helping them succeed. These users are valuable assets who generally have no problems expanding beyond their current tasks and taking on more responsibility.

Many people are not interested in employing skills which they deem unnecessary or time consuming. These people are very goal oriented, and want the path streamlined as much as possible. Attempting to teach the intricacies of a task to these users is seen by them as a waste of time. Success might be found in developing a solution which requires less interaction or is less time intensive. These types of solutions are generally an improvement for the business at large, as processes such as these consume everyone's time. Because of this, the goal oriented users become barometers for inefficiencies that, despite their reporting being mostly negative, instigate progress for everyone else.

Users also exist who learn and operate by rote. These users are typically very quick and efficient workers, as long as their environment is consistent. Large operating changes can unsettle users such as these, so a forewarning that a change will be implemented, along with a phone call or visit when the change takes place will ease them into their comfort zone. While the amount of personal attention required might seem overwhelming at times, remember that this is time invested. These users will not only pay you back with efficient operation, but they will act as red flags to system changes that even the keenest monitoring systems may fail to recognize.

This column would be incomplete without mentioning the infamous "bad user." Though lumped into one category, the reason a user earns this label varies a great deal. Either through actions or attitude, this user has shown that they do not want to contribute positively to the task at hand. Assuming that you do not have a supervisory role over this user, the best advice might try to make the best of a bad situation and work around them. I would advise against granting these users any more access than they need, regardless of technical prowess. Providing disgruntled users security clearance is asking for abuse. If the situation gets bad enough, it may warrant speaking to your supervisor about it. Under no condition should you speak with their supervisor, unless you both report to the same manager. The corporate hierarchy exists for a reason, and crossing boundaries will not endear you to the difficult coworker, their manager, or your manager.

Regardless of the tactic you are employing, the goal is to complete your task. Getting sidetracked into personal conflicts will deter you and hamper your schedule, not to mention lowering your concentration on the job at hand. Whenever interpersonal friction occurs, be the bigger person and move past it.

I've dealt with each of these users in my career, and I currently work with all of them with the exception of the "bad user." I have also mishandled each one of them at times. It was through trial and error that I found my mechanisms to best deal with them, as well as learning about the benefits that they all offer the group. I'm very fortunate to have my current group of users, even though I will most assuredly make mistakes dealing with them in the future. To live is to learn, as they say.

Administrative Policies

As a general rule, the policies that you write and put in place should exist to best enable the business of the company to proceed, while meeting the various legal and professional regulations.

Never adopt administrative policies for punitive measure, and do not allow your emotion to cloud your judgment. Policies provide guidance and protection, not punishment. A potential positive opportunity to educate users can easily become negative if handled incorrectly. A user or group singled out in a policy can quickly become bad users, which lowers the efficiency of the company, which in turn makes it more difficult to conduct business. Focus on preventing the undesired actions in the future and ignore the past when writing policy.

I think we're all guilty of retribution at some point in our lives, even if we're not proud of it. I'd be lying if I said that I didn't implement firewall rules at an ISP in direct response to user activity. The user didn't break the terms of service (TOS) in literal word, but in spirit, and I reacted to that by changing the firewall rules specifically to prevent them from doing what they were doing. What I should have done was discuss the situation with management, edit the TOS, and the firewall to reflect it. Fortunately in my job roles after that, I have reacted more maturely to issues such as these.

The Bottom Line

Managing users effectively is one of the most difficult tasks that your job can include. Experience is the best teacher, as long as you pay attention and learn from your mistakes. You don't need to be a "people person" to communicate well, you only need to try to see things from the others" perspectives and follow the golden rule. Treat other people like you would want to be treated. Respect and real communication will be your reward.

Change Management – What It Is and Why You Need It

02 March 2009

by Bill Holmberg

Bill Holmberg takes a practical look at Change Management, and list the steps you need to take. To make things simpler, he also provides a sample Word template you can use to follow through a complete change control cycle on a given project. Leaving nothing to doubt, he finishes with an illustration; a fictional exercise so you can see his guidelines in action.

What is Change Management?

Change Management, often called Change Control, is simply the process of managing most changes that can have positive OR negative effects on any environment- be it IT, Human Resources, or simply the way a small office works from day to day. Simply put, it is the act of pausing the decision making process before the implementation stage to check whether if it could be affecting more than the scope of the intended activity.

A simple example might be a decision to paint an office. A manager may tell an office worker to arrange to get the office painted, only to have the person schedule it during office hours, instead of on a weekend or overnight, with a consqequential interruption to the normal flow of productive work. A short change control meeting of the stakeholders – the office workers – could have avoided the work interruptions. Similarly, if the change control group also included the corporate office representative, they may have found out that the corporate headquarters had already scheduled a contractor to do all locations, and thereby received significant savings or ensured a standardized corporate "look."

Usually, **change occurs because a problem exists in a process**, and there is a clear idea for a better way to do it. While this sounds fairly easy, it often can be quite complicated and frustrating. Without a change control process in place, people will assume that any change is a good change – until it isn't – and that's when it becomes a problem, sometimes an expensive one, for companies.

Why Change Management?

Don't many organizations simply allow the various department heads to run their territory and then deal with issues that arise afterwards? Yes, many do, but the ones who actively manage the change process say that extra effort pays off.

Consider: Any organization that first stops implementation – until approval by a change management group – of any activity or purchase that affects productivity, system security, safety or budgets uses their hindsight far less because they exercised their foresight.

Change will inevitably happen.

However, in order to benefit from change we need to know how to manage it effectively, so perhaps we need to review some of the key points that will help us accomplish our goals.

On the flip side, when we actually have decided upon a change, we will encounter resistance. People in every organization protect the status quo, and will often resist any changes, no matter how slight they may be. Some folks are afraid of change in and of itself, while others are afraid of losing their control over a system or method – and some find no profit from changing a system that they see as working adequately.

In addition, some are threatened by automation in change, as a threat to their very livelihood, and will "silently sabotage" any efforts at improvement. This can be everything from playing dumb (can't learn it – it's too complicated) to actual overt and purposeful damage to a new system's data or hardware.

Creating a change control board

Item 1: Key People

This is a critical step in getting a results oriented group: Selecting the right people. The proper temperament of the individuals is important here. People that are averse to risk to a far degree will bog down the meetings, make them long and unproductive, and ultimately nothing will be done (or if so, will be done ages past when it is optimal). We call this "Paralysis by Analysis."

On the flip side, if the people involved do not have a knack for asking the right questions in their area, or are impatient and begrudging of the time they spend on these tasks, changes will be sped through without a real regard for process and critical thought. "Something must be done! This is something! This must be done!" leads to undesirable project outcomes.

Remember to clearly explain the role of the change control board to the prospective members: This is not where projects begin, it is where they become implemented. It is not for the change control board to be the final arbiter of what projects get blessed, but to see that the changes impacting the environment of the company are well thought out and understood.

In example, a marketing group may have made the decision to switch to a Macintosh based art creation system (or from one) for ancillary marketing and sales materials. It should be understood that the manager of this area has a deeper understanding of his departments needs, and those of his customers and employees, than the generalized people on the board. There may be good reasons to temporarily stall the project – say the network folks need time to look into implementing Appletalk across routers and it's impact on the Network – but this is not the time for someone's personal biases or opinions to be used to try to kill a project that has ostensibly already been blessed and funded.

Do you know what you are "looking" for?

That is not to say that some projects won't be found to be unviable as originally depicted at the change control stage – far from it, in fact.

A security plan to install cameras to allow remote security personnel to see the face of a person as they presented their corporate ID's was stopped by an HR person at a multi billion dollar corporation in the Midwest.
Why? Because the proposed project included a specific location of the cameras, which demanded that people bend over to place their badge onto a mirrored surface for reading while looking into the camera (They customarily wore them on lanyards around their necks).

When the installation was presented to the board prior to beginning the work the next month (which was designed, tested and created by the all male work force in the technical services department), they did a run through of the process using each of the members of the change control board – which included several women.

What was revealed was… revealing, to say the least. By assuming that particular posture, the security guards cameras were looking straight down the users necklines as they were admitted – which was noticed by all looking at the monitors when they watched each other go through the mock entrance they had made.

So, another project that needed some outside perspective before a potentially expensive re-modification was needed. Challenging our assumptions at the change control stage saves money, time, and potential embarrassment (or possibly even lawsuits!).

The Core Group

At the least, people should be available from the following departments (with some variety of expertise):

- Technical services
- Help desk
- Network
- Purchasing
- HR
- Marketing
- Sales
- QA/QC
- Security
- Maintenance
- Production
- Legal
- Specialty skills (In a trucking firm that would be mechanical, fleet ops, dispatch, etc. – in a Tech services firm that might be Lotus Notes developers, Oracle development or software assurance managers and Exchange Administrators or SAN managers).

In some companies it is not practical to have the entire set of personnel above meet on a regular schedule, so a core group may be whittled down – with an ancillary group, or alternates individually – available for consultation or attendance when warranted. Again, at the least these folks should be designated so that they are available for comment after the meeting where the change is first presented. In some companies certain individuals will fulfil one or more of these roles.

Item 2: Key Areas

Identification

The written procedure should cover the identification of the changed device, assembly, component, labeling, packaging, software, process, procedure, manufacturing material or any other related item or document. The change control form should have blanks for recording this data and other data discussed below.

Effective Date

The procedure should cover the effective date of the change – which is usually a completion date – or an action to be performed when a specific event occurs, such as "implement the change when the new router is installed, configured, and operational."

Responsibility

The change control procedure should state which department or resource(s) is/are responsible for each function to be performed. Be sure to specify which agency is responsible for the completed change control form. Also, be sure to point out if there is any extra level of management oversight required (and by whom, specifically) during the change.

Revisions

Which manner will be used to handle revision levels? It is common practice in many industries to use numerical revision levels during pilot production and planning, and transition to letters during implementation. This allows "At a glance" verification that you are on current doc sets.

Validation

Each changed device or process should be thoroughly tested or endorsed by the appropriate department(s). The results of the change – and all information related to the change – should then be reviewed by the change control board (or other designated review group).

Communication

The change procedure should cover the communication of changes to all stakeholders.

Updating Documentation

The change procedure should cover updating documentation such as network maps, user instructions, etc. It does little good to upgrade users to a new version of an office application if the intranet help resource still refers to old versions.

Documentation Distribution

Documentation should be distributed to persons responsible for the operations affected by the change and old documents removed wherever they may reside.

Follow-up Tasks

Change affects the system, and the system may need some tweaking in certain areas if it is to be beneficial in as many respects as is possible. The change control procedure should outline the steps required.

Regulatory Submissions

Some change may involve filing or even prior approval from governmental agencies. Modifications to devices or manufacturing processes should be made and covered under a QA/QC change control procedure.

Business Factors

Financial impact, the modification of marketing and sales collateral, update of products in commercial distribution, or a change in the way visitors or customers access systems should all be considered by the board before granting consent.

Quality Assurance Review

Any change also needs to be correctly implemented. Quality assurance personnel can ensure that the change is functioning as planned.

The Framework

The following are generally considered to be the proper steps or framework within which to study and implement changes.

Project Evaluation

Who will implement it and who is affected?

- What do we need? What is it worth to us? What does it cost?
- Why do we need it?
- When must it be in place?
- Where will it occur?
- How will it be done?

Challenge each assumption for due diligence.

The following areas need to be properly explored by the change control board.

- **Research:** Mapping your Path, planning the work, work the plan
  Does the proposed change seem adequately thought out? Don't be hesitant to send them back to get more data to ensure the proper research has been done to ensure the change has a manageable impact on the stakeholders.
- **Design and Scheduling:** Be sure what we are putting together goes together
  Check the assumptions – i.e., a Network that has Mac clients usually needs Appletalk on the routers for direct printing – is the proposed change of turning it off *really* what should be done?
- **Development/Implement:** Do the work
  Calculate the resources available versus time to implement the change. Are the proposed changes assumptions workable?
- **Training and support:** Give the users the tools and know-how to be successful
  The best change control boards are the ones that don't have to deal with unsatisfied users afterwards. Proper training on changes minimizes the grumbling.

Project re-evaluation

It is also a great idea to do a Post Mortem, or follow up on the change control items after they are implemented – whether successfully or not. Here are some points to go over afterwards.

- What variances from the expected outcomes were encountered?
- What did we do right?
- What could be done better for future projects?
- Were our initial assumptions correct?

Finalizing the documentation

Getting the docs together will save time and money when issues need to be addressed or the process needs to be duplicated elsewhere in the organization.

While at times this can seem an onerous or burdensome task, we really go through these steps anyway- just in an informal and often haphazard way. By forcing structure onto the process we can guarantee that we at least mitigated some of the negatives by ensuring as many as possible of the affected groups as is practicable are involved, and have the opportunity to voice potential issues or pitfalls- or even highlight some great advantages previously unforeseen.

Since it can be difficult, I have attached a pair of documents to this article. The first one is a template you can use to follow through a complete change control cycle on a given project. The second is an illustration of a fictional exercise so you can "see it in action," so to speak.

At its basic level, a great change control meeting can be as simple as bringing up a list of proposed items at a weekly meeting.

The attendees make a note to address any concerns they may have during the course of the next week, at which time they reveal their findings to the rest of the group, who all do the same. A short conversation is had on the topic, and a decision to go ahead, suggest more research of a specific type, or to table it until a future specific date is reached.

It doesn't have to be a long and drawn out, burdensome affair – in fact it can often be the highlight of a weekly status meeting, as the change control is the area at which the "Green Light" is often granted to a project, and things get done – or not – because of it.

Sample Change Management Documentation

Chewing the pencil?

No worries. Here with the article are two Word files. The first one is a template for a Change Control document, and the second is a Sample one made out to a fictitious project, using the template.

The Change control Template .

The Change control Sample Doc .

Manage Stress Before it Kills You

29 June 2009

by Matt Simmons

The key to a long career in IT is in learning how to cope adaptively with stress. Matt Simmons, like many, didn't realise it applies to everyone until he was rushed to hospital with what looked to be a heart attack. It was a salutary experience. He then discovered that there are many effective ways of unwinding, and it pays to be aware of them before you get the warning shot.

The warning shot

My heart beat against the wall of my chest. I could feel it pounding, even over my ragged breaths. I felt like I was suffocating, even though I was pulling in twice as much air as usual. I knew I was having a heart attack, and that realization was fuel for the fire which coursed through my veins. I made a decision. I was not going to die here, lying on a bench.

"We need to go to a hospital. Now."

The concern written on my friends' faces told enough for me to know that I looked as bad as I felt, and with their support, I made my way to the car. Shaun drove like a madman to the nearest emergency room. Immediately I was processed and evaluated, then wheeled into a staging area where I had electrodes placed all over my body so that the doctors could read the signals my heart was putting out. Blood was drawn, I was monitored, and cold stethoscopes were used.

It didn't take long for the results to come back. Less than an hour had passed before the nurse pulled open the curtain and stepped in.

"Mr. Simmons, we've determined that you're not having a heart attack."

I was stunned. Speechless, really. Every commercial I'd ever seen on television matched all of my symptoms. Every description I've seen before and since has matched the sensations I had that night. I had all of the outward signs, but none of the physical attributes of a heart attack. So what had happened?

All in the mind?

The answer, determined my cardiologist, was stress. Sure, it wasn't an easy diagnosis, or an immediate one. I underwent multiple stress tests, including one where I was injected with radioactive dye so that the three dimensional blood flow around my heart could be monitored by a machine the size of a small room. They tried very, very hard to find something wrong with me, but it was in vain. Structurally, I was fine. Mentally, I was stressed beyond the breaking point.

How is it that something that only exists in the indefinite recesses of a few hundred million neurons in my brain could wreak havoc and incapacitate my body in such a complete manner?

Stress management

It's because my stress management up to that point was nonexistent. My coping method would be to shove stress out of the way. Move past it, compress it into a little bottle and ignore it. That night, it refused to be ignored and it fought back. I should consider myself fortunate. The stress attack that night, and the tests that followed it, served as a warning. My mechanisms for dealing with stress were badly out of kilter, and unchecked, it could get worse. Much worse.

To investigate how bad it can get, lets examine how stress works.

Humans, like all animals, have a fight or flight reflex, keenly evolved over hundreds of thousands of years. Originally designed to protect us from physical threats, this instinct is now triggered by psychological threats as well. Anyone who has accidentally erased an entire directory's worth of important files can attest to this.

In addition to pure fight-or-flight stress, day-to-day situations apply mental strain and tax our body's resources. Work is a large source of stress, since many times we perceive our professional (and by extension, financial) statuses to be at risk. In addition, our personal lives and relationships are frequently a major cause of stress.

Stress is not, in and of itself, a bad thing. We're designed to be able to handle a certain amount of it. In fact, we require stress to operate at peak performance. Many highly dynamic people will tell you that stress provides them the fuel they need to accomplish their goals. Not everyone functions like that, though.

Stress has three phases on the body.

- The general phase, sometimes called the "alarm phase," is the initial response to the stress. Flight or fight. Some muscles tense, others loosen, and adrenaline enters the blood stream. In short-term stress, this phase passes quickly and our bodies return to normal.

- In longer-term stress, our body invokes coping mechanisms, and we start to adapt to the stress. This resistance phase taxes our body's physical and mental reserves, and the longer it goes, the more stretched out and weary we feel.

- The final phase is when our body is exhausted, and gives up fighting the stress. We're no longer able to resist whatever is causing our stress, and we stop being able to deal with it. The final stage is where we experience burnout. We're far past our peak performance and well into misery. Personally, when I get to this phase, I get irritable and hard to deal with. I snap at otherwise normal requests, I'm sullen, and generally not nice to be around. You probably know how you feel when you get to this phase, and I'm sure you don't like it any more than I do when it happens to me.

Extended periods of stress can do a lot more than strain personal relationships. It increases blood pressure, and over time, high blood pressure damages walls in blood vessels. When the walls heal, scar tissue is left, and these stresses cause blockages similar to cholesterol build-up. This is a very bad thing, because it can eventually cause a real heart attack. Fortunately, even the hard chargers and workaholics among us can lower their stress levels and live healthier (and longer) lives.

James Manktelow, author of "Manage Stress," advocates starting and maintaining a stress diary. A spreadsheet would be ideal for this: in the first column, rank your stress from one, being the least stressful, to ten, being the most stressful. Record what the stress was in the next column, and beside it, enter how the stress made you feel. Make a column for the number of times you feel that particular stress, and increase the number each time you experience it. In the far right column, record how you deal with that stress.

Sort your stress diary by the number of times you have experienced each situation, and make a note of the top entry. This line represents your most frequent source of stress, and gives you a target to focus on when removing stress from your life.

After recording your most frequent stressor, sort your diary by how stressful you rated each situation. When complete, the top entry will be that stressor which causes you the greatest trauma. Elimination of its source will also provide you the single greatest relief.

You should make it a priority to eliminate these two top sources of stress from your life. They cause the majority of damage to your health, and by eliminating them you will improve your overall stress level and well-being. Take action and use one of these three methods of dealing with them:

- **Action oriented** – Confront the root source head on. If your stress is work related, deal with it directly. Discuss the problem with your supervisor, explain how you feel, and be honest about it. Ultimately, if your employer is unable or unwilling to improve your conditions, leave. Your health is too precious to squander for a job that – in all likelihood – you don't like anyway.

- **Emotion oriented** – Attempt to change how you feel about the source of the stress. Personal relationships cause a lot of stress, but lots of problems can be effectively dealt with by trying to change how you feel about the problem, rather than changing the problem itself. Examine it from a different perspective, try to understand the underlying causes, and tolerate their effects. Attempt to be flexible in your outlook.

- **Acceptance oriented** – When all else fails, and you can't change the problem or how you feel, accept the source of the stress. Focus on surviving by building buffers between yourself and the source of the stress. Cope any way you can until you can improve the situation.

Stress is powerful, for good or ill, and effectively managing it is key to maintaining a healthy life. If you are having problems with stress in your own life, please contact a local stress counsellor or the human resources department of your company. Help is available for people who need it. You are able to defend yourself against the stresses of life. You only have to make the first step.

Hiring System Administrators

24 September 2009

by Matt Simmons

Hiring someone for a technical post is a task that should never be tackled half-heartedly. Matt Simmons provides some general advice on the subject, illustrated by a recent experience of hiring an IT administrator to help share the load. It opened his eyes to the real state of the economy.

It may seem that nowadays, when unemployment is high, and we're in a worldwide recession where millions of people are out of work and things may get bleaker before they improve, it may seem the wrong time to talk about hiring staff. Although many sectors are being hit hard, there are markets which have remained relatively unscathed, and some are even growing.

I believe that hiring should take place over three discrete, independent steps. Planning, Execution and Evaluation.

Planning

Tom Limoncelli, author of "The Practice of System and Network Administration" suggests that there are two types of job openings. One is for the position, the other the person. If you need someone to do a specific task, then you would hire someone who has the required skills. Other skills and experience that are possessed by that person are less relevant, because you need someone to perform the specified tasks, and it is essential to success that they are done well.

On the other hand, you may need someone to run a department, or to perform an array of less-well defined tasks. In that case, you would do better to hire a person who is capable of the flexibility that is required by such a role. Finding a qualified candidate is going to be less

clear-cut and will require more finesse from the hiring team. There is a good chance that the process will take longer as well, because it takes time to vet a candidate carefully. Your pool of applicants is likely to be smaller but more competitive, and salary requirements for this type of opening are generally higher.

Whichever sort of opening it is, one must define the position's parameters. Be specific about variables such as skill level, and ensure that you have explicitly listed the achievements and qualifications that you require from the candidates. This includes certifications, degrees, experience in a similar position, and even experience in an organization of a certain size if that is important to you. If these are outlined and decided upon early, the next step becomes much easier. Don't make the common mistake of assuming that any position seeking an inexperienced person is unimportant. Every position is vital, and should be treated as such.

In my own experience, pre-defining the range that you're willing to pay is a two sided affair, particularly if you are upfront about the range when you create the job posting. I consider it preferable to post the salary range, with the caveat "commensurate with experience." Unfortunately, being forthright in this manner has the drawback that individuals with more advanced skills may overlook your posting when they might be otherwise tempted to apply. While your budget is likely to be fixed, skilled candidates may be lured by other benefits: At the same time, it helps to prevent disappointment from overqualified prospects who apply if you dictate the range early in the process.

Determine who will participate in the hiring process and at what stages. There should be a defined interview team who remains somewhat constant throughout the process. This is essential to ensure that you compare apples to apples. Changing this team in the middle off the process may do a disservice to the candidates, who are almost certainly going to be viewed differently by different people. The number of people on this team, and the departments they represent, will vary according to your organization, but for positions that affect the company as a whole (and there are few who don't), buy-in from stake holders is essential.

Execution

Work with your hiring team to develop a job description that will be distributed to the various job boards and websites. In my case, the description simply consisted of the required skills and the responsibilities that the successful candidate would assume when hired. Many organizations use far more complex documents. Your Human Resources representative will be the best person to contact for more information.

Once you've got a completed description for the position, you can submit it to the job sites. but you need to make sure that you use the appropriate sites. The technical skill level of the ideal candidate should be taken into account. While an entry level candidate could be found easily enough on Monster.com or even Craigslist, a more precise search would improve the signal to noise ratio for highly specialized technical appointments. In the past, I've recommended the job boards on various professional organizations (IEEE, USENIX/SAGE, LOPSA) as well as high profile commercial sites. The serverfault.com job boards in particular attract very intelligent people, for the most part.

Once you commit your posting, applicants will begin to pour in immediately. Even in the best of times, it's possible to get deluged with applications, but with a large segment of the workforce displaced, the onslaught can be especially overwhelming. I would recommend that you stick to your guns and use the requirements you put together in the previous step (as well as any personnel you have at your disposal) to filter through the applications. How many are you likely to receive? Of course, it depends on your local area and the available position itself.

After the filtering has been completed, you presumably have a pool of candidates to select from, and chances are very good that you've got more candidates than you could reasonably interview in a short span.

I have several interviewing guidelines, but no real instructions on how to interview a candidate, because I feel like it's a very individual sort of endeavor. Your company's culture and your team will dictate the particulars of the interviewing process, just as mine did, but if I can pass along a few general tips, the first and most important would be to listen more than you talk. Ask open ended questions and see where the candidate goes with it. Don't interrupt unless they're floundering.

A great bit of advice I got from "Hiring Smart," by Dr. Pierre Mornell, was that a few minutes before the end of the interview, give the candidate a warning that their time is almost up. This gives them a chance to air anything that has been nagging at them, and can sometimes be very enlightening. Dr. Mornell related a tale of one applicant who waited until the five minute warning to let the interviewer

know that he was unable to take the job! As late as that is, it's certainly better to find out before the offer letter phase.

After the interviews have been completed, meet with the hiring team to discuss and decide upon a candidate. Review your interview notes, as well as the requirements you laid out in the planning phase. Once you've decided on a candidate, call them to congratulate them, and while you're on the phone, verify their physical mailing address to ensure that the offer-letter reaches them.

An Experience in Hiring

In part because of the state of the financial services sector and in part because of the organic growth of the company, I recently found myself in the position of needing to hire another IT administrator to help share the load. The hiring process opened my eyes to the real state of the economy.

I explicitly stated that this was an entry-level position, and that I was looking for candidates who didn't necessarily have administration experience. In spite of that, I had someone apply who had spent 20 years at their previous company. Many people spent over 5 years as senior administrators at their last company. That's the sort of thing that happens in this economy.

In my case, I knew that I was looking for a certain type of person. My budget dictated that whomever I hired had to be inexperienced. This made the selection process more difficult, as "systems administrator" would be a new role to most of my candidates.

I received several hundred applications in the span of only a week or two. It started only 5 minutes after the posting went live and didn't stop until a week after the posting expired.

I selected individuals having experience supporting technology in some measure and I sought out people who had experience using platforms similar to those my critical systems are built on. Most importantly, I selected those people who demonstrated a thirst for knowledge and for trying new things. It was from a pool such as this that I selected the junior systems administrator.

I emailed each of the potential candidates and scheduled a telephone interview with them. It wasn't a long, detailed interview; ten minutes was sufficient to tell me whether I wanted to bring a person in to interview them in person.

I was thankful that I took the time to conduct the telephone interview. There were a few cases where I talked to people who had impressive resumes, but upon questioning admitted things, such as, "well, no, actually I didn't implement that VPN connection, but I was there when someone from the vendor came and did it." These are things it's good to know before you waste several hours interviewing someone in person.

My interviews took, on average, two and a half hours. This seems like a long time, but it took about a half an hour before the candidates really relaxed and their personalities showed. At the same time, I wanted to extract a large sampling of their knowledge, and more importantly, introduce subjects new to them, to observe their learning process. Existing knowledge wasn't as important as the ability to acquire more of it, for my position.

I was upfront in the job posting with the salary range that we were prepared to pay for the candidates. When I performed the phone interview, I made sure to discuss pay with the candidates, and to get their thoughts on where in the posted range they fell. I was really surprised to hear that most of the candidates I talked to evaluated themselves very honestly, and for the most part, they suggested pay rates which were close to my mental target. During the interview, I openly discussed the question of pay with the candidates, because I wanted the major negotiation to happen in person. Your organization may have different customs and policies, so again, make sure to discuss your plans with Human Resources beforehand.

Evaluation

The hiring process does not end when someone is hired. It doesn't really end until the person's first annual review. All of the processes that we take part in can stand to be reviewed, and hiring is no exception.

Although this length of time hasn't passed for me yet with the IT Administrator, I intend to review the hiring process with my interviewing team after the new admin has been with us for a quarter. I want the emphasis to be on the process, and whether or not we succeeded in our goals, and what we could have done better. These sorts of occasional self-evaluations can lead to increased performance the next time I, or one of my teammates, is in the position to hire another employee.

Time will tell, but from what I've seen, everything is going to be great.

Increase Your Value as a Professional in the Technical Industry

17 November 2009

by Dr. Masha Petrova

> ***It has never been so important to enhance your employability as it is today. Job security can never be taken for granted. Employability, increasing your professional value, means far more than just collecting qualifications, as Dr Masha Petrova explains: It also involves communicating, writing, and participating in communities.***

How certain are you that you are keeping your job this year? Three or four years ago, if you lost your job and had a technical degree, it was only a matter of weeks until you secured another position, probably with even better pay. I remember that in my undergraduate engineering class, fellow mechanical engineering students with GPAs less then 2.0 (out of 4.0) and no internship experience were getting jobs in the industry right out of college. In contrast, this year when I gave a talk at the University of Delaware, students with graduate degrees and multiple publications were telling me that they were unable to get jobs and were going back to school to get their second Ph.D.

The world has changed in the past several years. Whether you are looking for a job in USA or UK, your options have become much more limited. These days, your technical degree, whether computer science or engineering, no longer guarantees your employment.

What can you do if you were swept up, along with many others, by the world-wide recession? What should be your plan of action if your one month layoff stretches to six months or a year? How can you make yourself more qualified and more appealing to future employers, while you're not technically working?

If you are one of the lucky ones to have kept your job, how can you increase your chances of keeping it? For those of you who need answers to these questions, here are some steps that you can take now, in order to increase your professional value, whether or not your are currently employed.

Join Several Professional Associations

In the world of IT, software and data management these are organizations like LOPSA (Association for Professional System Administrators), PASS (Professional Association for SQL Server), NWCET (National Workforce Center for Emerging Technologies), SIIA (Software and Information Industry Association) and the list goes on. If you are already a member of one of these associations, join one or two more. A few years ago it was enough to have a good resume and a technical degree in on order to get a job in the industry. Now days, you need to know people to secure employment.

Memberships with these societies will allow you to meet people in your industry, find out about job openings and professional opportunities, and possibly get recommendations and references. If you are unemployed, professional associations can provide some of the benefits of a work place – a technical community of like-minded individuals, a place to learn and to obtain referrals. To find out more about what associations you should join, search through the Internet Public Library (http://www.ipl.org/div/aon/). For USA associations, take a look at "National Trade and Professional Associations of the United States" book, which can be found at your local library.

Become an ACTIVE Member of those Associations

Serving on committees, chairing technical sessions, and volunteering for local events will provide you with lot of opportunities for networking and obtaining letters of recommendation. Being actively involved with professional organizations means that you might find out about new job openings while staying in touch with current research in your industry. Visit the websites of those associations, find a local meeting and attend it. The simplest way to get started is to contact your local association chair and let them know that you would like to get involved in some volunteering activities. Many of the association websites are starting to incorporate blogs to provide more valuable information to their members. Blogs are typically updated much more frequently than newsletters, so the associations are usually happy to invite guest authors to contribute to the blogs.

Why do you want to get involved? It's simple. The more people know you as a professional in your field – the more doors open up to you in terms of new jobs and professional recommendations. In addition, when you give away your expert knowledge in the form of articles, blog posts, and committee participation, you are building your credibility as an expert. If you keep at it, companies might actually hear of you through your professional community work, before you even apply for a job!

Join an On-Line Professional Community

This step is easy and, at the same time, can be extremely beneficial for helping you land your next job. My preference for a professional on-line community is www.linkedin.com. I have gotten everything from answers to technical questions, to referrals for web designers, to recommendations of my consulting work from this site. If you are already on Linkedin, become an active member – join discussion groups or create your own group in your area of expertise. As a word of caution, there is a wrong way for using an on-line community and that is to post your resume along with a desperate call of "I am looking for a job!" on every discussion board. Everyone is looking for a job these days. Your resume blasts are not going to impress anyone enough to get you an interview.

Give Talks

As professionals in a technical industry, we are often taught that explaining things in a clear and understandable manner should not be high on our priority list. In fact, in many cases, we are taught that the more confusing we sound, the more we will be perceived as intelligent and important individuals. When was the last time that you attended a talk at technical conference where at least a few of the attendees were not fidgeting with their laptops or just plain asleep?

The ability to present complex technical ideas in a clear and understandable fashion is an extremely valued skill in the corporate world. It is also a skill that is immediately apparent in an interview. Learning to be a good communicator and speaker will put you miles ahead of other IT professionals searching for a job.

A good way to start to develop those communication skills is join Toastmasters International: www.toastmastersinternational.com. This is a wonderful organization that can help develop your public speaking, communication and leadership skills. For a nominal yearly fee, you can attend weekly meetings, get feedback on your speaking ability from other members and, most importantly, get plenty of practice speaking in front of a group of people.

Write

Finally, I recommend that you develop your written communication skills in the same way that you develop your public speaking abilities. To help you get started on improving you writing skills, here is another very simple idea. Create a blog and contribute to it every week. It is very easy to set up and it's free. Visit the Google blog service, www.blogspot.com to get started.

Here is another idea. You are reading www.simple-talk.com an on-line journal for SQL Server and .NET developers and administrators. You might notice that there are a number of other authors in this journal whose article you enjoy reading. Why not become one of them?

A more advanced step for those of you, who have built an established blog and feel comfortable with writing, is to find on-line journals or other blogs in your area of work and ask if the editors might be interested in some of your article ideas. They very well might be, in which case, I look forward to reading your piece in Simple-Talk in the future!

You are an expert in the technical field. You have excellent problem solving abilities. As do most of the other software scientists, system administrators and IT professionals looking for a job. Now is the time to start developing the other aspects of your professional skills if you want to stand out of the crowd job seekers. Develop your professional relationships, evolve your communication skills and become more involved in the technical community. Do this, and this recession might just turn out to be the best thing that ever happened to you.

For more information on Dr. Petrova's Steps to Increase Your Professional Value, please visit: www.SuccessfulUnemploymentToolkit.com.

Images 1 and 2 courtesy of Pat Wright taken at the 2009 PASS Summit.

The Art of Dealing with People

15 December 2009

by Dr. Masha Petrova

> ***Technical people generally don't easily adapt to being good salespeople. When a technical person takes on a customer-facing role as a support engineer, there are a whole lot of new skills required. Dr. Petrova relates how the experience of a change in job gave her a new respect for the skills of sales and marketing.***

A few years ago, I was sitting in my cubicle staring intently at my computer screen. I was working for an engineering software company, and have been trying to figure out why a subroutine in the software that my company was developing, kept crashing. I mostly worked on numerical engineering research, but every once in a while, my projects would involve writing and de-bugging code.

I enjoyed those times because I could sit uninterrupted in my cubicle for hours at a time, living in my own little microcosm of numerical algorithms. But that day, my peaceful bubble was about to burst. The Sales Guy infiltrated the personal space of my cubicle. "Dr. Petrova!! Heeeeeey!! How are things?!!" – he roared with such enthusiasm, you'd think he just downed a whole bottle of Prozac.

Every few days, the company Sales Guy would make his rounds through the engineering department, interrupting work, shooting the breeze, and in general spreading the blessings of his persona throughout the R&D department. We would all politely smile, answer his questions, and try to make small talk, while secretly wondering how the heck the Sales Guy manages to get paid for doing absolutely nothing.

As it turned out, there was quite a bit to learn from the Sales Guy. Fast-forward three years. I have joined the sales and marketing team as a support engineer because I discovered an extrovert side in me that enjoyed working with people. The problem was that all of those years in front of the computer screen, gave me little time to practice my people skills. Those skills happened to be crucial if you need to work with, well, people.

Even if you are not planning on working in sales, mastering the art of dealing with people is essential to any professional career. Does that mean that you have to force yourself to make small talk with everyone at your company instead of finishing a project before a deadline? Absolutely not. There are much classier and more comfortable ways to bring out the extrovert in you. I would like to share some of the lessons that I learned from the Sales Guy in my journey of mastering the art of dealing with people, adapted for technical professionals.

Lesson #1: Get a Mirror

After the first couple of weeks of working for the "dark side" (clever term bestowed upon the sales and marketing by the engineers), the Sales Guy, and now my boss, walked into my office. "Masha," he said lightheartedly, "Do you often make people mad?"

By that time I figured out that sales people have a different way of talking to you, so I was not instantly offended.

"Not that often," I said reluctantly. "Why?"

"Well," said the Sales Guy, still beaming, "Do you know what your face looks like when you talk to people?"

He stumped me with that one. That was a weird question even for a Sales Guy. "It looks like…I am…in deep thought..?" I said uncertainly.

"It looks scary," said the Sales Guy, smiling as if he was watching baby bunnies frolic in the meadow. "If you want to have any hope of relating to people, get a mirror and put it on your desk. When you're talking to anyone who comes into your office, glance into the mirror and see what the other person is seeing."

With that he was gone and I was left wondering why in the world I left my cozy engineering R&D position for this. In the spirit of learning, I decided to give the mirror method a shot and was amazed at what my face was telling people. The first time I glanced into the mirror, while I was talking to a coworker, I saw the face of an axe murderer staring back at me.

Up to that point, I was sure that no one cared about what my facial expressions looked like, but once I started using the mirror and adjusting my facial expressions, the results were astonishing. People were more willing to talk to me, give me needed information and in general seemed to like me better. Being aware of what your body language and facial expressions are saying is an important first step in relating to people.

Lesson #2: Take them to lunch

One of the first things that the Sales Guy recommended I do in order to develop my extrovert skills, was to start taking people to lunch. That turned out to be a great learning experience for me. Many of us working in the technical fields perceive spending a lunch hour with people who are not our immediate friends as an hour wasted. In reality, you can learn more in that one hour, than during many more hours in front of a computer.

The first few lunches with various people in my company were a bit challenging, but practice makes perfect. It became easier and aside from learning more about each coworker, I was able to gather information that was useful to my job as well as establish myself as someone who was genuinely interested in people. That, in turn, encouraged others to help me, when I needed information or resources. I was forced to bring out the extrovert in me during each of those lunches and before I knew it speaking to random people at conferences and professional events was a snap.

If you are an introvert looking to grow yourself as a professional, one of the first things you should to be working on is training your internal extravert to come out at your command. Asking various people at your place of work to have lunch with you is a great way to start. It might be uncomfortable and probably even scary at first. I recommend starting with some fellow programmers, IT people, or engineers at your company. Chances are they will be even more introverted than you, which would make for a comfortable and quiet lunch hour.

Then move on to the sales and marketing guys. Take this to be an exercise. You don't have to try to be their best friend, but you are looking to learn how to best deal with different types of people.

Lesson #3: It's not your Thesis Defense – stop proving how smart you are

During the second week of my sales career, the Sales Guy told me that I will be sitting in on a call with a potential client. I was terrified and thrilled at the same time. Finally, a chance to showcase my brilliance!

During the call I tried to answer every question the potential client had in gruesome detail. I was so proud of how well I was doing. The Sales Guy kept shooting me dirty looks, but I figured it was because he was just overwhelmed with my knowledge. The call ended and I was thrilled. "That was a great call!" I said.

"Yeah," mumbled the Sales Guy, "We'll never hear from him again." "Why?" I said, appalled.

"Tell me something," said the Sales Guy," How big is this engineer's group?"

"I don't know. He didn't say."

"Can you tell me what the role of this engineer in his department is?"

"Well, I didn't ask."

"Can you at least tell me if they can afford our software?"

How am I supposed to know that?" I said irritated. "He didn't tell me!"

"Interesting," said the Sales Guy. "You just spend an hour of company time on this call and did not gather a single piece of information that would actually help us make a sale. In addition, the engineer doesn't think that you care about him or the problems he is trying to solve."

"Why?" I was taken aback, "I gave him a lot of information about our software!"

"Yes. You took up his time to demonstrate how smart you were instead of asking him questions that would show him that you are interested in helping **him** solve **his** problem."

The engineer never called back. Trying to prove how smart you are during conversations makes people not want to talk to you. People like to feel that they are being heard. Ask a question and then be quiet and just listen. Figure out a way to be interested in what others are telling you. You might have to fake it in the beginning; especially if you are not that interested in a particular conversation topic. Beware, faking it all the time does not work, others will be able to tell and will be offended.

If you would like to better master the art of dealing with people, take some cues from my Sales Guy. Be aware of what your face and body language is saying to people, practice bringing out your inner extrovert, and stop focusing on your own brilliance. With those skills you will be well on your way to a more prosperous career.

Virtualization

Virtual Exchange Servers

20 November 2008

by Jaap Wesselius

> ***Microsoft now supports running Exchange Server 2007 in server virtualization environments, not just on Hyper-V, but on any virtualizing solution that is validated in the Server Virtualization Validation Program. Before virtualizing Exchange Server, you will need to be clear about the business advantages, and consider the size of the installation, otherwise the results can be a disappointment***

Virtualization is a hot topic these days. Microsoft has released Windows Server 2008 Hyper-V just before the summer and Hyper-V is Microsoft's answer to the supremacy of VMware ESX Server. Customers have been running applications like Microsoft Exchange Server on VMware for years now, despite the formal support statement from Microsoft that this wasn't supported. Okay, there was a "commercially reasonable effort" support, but that was basically it.

This has changed since Hyper-V was released. All server virtualization products that are validated in the SVVP (Server Virtualization Validation Program) are now fully supported by Microsoft. Needless to say that all major virtualization vendors have signed up in this program, so all recommendations regarding virtualization are not only valid for Hyper-V but for all vendors in the SVVP program.

Microsoft has a list of all Microsoft applications that are supported in a Virtual Machine, this list is published via kb article 957006 (kb article 957006). There are some remarks in this document, and one of them is for running Exchange Server in a virtualized environment. And that's exactly what I want to talk about.

> ***Note***
>
> *The official Support Policies and Recommendations for running a virtualized Exchange server can be found here:*

Windows Server 2008 Hyper-V

Hyper-V is a hypervisor product. A Hypervisor is a very small software layer that's situated between the Operating System and the hardware. In terms of "small," Hyper-V is less than 1 MB in size.

Install Windows Server 2008 (X64 only!) on "Designed for Windows" hardware. Make sure that this hardware supports hardware virtualization and have a correct BIOS version (including the Execute Bit Disabled option). All Class-A servers should be capable for Hyper-V, but be careful with the low-end servers. I have seen budget servers with Pentium IV processors that are not capable of running Hyper-V! But I have also seen laptops with for example an Intel Core2Duo processor that are capable of running Hyper-V (don't expect a real-world performance though). The server should have enough internal memory and disk space to facilitate the use of Virtual Machines.

After installing Windows Server 2008 you have to install Hyper-V. Hyper-V is a server role and can be installed using the Server Manager. The Hypervisor slides between the Operating System and the hardware, and after a reboot the system is ready. Note that the original Operating System has now become a Virtual Machine as well! This Virtual Machine is called the root or parent partition. This is a special partition since it controls all other Virtual Machines on this host. These VM's are called child partitions. Special care should be taken for the parent partition, and no applications should be installed on it. The best solution for the parent partition is to use Windows Server 2008 Server Core which only has a command line interface. There are a few graphical tools like timedate.cpl for setting the timezone and time/date information, and there's also Notepad for creating batch files. But the general interface in the command line. Windows Server 2008 Server Core has low overhead, it has a small memory footprint and it has a small attack surface. It is more difficult to manage though, especially for the average Windows administrator.

There are some additions to the host Operating System when Hyper-V is installed. A VSP (Virtual Service Provider) is installed. This is a piece of software that gives a Virtual Machine access to hardware resources. The VSP is connected to the VMBus, an in-memory bus used for communications between parent and child partitions. Every child partition has its own VMBus for safety reasons.

Also VMWorker processes are installed on the parent partition. These are used for non-native Hyper-V Virtual Machines, they cannot use the VMBus or VSP interfaces. Hardware resources are emulated on the host in the VMWorker processes for these types of Virtual Machines.

Native Hyper-V Virtual Machines offer the best performance, whenever possible try to use native Virtual Machines. Supported server operating systems running as a native Virtual Machines are:

- Windows Server 2008.
- Windows Server 2003 SP2 and higher.
- Windows Server 2000 SP4.
- SUSE Linux 10d.

To integrate the native Virtual Machine with Hyper-V special drivers need to be installed in these Virtual Machine. These drivers let the Virtual Machine use the new VMBus structure as can be seen in REF _Ref212978047 \h Figure 1. These drivers are called Synthetic Drivers and are part of the "Integration Components" of Hyper-V.

Other operating system run very well in Hyper-V as well, but they use the hardware emulation instead of the new VMBus structure.

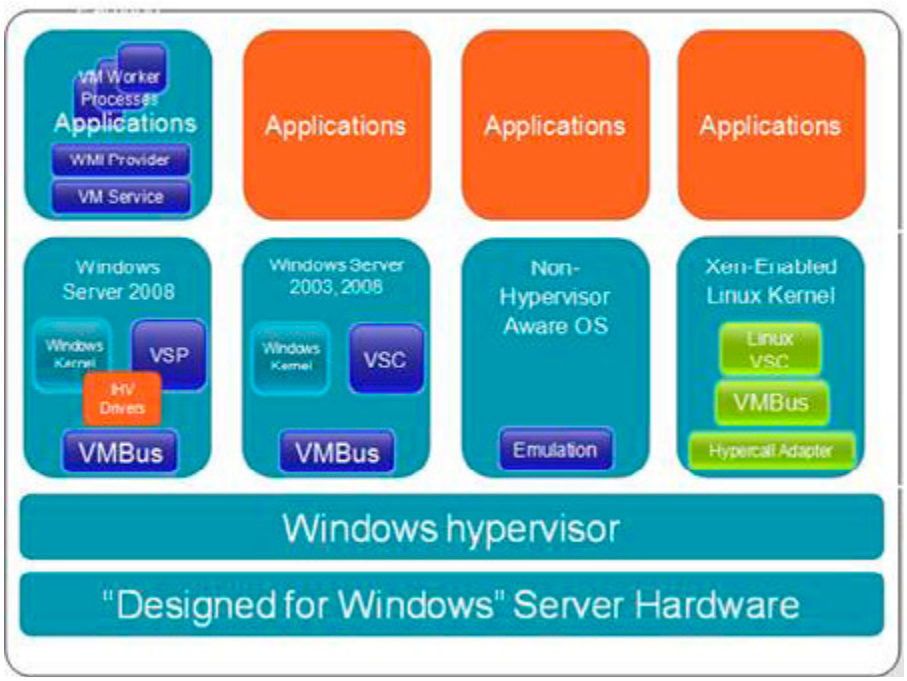

Figure 1. Windows Hypervisor structure with multiple VM's. Note the Emulation part in the 3rd VM.

So, the original Windows Server 2008 image becomes a Virtual Machine as well, but this one should not be used for production purposes. For running Exchange Server in a virtual environment we need to install a new "native Virtual Machine" that's capable of using the new VMBus structure. Never install Exchange Server in the parent partition.

Virtual Exchange Servers

That being said, we want to create a fully supported virtualized Exchange Server environment. As stated before, the official support policies can be found ON THE MICROSOFT WEBSITE.

The first thing is that the only version supported on Hyper-V is Exchange Server 2007 SP1 in combination with Windows Server 2008 as the Child Operating System. This is mainly because Windows Server 2008 performs better than Windows Server 2003 using multiple virtual processors. The Exchange product team didn't want to spend any time on testing Exchange Server 2003 running on Hyper-V, so the official standpoint didn't change here. I'll refer to that later in this article.

Virtualizing Exchange Server is like designing a bare metal Exchange Server environment, it's all about performance. All design guides available for designing bare metal Exchange Server 2007 environments should be used for virtual Exchange Server 2007 environments as well. This goes for all Exchange 2007 roles, except the Unified Messaging role. This role is not supported running under Hyper-V due to the real-time speech recognition in the Unified Messaging role.

An Exchange Server 2007 running under Hyper-V can look something like this:

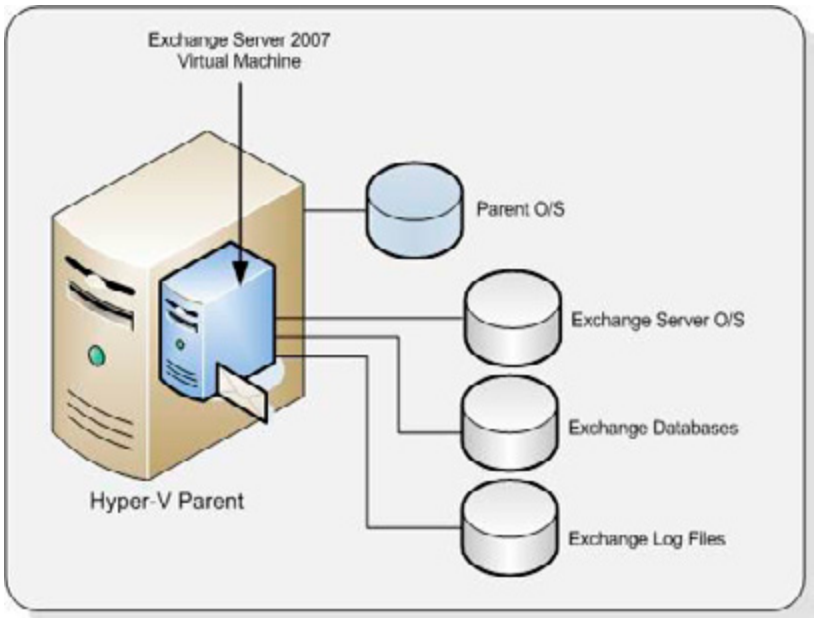

Figure 2. Exchange Server 2007 SP1 runs on its own disks.

Before installing Windows Server 2008 as the child partition, a Virtual Hard Disk must be created. This can be fixed size Virtual Hard Disk (.VHD file), remember that dynamic .VHD files are not supported for Exchange Server 2007 SP1. A better solution is to use a dedicated or pass-through disk. This is a dedicated disk that can only be used by the Virtual Machine. Tests have shown that a dedicated disk under Hyper-V has a similar performance characteristic as a native disk in a non-virtualized Exchange Server.

Our server has two physical disks. Disk one is used for the Parent Partition which is of course running Windows Server 2008. The second disk will be dedicated for our Windows Server 2008 Child Partition where we install Exchange Server 2007 SP1.

First bring the 2nd physical disk offline, you can do this in the Disk Management option under Storage in Server Manager. By placing it offline it is not available anymore for other operating systems.

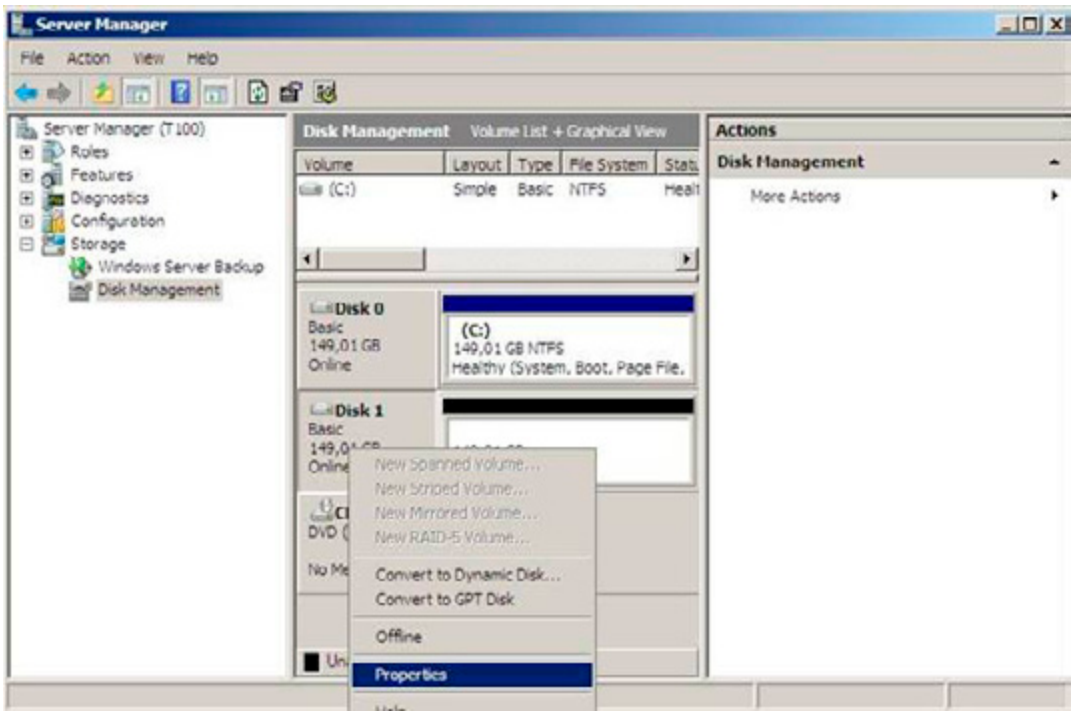

Figure 3. Bring the disk in offline mode in Server Manager.

- Now create a new Virtual Machine:
- Open the Hyper-V Manager in the Administrative Tools menu.
- In the tasks pane click New and select "Virtual Machine." The New Virtual Machine wizard starts and after the welcome screen you have to enter all kinds of information regarding the new Virtual Machine:
  - specify a name and location to store the Virtual Machine.
  - assign it a certain amount of memory, for example 4 GB.
  - bind it to a publically available network.
- At the "Connect Virtual Hard Disk" windows select the "Attach a virtual hard disk later" option.
- Specify the ISO that will be used to install Windows Server 2008 x64;
- After reviewing the summary screen click Finish

Do not start the Virtual Machine at this point.

When finished open the properties of the Virtual Machine and go to the IDE Controller 0. This is the controller where the boot disk should be attached to. Check the "Physical Hard disk" option and select the physical disk we've put offline earlier. Click OK to save the configuration.

Unfortunately it is not possible to add a SCSI controller and the disk to this SCSI controller. The Hyper-V SCSI controller is part of the Integration Components and thus runs against the VMBus structure. A very flexible solution, but only available when the Virtual Machine is up and running.

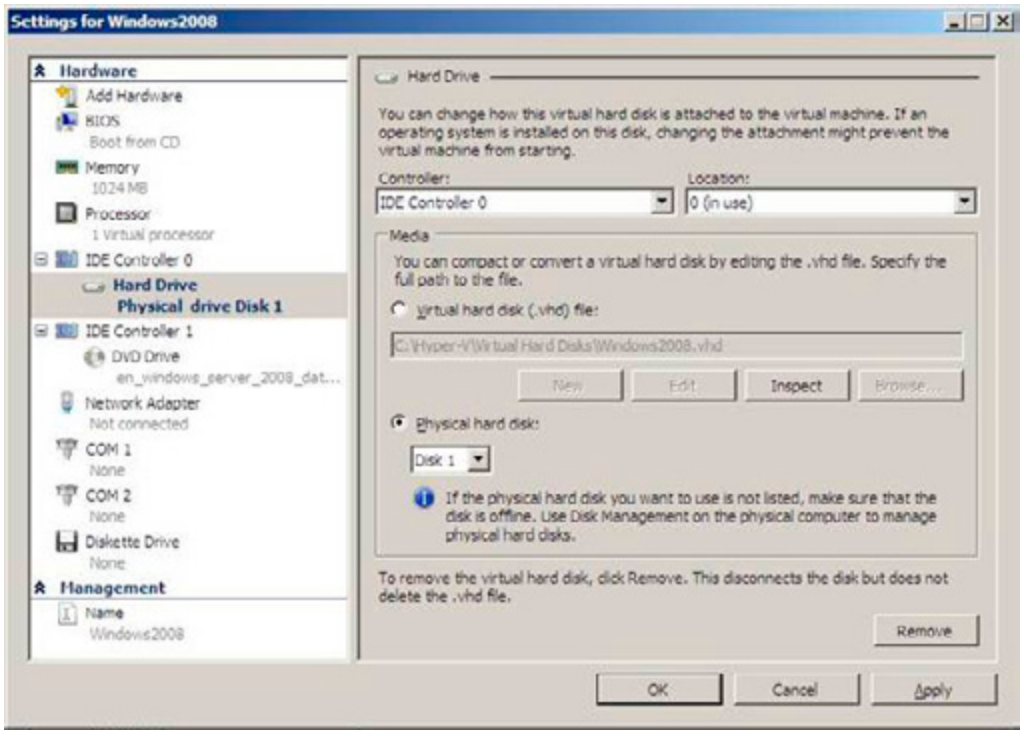

Figure 4. Add the physical drive to the IDE Controller 0.

The new Virtual Machine is now ready to be installed, just power it on and install Windows Server 2008.

After installing and configuring the server the server needs to be updated. The RTM version of Windows Server 2008 comes with a beta version of Hyper-V and this needs to be updated as soon as possible. When updating the Virtual Machine make sure that you install Microsoft hotfix kb950500, this hotfix brings the Virtual Machine to the RTM level of Hyper-V and thus includes the latest version of the Integration Components.

Installing the Integration Components in a new Virtual Machine can be challenging. Before the Integration Components are installed there's no network available, and without a network it's impossible to copy drivers and updates to the Virtual Machine. It is a good idea to create an ISO image with the most important updates, including the Hyper-V RTM hotfix and install this on the host server. This way you can always mount the ISO images and install the latest Integration Components.

After the Installation Components are installed and the server is brought up-to-date we have a fully functional Windows Server 2008 server. Assuming that you already have an Active Directory available on your network you can continue installing Exchange Server 2007 SP1.

Exchange database and Log Files

An Exchange server running on a server virtualization platform doesn't differ from a bare metal Exchange server and as such you need to place the Exchange database and the log files on separate spindles. This is both from a performance perspective as well as a disaster recovery perspective.

There are three ways to configure the database and the log files:

- Use a fixed Virtual Hard Disk – This is a preconfigured .VHD file with a fixed size placed on a separate disk.

- Use a dedicated or pass through disk – this is identical as the disk we just installed Windows Server 2008 on. However, since we already have a running Windows Server 2008 Virtual Machine we can add a SCSI controller to the Virtual Machine and attach the pass through disk to the SCSI Controller. This is the preferred and recommended solution.

- Use iSCSI LUN's – using the Windows Server built-in iSCSI initiator we can access LUN's on a storage device and place the database and log files on separate LUNs. Although using iSCSI within the Virtual Machine is a fully supported configuration the performance is less than exposing an iSCSI LUN from the parent partition as a dedicated disk. This is due to the networking overhead within the child partition.

Again, there's no difference between designing a bare metal Exchange server and a virtual Exchange server. Always design your Exchange server with the best performance in mind!

It is not yet possible to use a fiber channel solution with Virtual Machines natively. The HBA's (Host Based Adapter) are not yet available for use with the VMBus structure. HBA vendors are working on this however, but it is unknown yet (as of October 2008) when this will be available. It is possible however to use a fiber channel solution on the host system and expose LUN's on a SAN as disks that can be used using the pass through mechanism.

Backup and Restore

As explained in earlier articles on http://simple-talk.com backup and restore is very important on Exchange servers. Running Exchange servers on Hyper-V make backups look very easy, just create a backup of your Virtual Hard Disk and that's it. Microsoft Server 2008 even supports VSS (Volume Shadow Copy Service) backups of the Virtual Hard Disk files.

Although this is true you still have to be very careful with backing up your Exchange Server under Hyper-V. Not only do you have to backup your data, also the Exchange server's database needs to be checked for consistency and the Exchange server's log files need to be purged.

The Hyper-V VSS writer that's part of Windows Server 2008 is communicating through the Integration Components with the Exchange writer in the Virtual Machine. Windows Server 2008 Backup (installed as a separate feature) can create VSS backups of the Virtual Machine running Exchange Server 2007 SP1. When checking the Exchange Server after a Hyper-V VSS backup is created the Exchange database header show backup information and also the log files are purged.

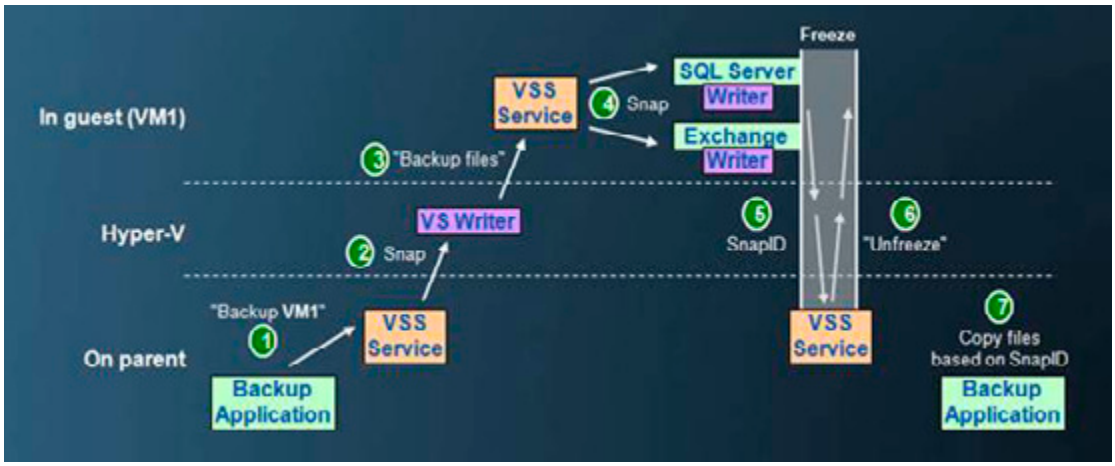

Figure 5. Creating a Hyper-V VSS backup does interact with the Exchange server running in the Virtual Machine.

Although it is fully functional, it is not a very user-friendly solution. Microsoft is also offering a complete backup solution, Microsoft System Center Data Protection Manager (DPM) 2007. With the upcoming service pack 1 release of DPM Microsoft is going to support VSS backups of Virtual Machines. The expected release of DPM 2007 service pack 1 is in the first quarter of 2009. 3rd-party vendors like Symantec offer VSS backups of both Hyper-V as well as VMware Virtual Machines in BackupExec 12.5 which was released early October 2008.

At this moment (November 2008) the recommended way to backing up your virtualized Exchange Server 2007 environment is within the Virtual Machine itself. Install a DPM or other 3rd-party backup agent in the Virtual Machine and back it up from there.

Snapshots

A snapshot is a point-in-time copy of your system, in this case the Virtual Machine. After creating a snapshot of our Virtual Machine it is always possible to return to the state of this Virtual Machine at the moment of creation of the snapshot.

The following takes place during a snapshot:

- A copy of the configuration file is created.
- The Virtual Machine's memory is flushed to disk.
- A differencing disk is created and the original disk is set to read-only.
- The Virtual Machine resumes operation.

While this is a great technology to return to a known state at a certain point in time it is not supported for running Exchange Server 2007 SP1 in a production environment.

Exchange Server 2003

According to the earlier referenced Microsoft article, the only Exchange version officially supported under Hyper-V is Exchange Server 2007 SP1. If you want to virtualize Exchange Server 2003 the only official supported way to achieve this is to run it on Virtual Server 2005 R2 SP1.

Although not officially supported, Exchange Server 2003 runs great in a Virtual Machine under Hyper-V.

Conclusion

Microsoft now supports running Exchange Server 2007 in server virtualization environments. This is not only Hyper-V, but all vendors that have their solution validated in the Server Virtualization Validation Program (SVVP) are fully supported.

Important to remember is that all design guidelines that are valid for a bare metal Exchange environment need to be used for a virtualized environment as well. Use disks that are capable of handling the expected IO load, it is very likely that dynamic disks will not meet these requirements. Therefore the use of dynamic disks is not officially supported by Microsoft when running Exchange Server 2007 SP1 under Hyper-V.

One should always take the business requirements into account, why do you want to virtualize and what are the project goals against which costs? I have seen customers returning from the virtualization twilight zone being very frustrated and ending up in a bare metal

environment. For larger environments I have seen a lot of large implementations where the Exchange Server 2007 Client Access Servers and Hub Transport Servers were running under VMware, but where the mailbox servers were running on bare metal. I have to admit though that these environment were larger than 4000 mailboxes. For these environments virtualizing the mailbox servers is questionable, but this might change in the future and this is something that is hard to foresee.

Virtualizing Exchange: points for discussion

20 November 2008

by Nathan Winters

With the increasing acceptance of the use of Virtualization as a means of providing server infrastructure, this technology is being applied to production Exchange servers. This is a solution that is not just limited to the small shop, Nathan Winters discusses the pros, cons and challenges that lay ahead in providing a flexible and highly available email system.

Introduction

The subject of virtualization has been a pretty hot topic for the last couple of years. It seems to offer a massive amount, for example: consolidating underused servers and therefore providing cost savings on hardware, rack space, electricity and cooling. On top of that, virtualizing your server infrastructure brings other benefits such as the ability to easily move a virtual machine from one piece of physical hardware to another and also to rapidly provision new servers where required.

Although this sounds great, there has been somewhat of a problem with virtualising Exchange 2007 because, Microsoft has not, until very recently, supported Exchange 2007 on any virtualization platform. To be fair though, that has not stopped a lot of people from putting their Exchange servers on a virtual platform. In fact it is something I have, myself, done for a handful of clients utilising the VMware ESX platform. However, as mentioned, the support situation has now changed, and these changes form the topic of the next section.

Is it supported?

In this rapidly changing market place, there is no straight-forward answer to the question of support. Essentially, as far as Microsoft is concerned, this all boils down to what they have tested. Exchange 2003 has long been supported on the Microsoft Virtual Server platform, although only if you had a Microsoft Premier Support Contract. I have seen almost nobody using Virtual Server as anything other than a test platform, and other products like VMware Workstation were rather better for that. So why is Exchange 2007 different? It is because Exchange 2007 is the first Microsoft Server application which required a 64 bit (x64) operating system (OS). At the release of Exchange 2007, Microsoft still only had Virtual Server as their virtualization offering, which does not support the 64 bit OS required to run Exchange 2007. Therefore, Exchange 2007 was not supported in a virtual environment as Microsoft would have been reluctant to test on another company's virtualization platform!

Since July 2008, when Microsoft released Hyper-V, Microsoft's new hypervisor based virtualization platform, it now has a virtualization platform capable of supporting 64 bit operating systems. Therefore it came as no surprise that Microsoft issued a new support statement in August this year which is found at the link below:

Microsoft Support Policies and Recommendations for Exchange Servers in Hardware Virtualization Environments

Some of the key points are:

- The virtualization software must be one or other flavour of Hyper-V, or if 3rd party, must be listed on the Windows Server Virtualization Validation Program (SVVP).
- Exchange must be running on Windows Server 2008.
- The Unified Messaging role is not supported, although all others are.
- Virtual disks that dynamically expand are not supported by Exchange.
- Virtual disks that use differencing or delta mechanisms (such as Hyper-V's differencing VHDs or snapshots) are not supported.
- Snapshotting Exchange server virtual machines is not supported, as this is not an Exchange aware backup mechanism.
- Both CCR and SCC are supported in hardware virtualization environments provided that the virtualization environment does not employ clustered virtualization servers.

The final issue listed here deserves some clarification as it has already caused some confusion. Essentially you can do one or the other of the following where the preference would be the second option: (1) cluster the hypervisor roots, or (2) cluster the guests (using multiple roots, for example CCR where one node is a VM on one root and one node is a VM on another root). You cannot, however, combine the two, which suggests that using technology like VMware's VMotion and HA/DRS is not supported in conjunction with CCR or SCC. Perhaps this again will change when Microsoft release they Live Migration technology in R2 of Hyper-V although having said that, Microsoft currently has Quick Migration which also isn't supported when using Exchange clusters. We will simply have to wait and see.

Although the use of hypervisor high-availability techniques, in conjunction with Exchange clusters, is not actually supported, using VMware VMotion to move the active node in a CCR cluster does actually work (on a test system at least), and doesn't seem to cause a failover either! You do get a bunch of Event ID 1122 and 1123 messages telling you about the lost network connectivity, but things appear to keep working. Of course this may well not be true for a heavily loaded system as depending on the amount of time required to VMotion, a failover may be triggered. All in all, it isn't supported and frankly, going to the trouble of setting up an high availability system, only to run it in an unsupported way, seems rather perverse to me!

Alongside the support announcement discussed above, Microsoft also changed some licensing conditions as the below quote from the Exchange team blog describes:

> *Microsoft is waiving its 90-day license reassignment policy to enable customers who virtualize Exchange to move their licenses between servers within a data farm as often as necessary.*

So all in all this means that it is very much supported to run Exchange under the conditions stated on the Microsoft Hyper-V platform and to have flexibility to move virtual machines between physical hosts. What I guess a number of you are wondering is what does this mean for VMware support? Well to answer that we have to take a look at the Server Virtualization Validation Program or SVVP.

The website for SVVP can be found at the link below:

Server Virtualization Validation Program

Essentially this is a way for Microsoft to certify whether 3rd-party virtualization products can adequately serve Microsoft Windows and the Microsoft software which runs on Windows. VMware is on this list as supported however, what is critical to note with regards Exchange is the supported processor architecture type. The link below lists the platforms which have met the requirements of the program for x64 processors:

Server Virtualization

Interestingly this list changed during the writing of this article (Oct 2008). When I started the article, VMware was not on the list! However, during the first revision process, VMware became a supported x64 platform with the only caveat being that virtual machines can only have 16 GB of RAM allocated to them.

I hope having read the above you have an understanding of the issues surrounding the support of Exchange on a virtualization platform. Before we move on, I would simply like to add that I and many other people have been running Exchange 2007 on various virtual platforms and it does generally work very well. The support issue is a risk, but often it is not a big enough risk to stop people making use of the benefits or virtualization technology. Having addressed the support issue let's now move on to take a look at virtualizing Exchange in practice.

Virtualizing Exchange in practice

In this section I will look at the pros and cons of virtualizing Exchange. I must make it clear at this point that I have not yet had the opportunity to deploy Exchange 2007 on Hyper-V, although I have deployed Exchange 2007 on VMware ESX. Therefore, my comments on Hyper-V are based on numerous discussions with trusted colleagues who are specialists in the Virtualization field.

Benefits of Exchange Virtualization

As you will know if you run an Exchange organisation, Exchange is not a simple application. Exchange is a massive product and getting familiar with it all is not easy, therefore, the ability to have a replica of your production environment is extremely helpful when it comes to the testing and validation of upgrade and migration work. Virtualization makes this extremely simple as copies of existing physical, and virtual machines, can be taken, and then run in isolation from the main network.

The introduction of multiple roles in Exchange 2007 has been very helpful in allowing Exchange to scale well. However, it has also pushed up the server count dramatically. Virtualizing the Exchange environment allows these multiple roles to still run on separate virtual machines which can be tuned accordingly whilst keeping down the number of physical machines required.

Following on from the flexibility of providing a test lab, it is common nowadays, that this lab solution be provided on equipment available for disaster recovery. The ability to move a Virtual Machine from one piece of hardware to another means that in a disaster, getting the Exchange services up rapidly is much simpler than when relying on the correct physical hardware being available.

Interestingly, and one area which I did not expect to put in the benefits section is performance. It would appear that when sized according to Microsoft Exchange sizing guidelines, a virtualized Exchange infrastructure can perform almost as well as a physical one. This was discussed at VMworld this year. In particular, it would seem that message throughput is actually slightly better on a virtual Hyper-V platform than on physical hardware.

Another area which could be considered a benefit and a possible problem, is management. The reason I mention it as a possible problem, is simply that it is another layer of management technology, however, once you accept the necessity, then management options are a definite benefit. It is one area in which Microsoft excel. Although VMware have their Virtual Center console which is not a bad solution, the Microsoft solution, System Center Virtual Machine Manager (VMM), to give it its full name, gives you cradle to grave, hardware to application management. VMM 2008, like Virtual Center, has to be purchased however, what is brilliant is the way it integrates with System Center Operations Manager 2007 using a connector/management pack. That gives you expertise on your infrastructure that's built into the network. Want to know which servers have spare resources to be potential hosts? Want to know which machines should be

converted to VM's (and then P2V them). Want performance/health information in a single integrated management infrastructure (System Center)? This complete management solution is something the competition struggles to match.

Problems with Exchange Virtualization

Of course there are some problems with virtualizing Exchange. It is important not to think that virtualization gives endless resources. It is absolutely critical to size the servers just like you would before. On top of this, running a virtual platform gives an added layer of complexity that must be understood and managed carefully so as to provide a good platform for the virtual machines it supports.

Having mentioned performance in the benefits section, I think it is worth entering it here too. Why? Because I still feel that putting another layer underneath something that is already IO and memory intensive isn't necessarily the greatest idea. After all, if you only run one VM on the box why not just use physical hardware? I feel that this is particularly true when running the Exchange Mailbox server role on a virtualization platform. Whilst it is true that when using pass through disks performance is often within 1% of the physical hardware it is however, still true that when load in particular on the network cards increases performance problems can occur. There will be improvements in this area soon, as new network cards increase throughput by implement virtual switches in hardware. Whatever, it is not recommended to virtualise more than one Mailbox server on a single virtual host machine.

Still there are benefits to virtualization so at this point it is perhaps a question of whether these outweigh the fact you may get less users on a VM than when using a physical machine. Looking at non mailbox roles, Unified Messaging is simply not supported and really doesn't scale well even if you try to put it in a virtual environment, as the audio playback can become rather choppy! A possible barrier to virtualizing the Client Access and Hub Transport roles is the implementation of Windows Network Load Balancing (WNLB) on the virtualization platform. This is something that I have struggled with however; it would seem that it is possible but that problems can occur unless things are configured correctly as described at the links below. To be fair this is no different when in a physical environment.

Microsoft NLB Not Working Properly in Unicast Mode.

Exchange 2007 Unicast NLB Issue on Hyper-V

Although virtualization brings some benefits for disaster recovery as mentioned above, one can't get away from the fact that by virtualizing you are putting a number of servers on a single physical piece of hardware. Obviously there is a need to mitigate this single point of failure and one method is to use redundant hardware including PSUs and NICs. This is one area to be particularly careful of on the Hyper-V platform, as under Hyper-V physical NIC teaming is not currently supported, so if you lose a NIC, everything on that machine loses connectivity!

Summary

Microsoft Hyper-V, VMware and other virtualization platforms provide a great platform for Exchange especially as you now have the comfort of knowing that the solution is supported. Realistically it is very likely that as virtualization becomes more and more accepted as the normal way of providing server infrastructure, production Exchange servers will be virtualized.

In my opinion, there are clearly a few challenges to be faced when virtualizing Exchange, but so long as the challenges/limitations of the infrastructure are understood, the guidelines laid out by manufacturers are followed and you thoroughly test the performance of the implementation before rolling it out in production, virtualizing Exchange can be very successful.

What is very clear is that virtualizing Exchange is no longer just an option for the small shop, but is now a solution for even the largest Exchange deployments, looking to provide a flexible and highly available platform for their email system.

Build Your Own Virtualized Test Lab

05 January 2009

by Desmond Lee

Desmond Lee explains the fundamentals of building a fully functional test lab for Windows Servers and enterprise applications (e.g. Exchange and Office Communications Server) from scratch with Hyper-V on a supported x64 machine. This will simulate a majority of test scenarios. He highlights the key points, and suggests tips and tricks to help to make the journey smoother

If you are involved in supporting test, training or development servers, then you will find that server- or host- based virtualization will solve many of your problems. Instead of having to stock up several physical machines, each dedicated to certain roles and functionality, you will be able to consolidate them onto a number of powerful boxes with speedy multi-core processors, massive RAM memory and fast disk drives with huge capacity. This will simplify administration, lessen the use of floor space, and reduce both costs and energy consumption.

In this article, I'll use Microsoft's host virtualization solution, Hyper-V, to build a representative test installation. Hyper-V is available as part of some editions of Windows Server 2008, or more recently as a free standalone product similar to VMware ESXi.

The Big Picture

Typically, a corporate infrastructure contains many servers providing core enterprise services such as network, directory, security, messaging and file and print. Of late, the concept of Unified Communications (UC) has gathered a lot of attention within the industry. Unified Communications is a term for the technologies that aim to integrate and coordinate the systems that provide telephony, email, instant messaging, video, voicemail, and group chat, and ensure that the communications reach the recipients in the most effective and timely way.

One beauty of using virtual machines (VM) is the ease with which you can add, remove and edit virtualized hardware components such as network adapters, RAM memory and hard drives at will. You can also roll-back an installation should something go wrong, as long as you have taken point-in-time snapshots. The test lab is designed to run satisfactorily on up-to-date machines equipped with 2GB or more RAM memory and plenty of disk storage on the host machine. You can then scale up the hardware to accommodate additional system demands and workload. Depending on the test scenarios being played out, the number of active running VMs can vary. Technically, there is no restriction imposed but the licensing model may limit the number of VMs that can be deployed, depending on the Windows Server 2008 edition chosen. As a starting point, all VMs are created with the standard configuration of default virtual devices, with the addition of the legacy network adapter setup to connect to one common internal, local-only (private) virtual network switch.

Although migrating virtual machine configuration (.vmc) and virtual hard disks (.vhd) created in Virtual PC or Virtual Server to Hyper-V is relatively painless, I shall instead cover the fundamentals of building the entire test lab from scratch with Hyper-V on a supported x64 machine. We'll have to assume that you know how to run setup and accomplish basic tasks on your own without much hand-holding. I'll try to highlight key points, and suggest tips and tricks to help to make the journey a lot smoother.

All guest server VMs will be built using the SP1 version of Windows Server 2003, as it is so widely deployed. Hyper-V will support both x86 and x64 guest VMs (or child partitions) running side-by-side, but we'll choose the x86 edition so as to keep the overall memory requirement to a manageable level with minimum overhead. All this is done without the risk of losing functionality in the sort of application that most organizations will be working with.

You can find a checklist summarizing the test lab machines in Table 1.

name	ram (mb)	ip/24	roles / services	configuration
LABDC01	384	10.0.0.11	Windows Server 2003 SP1 x86 DC, DNS, GC WWW only Certificate Services CA and Certificate Services Web Enrollment Support	FQDN = testlab.local Raise domain functional level to "Windows Server 2003" Enterprise root CA = TestLabRootCA DNS SRV record = _sipinternaltls._tcp. domain.com:5061
LABEX01	768	10.0.0.12	Windows Server 2003 SP1 x86 WWW only NET Framework v2.0 Redistributable Package x86 Windows PowerShell 1.0 RTW x86 KB926776 A hotfix rollup package for 913297, 913393, 918995, and 924895 December 2007 cumulative time zone update for Microsoft Windows operating systems (KB942763) NET Framework 2.0 SP1 runtime x86 Windows Server 2003 SP2 x86 Exchange Server 2007 SP1 x86	Exchange organization (native) = TestLabExchange compatible with Outlook 2003 (public folder)
LABOC01	512	10.0.0.13	Windows Server 2003 SP1 x86 WWW only VC++ 2005 Redistributable .NET Framework v2.0 Redistributable Package Execute step1 to 8 in OCS setup. Windows Server 2003 adminpak.msi	testlab.local forest / Global Properties / Meetings tab / Default Policy / Edit / Check "Enable web conferencing" and "Enable program and desktop sharing"
LABXP01	384	10.0.0.51	Windows XP SP2 x86 Outlook 2003 Office Communicator 2007 client Live Meeting 2007 client Office Outlook Conferencing Add-in	

name	ram (mb)	ip/24	roles / services	configuration
LABVT01	448	10.0.0.52	Windows Vista Ultimate SP1 x86 Outlook 2007 Office Communicator 2007 client Live Meeting 2007 client Office Outlook Conferencing Add-in	

The Journey Begins

Active Directory (AD) provides resource authentication and authorization for third party solutions as well as all the Microsoft product family. You will need a solid Domain Naming Services network service to ensure problem-free operations across the entire AD forest. Therefore, the very first server (VM) should be set up as a Domain Controller (DC) and named as LABDC01 in our single forest, single domain infrastructure. For many good reasons, we must install the Microsoft DNS service and integrate it with Active Directory.

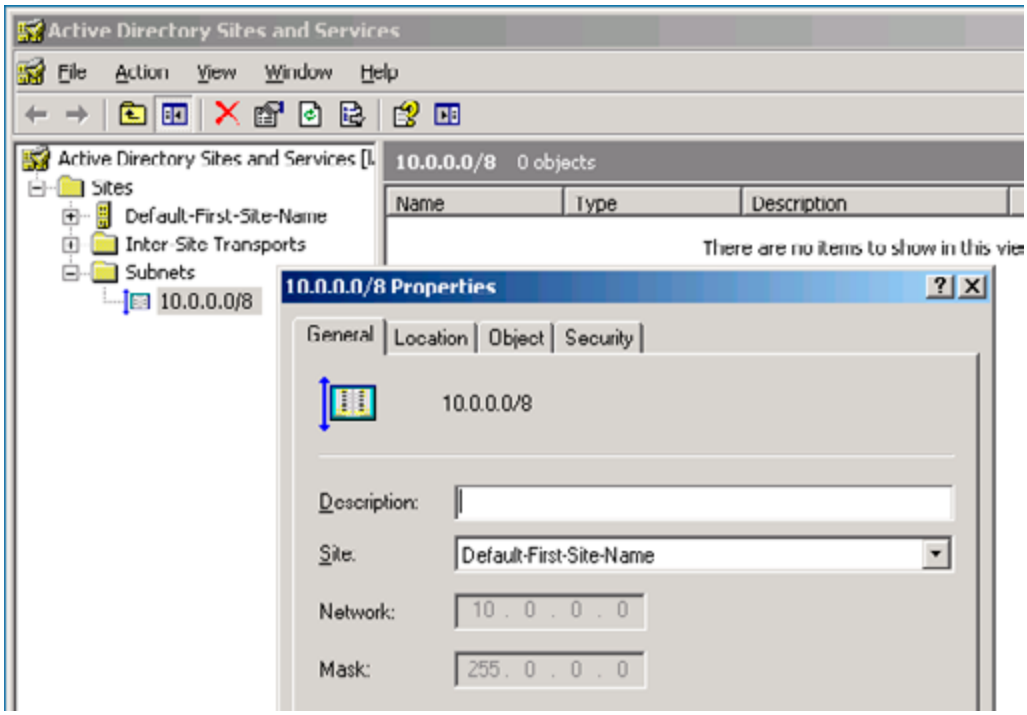

Figure 1: Configure and associate IP subnet with AD site.

We use the fully qualified domain name (FQDN) of testlab.local to scope the authoritative DNS and AD name space to the private intranet. Being the first DC in the domain implies that the Global Catalog role will also be automatically enabled. As a best practice, create an IP subnet object (10.0.0.0/8) and associate it with the Default-First-Site-Name AD object from which all applications that use AD directory services will gain optimal traffic routing (see Figure 1).

The default AD domain and forest functional levels are set at that of "Windows 2000 mixed" and "Windows 2000" respectively. In order to introduce Exchange 2007 as the messaging service, the target domain that holds one or more Mailbox server role must be at the "Windows 2000 native" mode domain functional level or higher. Similarly, any AD domain with mailbox-enabled accounts (user, contact, etc.) – even with no Exchange 2007 Mailbox server role present – must be raised to this minimal domain functional level. To support Office Communications Server 2007, you will need to raise the domain functional level all the way to "Windows Server 2003'.

Messaging Service

The process of setting up Exchange Server 2007 with SP1 is a bit more tedious and time consuming. As a first step, you will need to install the WWW service from the Windows Server 2003 SP1 media. Unlike its predecessor, the SMTP service is not a requirement to setup Exchange 2007.You will save time and avoid frustration if you next download items 3 to 9 listed in Table 1 from your administrative workstation.

When you are ready, install the files in the order listed onto the VM named `LABEX01` that is joined to the domain. You will require SP2 for Windows Server 2003 in order to run Exchange Server 2007 SP1 on this version of Windows. Because the Microsoft Management Console 3.0 is shipped as an integral part of Windows Server 2003 SP1 and above, you won't need to install it separately. Likewise, all the pre-requisites would have been met and you can straightaway launch step 4 as soon as you run the Exchange 2007 with integrated SP1 executable.

The typical Exchange Server installation option offers the minimal server roles that are compatible for setup and running Exchange on the same machine, namely Client Access, Hub Transport and Mailbox server roles. Throughout the installation process, accept all the proposed installation options except one and name the Exchange organization `TestLabExchange`.

Figure 2: Configure download client compatibility.

To maintain compatibility with Outlook 2003 and earlier or Entourage clients, confirm that the Yes radio button is selected at the "Client Settings" page (see Figure 2); then you won't have to worry about any extra configuration steps later on. Updates to the Active Directory schema and configuration is part of the Exchange installation process so there is another one less thing to be concerned with.

When you've completed setup and reboot the machine as prompted, your messaging platform is almost ready for service. Start Exchange Management Console on LABEX01 and create 2 new mailbox-enabled user accounts, say Alice and Bob. We have not yet mentioned the application clients which we shall do so next.

The Client Side

To model real world usage as far as possible, I suggest that you build two different client machines that run Windows XP SP2 and Vista SP1. Assign the names `LABXP01` and `LABVT01` and join them to the domain. On the respective systems, install the Outlook 2003 and 2007 clients. You may choose to set up other additional applications in the Office 2003/2007 suite. This will allow the Office 2007 applications to provide "click-to-dial" and to supply information about the users" presence to the Unified Communications system.

Next, log in to the XP desktop as Alice and Bob on Vista. Outlook 2007 on Bob's desktop will attempt to automatically locate, connect and configure appropriate settings with the correct Exchange home server (`LABEX01`). In contrast, you will have to manually specify Alice's SAM account or display name, and the name of the home Exchange server, to achieve the same results. With this out of the way, you can then proceed to run functional tests for mail, calendaring and scheduling: Then you can convince yourself that the core messaging services work as expected.

Extend Your Reach

Up to this point, the setup procedures for the basic test lab have been fairly straightforward. In order to provide Unified Communications functionality, you will have to face a lengthy setup process with Office Communications Server 2007.

Before we even begin, an internal Public Key Infrastructure (PKI) must already be present unless you are going to use digital certificates from public root Certificate Authority (CA) such as Verisign and Thawte. This strict requirement guarantees that all OCS server-to-server and server-to-client communication channels are properly secured. You will usually save cost with this approach if you have a large number of servers or clients that will not be directly accessible by or communicate with the general public from the Internet.

For the test lab, it is sufficient to install the Certificate Services configured as an Enterprise root CA. Use the name TestLabRootCA with the default settings on the DC itself (see Figure 3 and 4). End-user or machine certificate application and enrollment can be greatly simplified through the use of the web front-end. For this reason, the WWW feature must be installed on the DC. In real practice, a separate member server will fulfill this special role.

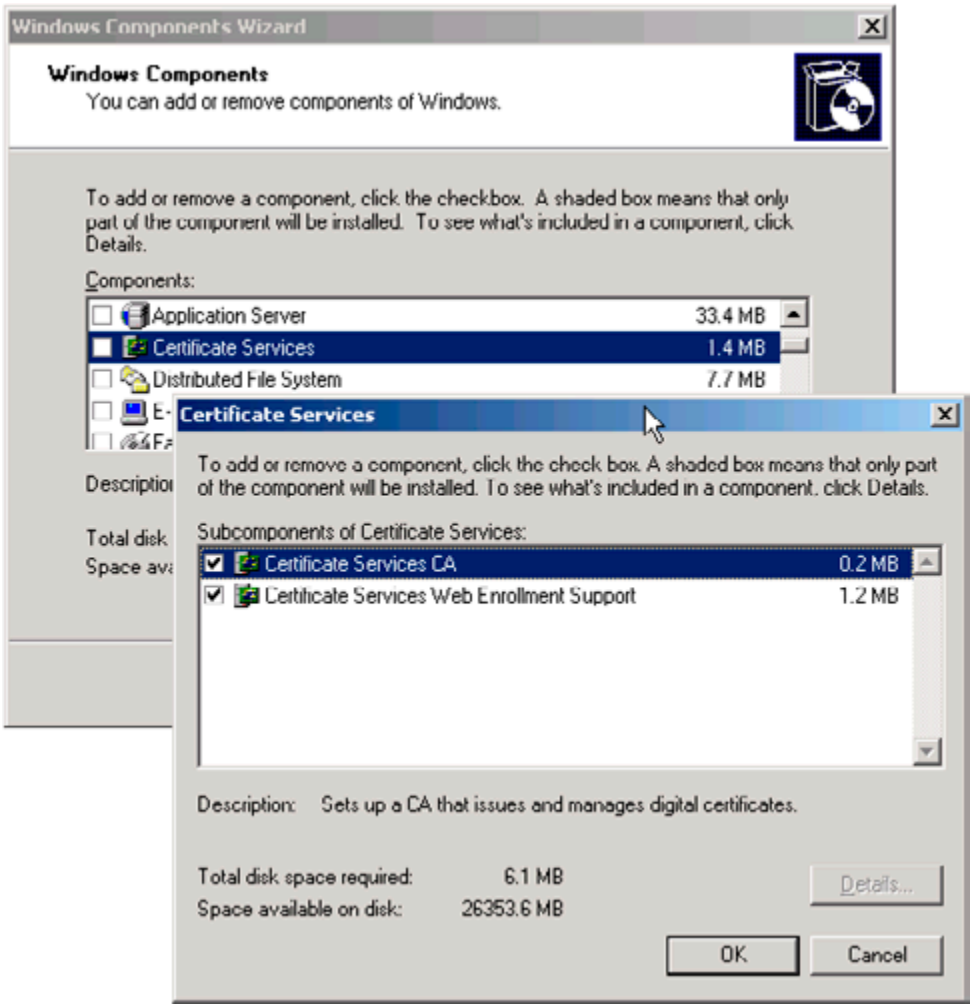

Figure 3: Setup Certificate Services.

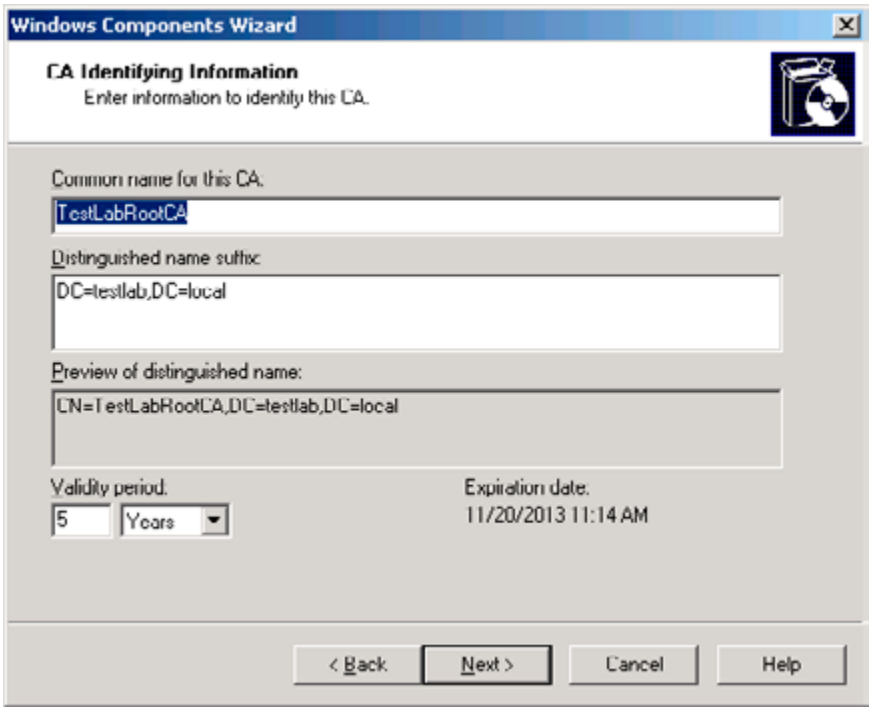

Figure 4: Configuring CA identifying information.

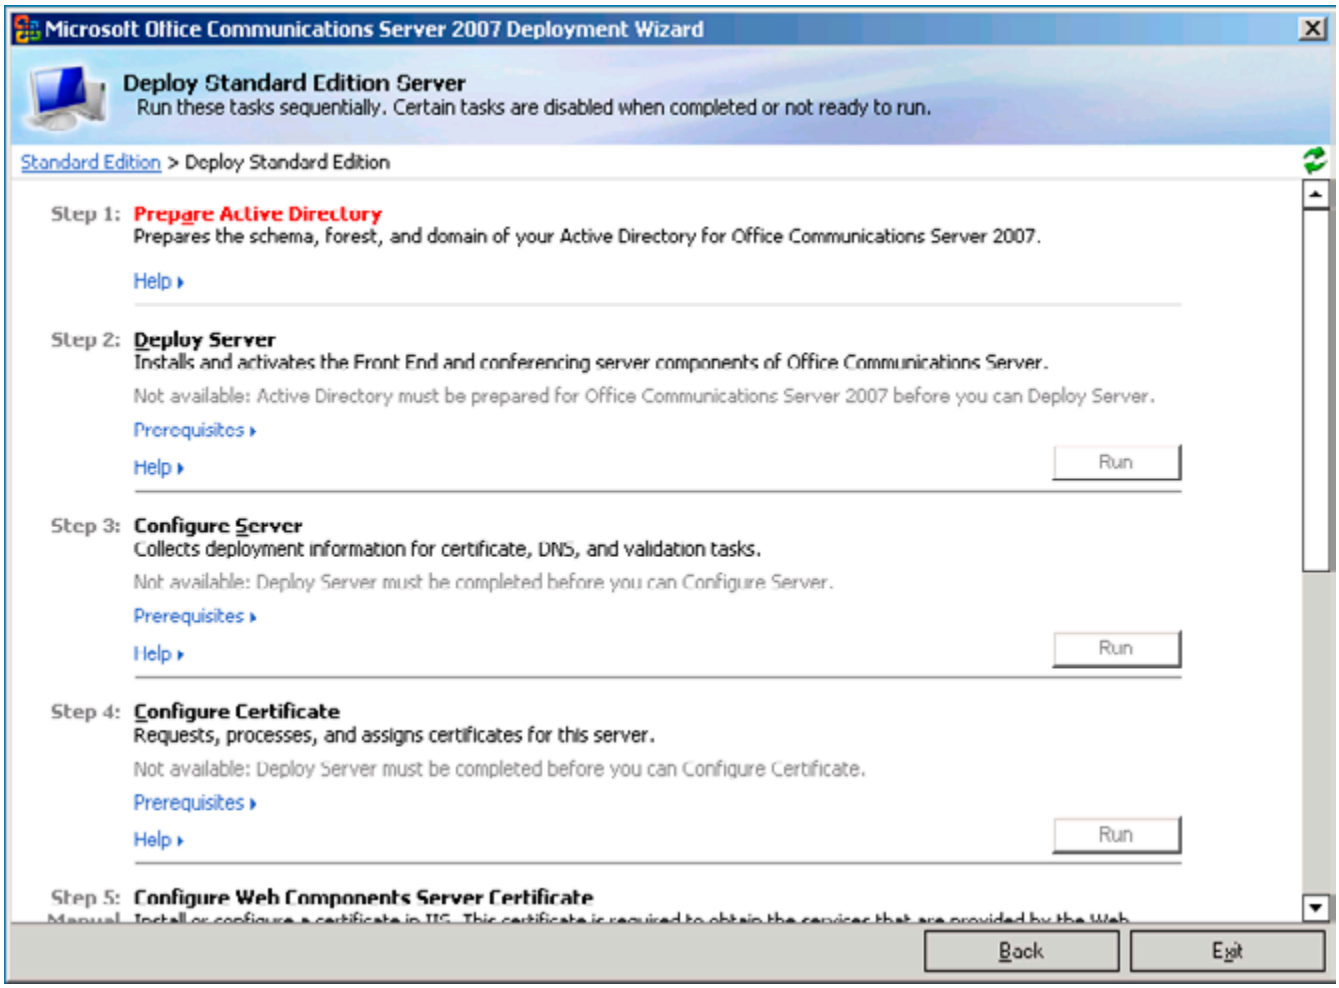

Figure 5: Deploy OCS Standard Edition (steps 1 to 8).

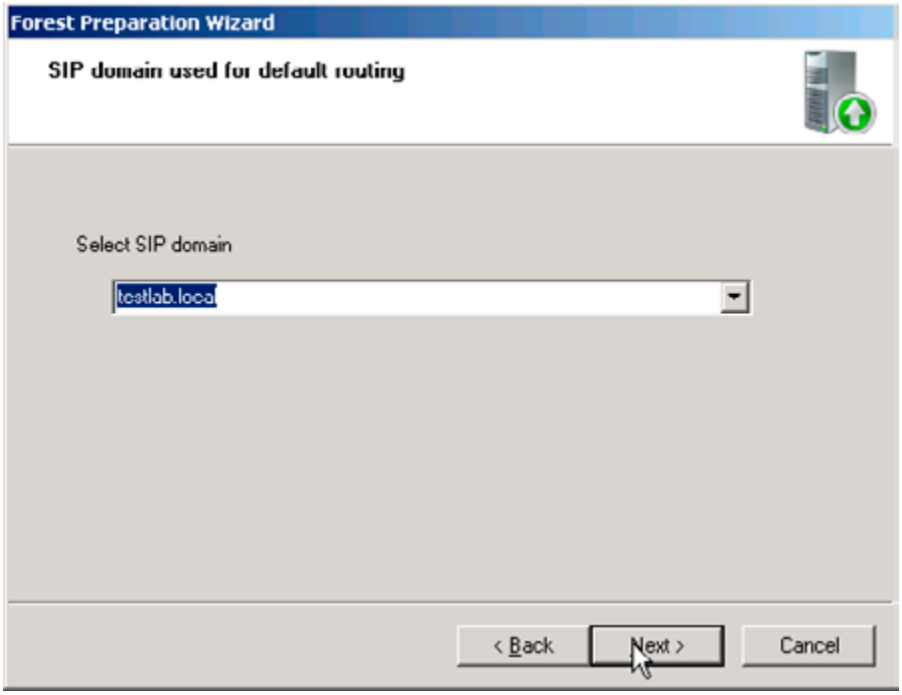

Figure 6: Configure SIP domain.

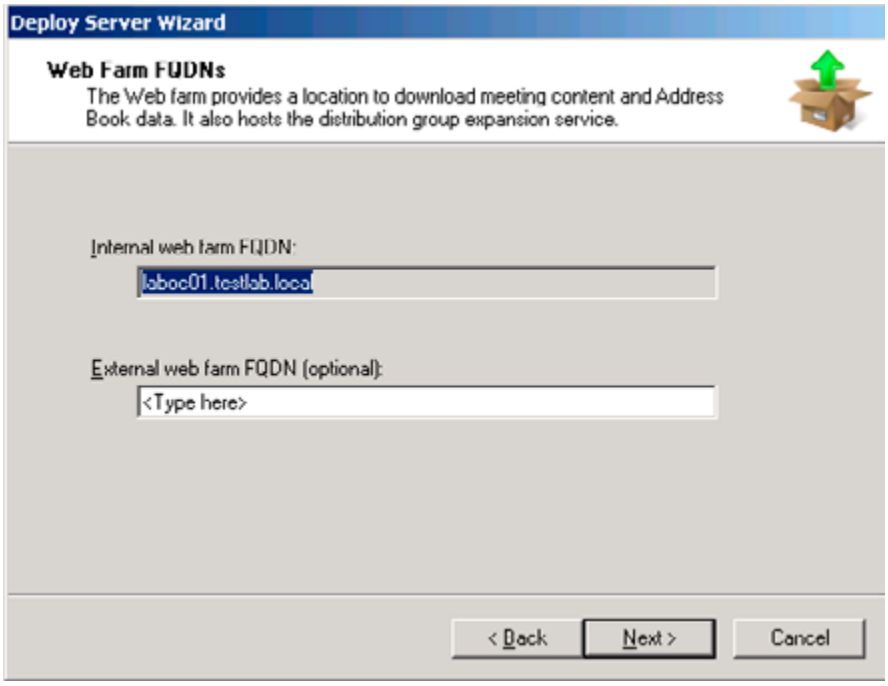

Figure 7: Specify Web Farm FQDN.

For a simple OCS standard edition deployment scenario for internal user access, the core set of features such as Instant Messaging, Audio/Video and Web Conferencing all reside on the same machine. Again, the WWW service is featured prominently and must be installed ahead of time. Unlike Exchange Server 2007, OCS ships with all the pre-requisite applications. Once you start setup, you will have the opportunity to install the Visual C++ 2005 Redistribution Kit as well as the .NET Framework runtime. Similar to the other installations, you should login to `LABOC01` with the built-in Administrator account in the domain. By default, this account has membership in the Domain Admins and Schema Admins groups. This permits you to execute steps 1 through 8 in sequence without having to switch to and fro to the DC that holds the critical Schema Masters and other FSMO roles (`LABDC01`); see Figure 5. I recommend that you use `testlab`.

`local` as the SIP domain and `laboc01.testlab.local` as the internal web farm FQDN to keep things simple (see Figure 6 and 7). Otherwise, it is enough to accept all the default settings during the OCS setup. Note that AD Domain and Forest preparation steps are an integral part of the OCS installation process.

At the termination of validation step number 8, you can expect errors to surface because certain roles and features are not installed for a standard OCS setup. This can be safely ignored . Following that, pop in the Windows Server 2003 media and execute \i386\adminpak.msi. This is essential to enable user management and administration of OCS profile settings on `LABOC01`.

> *At the time of writing, Microsoft announced that OCS 2007 R2 will be released to manufacturing shortly. Like Exchange 2007, OCS R2 is supported only on the x64 platform in a production environment. What remains unclear is whether an x86 version will be made available for testing, evaluation and administration , as was the case with Exchange. Nevertheless, it should be easy to introduce OCS R2 into the OCS 2007 environment we are building here.*

The Final Pieces

So that the Microsoft Office Communicator (MOC) client can automatically locate available OCS server pools without manual configuration, it is essential to setup the necessary SRV resource records in DNS. At a minimum, configure the _sipinternaltls._tcp.testlab.local to listen on TCP port 5061 on LABDC01 (see Figure 8). Once this is in place, go ahead and install items 3 to 5 in Table 1 on both the LABXP01 and LABVT01 clients. The Live Meeting and Outlook Conferencing Add-in clients are extra components responsible for the web conferencing feature. In OCS R2, basic service or help desk in the form of a full remote desktop client control feature is built right into the updated MOC client. The standalone Live Meeting client is still required for the full, rich conferencing experience.

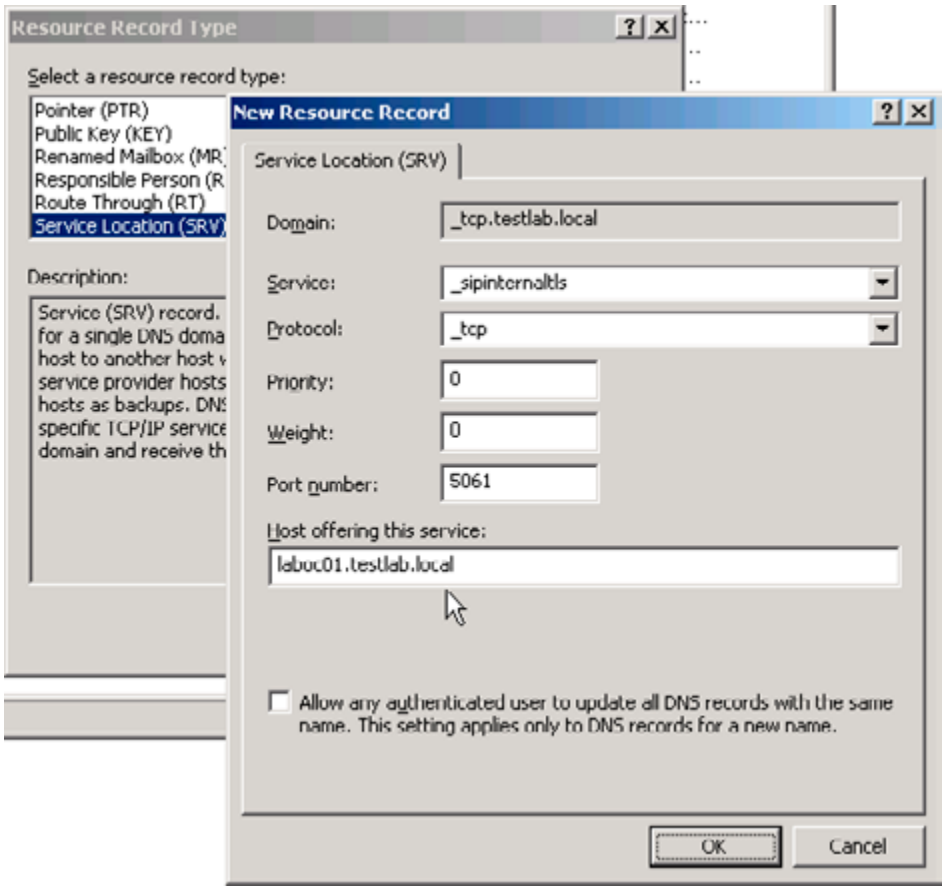

Figure 8: Configure SRV resource record.

You may want to restart the Windows client after installation completes. Although this is generally not required, you will then have the chance to identify any issues that typically surface only after a clean system shutdown and restart cycle is followed through. With a successful login to the domain, MOC queries DNS and your credentials will transparently be used to connect and sign-in to OCS.

To conduct the functional tests, observe that rich presence information (Available, Busy, etc.) is shown in MOC, Outlook as well as locally installed Office applications (See Figure 9 and 10). You can start an Instant Messaging conversation between Alice and Bob and escalate this to a collaborative web conferencing session. Using Outlook, it is possible to have appointments set up as ad-hoc or scheduled web conferencing sessions. For this to work, you must make the configuration changes on LABoC01 as described in Table 1.

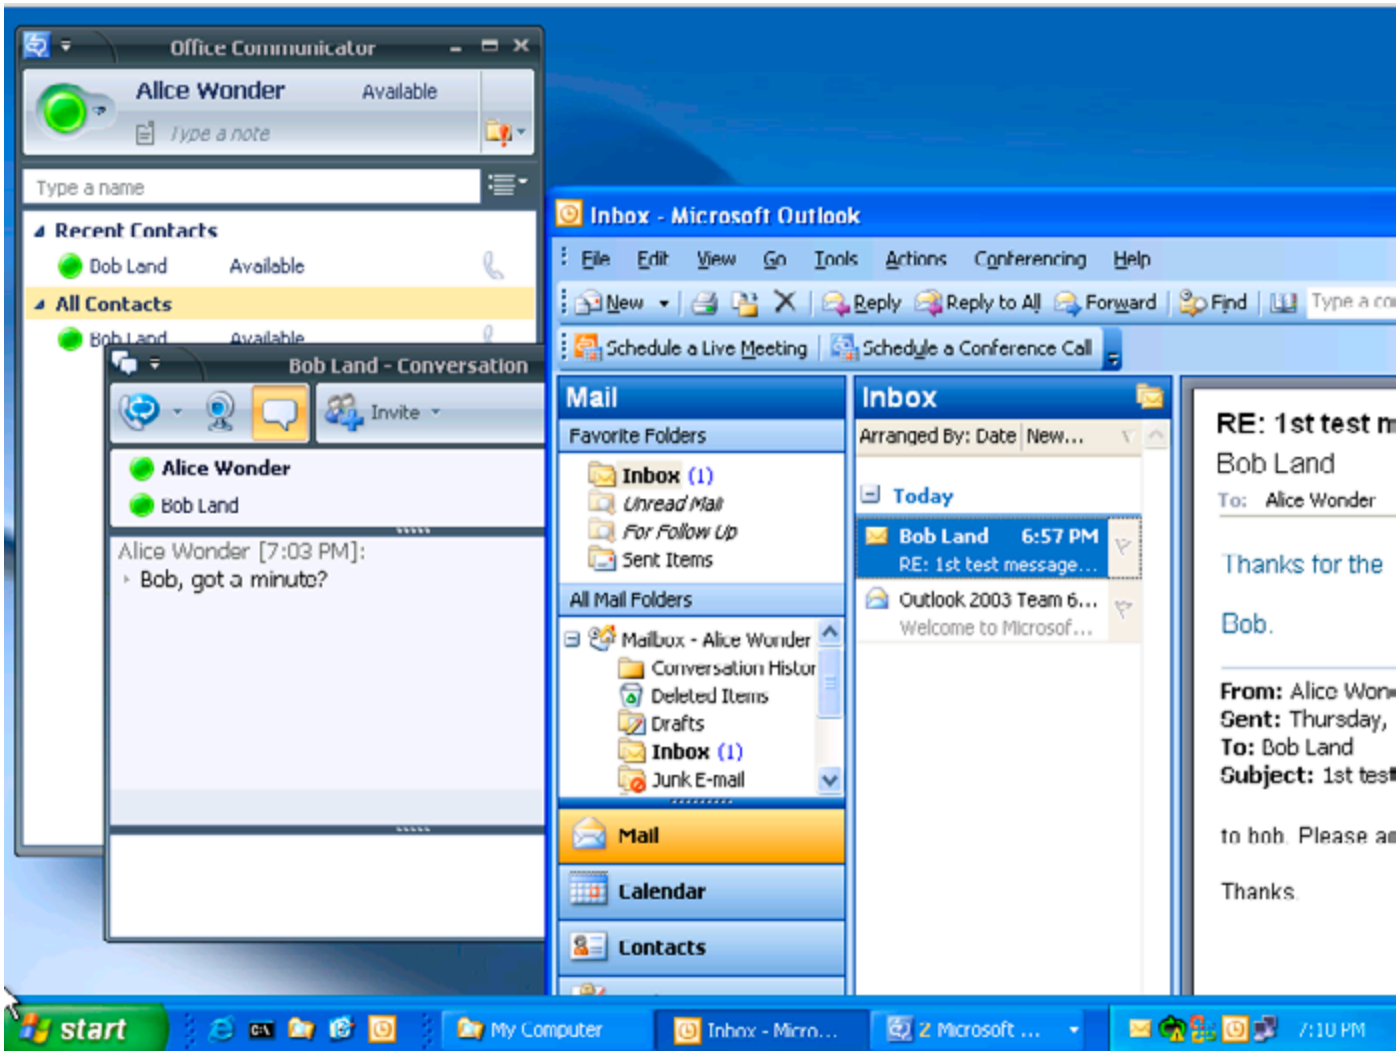

Figure 9: MOC and Outlook 2003 on XP client.

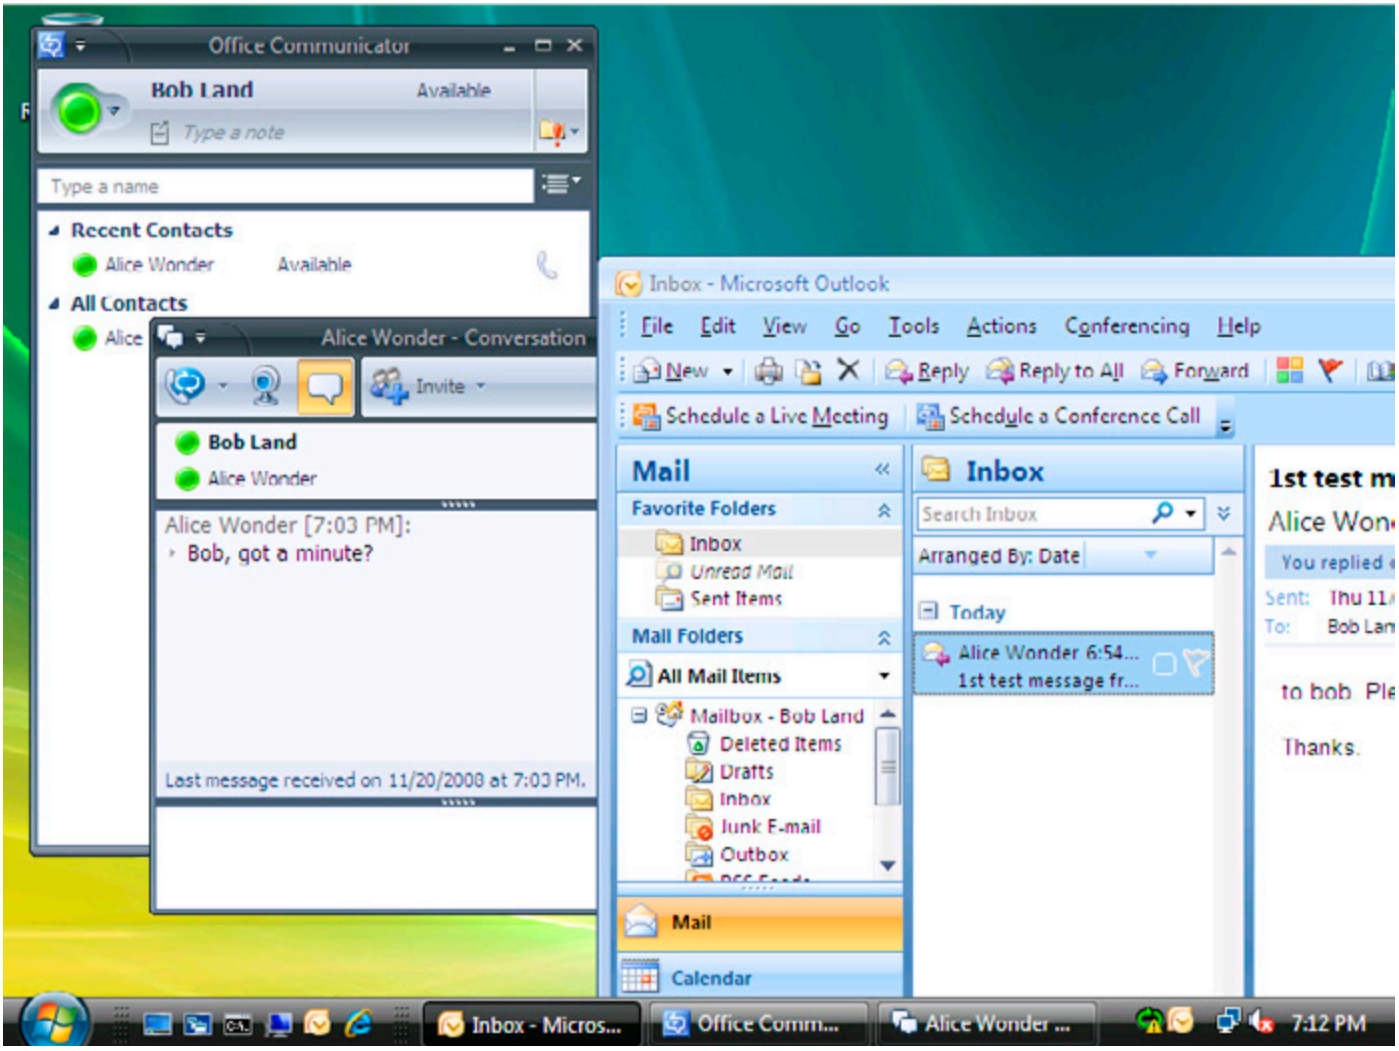

Figure 10: MOC and Outlook 2007 on Vista client.

Wrapping Up

Generally, you should install the Hyper-V Integration Components into each guest VM. This is carried out to boost performance and enhance host-guest integration. Other than the machines with Vista SP1 and Windows Server 2003 SP2 in our test lab, all the other Windows operating systems do not fulfill the pre-requisites to support the integration components. Therefore, the integration components are not installed in our test lab at the expense of VM performance improvements. This is a deliberate decision and does not affect the lab infrastructure.

Not surprisingly, there are already a number of patches and hotfixes since the official RTM of the various products. After you determine that functional tests of the enterprise services (directory, messaging, unified communications, etc.) perform as expected, consider applying the updates to bring the applications current. Before you do that, it is a good idea to save a working set of all the VMs by using Hyper-V's export feature. Subsequently, you can take a snapshot for each VM to enable quick rollback to a known state.

Using Hyper-V as the host virtualization solution is not without its shortcomings though. Because there is no native support for built-in sound card and USB devices, you cannot extend the test lab to include scenarios that involve enterprise voice communications (VoIP), Exchange 2007 Unified Messaging, fax integration or any hardware dependent services. Moreover, Novell SUSE Linux Enterprise 10 is the only non-Windows operating system supported in Hyper-V. You will have to look elsewhere if you are considering building integration solutions based on other Linux distributions or operating systems. One good example is the recently released Cisco Unified Communications Server 7.0 family of products.

To overcome such constraints, you can set up another physical machine to run VMware Server 2.0 or VMware Workstation 6.5 hosted on Windows 2003/2008 x64 editions. Unfortunately, these products will not even install when Hyper-V is detected on the host machine, therefore this extra hardware investment is essential.

Parting Words

By now, you should have a fully functional test lab to simulate a majority of real world scenarios. All traffic is confined to the internal network (local-only, private virtual network) and it does not take much to extend the test lab to cover external network segments. Subsequently, more VMs can be added to provide important services such as security, system configuration, monitoring and patch management. We'll look at these in a subsequent article.

A Beginner's Guide to Virtualizing Exchange Server – Part 1

05 May 2009

by Brien Posey

The advantage of virtualizing your servers is that it helps you make better use of your hardware resources, and reduces some of your licensing costs, However, there are disadvantages: With Exchange server, it isn't always obvious as to which server roles are suitable for virtualization, and it is rather hard to work out what system resources a virtual server is actually using. Brien Posey explains.

Virtualization is one of the hottest trends in IT today. Virtualization allows administrators to reduce costs by making better use of underused server hardware. Often times, virtualization has the added benefit of reducing server licensing costs as well. For all of its benefits though, there are some down sides to virtualizing your servers. Capacity planning and scalability become much more important in virtualized environments, and yet measuring resource consumption becomes far more difficult. In this article series, I want to talk about server virtualization as it relates to Exchange Server.

Microsoft's Support Policy

Before I even get started, I want to address the issue of whether or not Microsoft supports virtualizing Exchange Server. There seems to be a lot of confusion around the issue, but the official word is that Microsoft does support Exchange Server virtualization, but with some heavy stipulations.

I don't want to get into all of the stipulations, but I will tell you that Exchange Server 2003 is only supported if you use Microsoft's Virtual Server as the virtualization platform. Microsoft does not support running Exchange Server 2003 in a Hyper-V environment.

I could never, with a clear conscience, tell you to deploy an unsupported configuration. I will admit, however, that I virtualized my production Exchange 2003 servers using Hyper-V before Microsoft announced that the configuration would not be supported. To this day, I am still using this configuration, and it has been performing flawlessly.

Exchange Server 2008 can be virtualized using any hypervisor-based virtualization platform that has been validated under Microsoft's Windows Server Virtualization Validation program (http://go.microsoft.com/fwlink/?LinkId=125375) including Hyper-V and VMWare ESX. There are a number of stipulations that you need to be aware of though. These stipulations are all outlined in Microsoft's official support policy, which you can read at: http://technet.microsoft.com/en-us/library/cc794548.aspx. .

Server Roles

One of the biggest considerations that you must take into account when you are planning to virtualize Exchange Server 2007 is which server roles you want to virtualize. Microsoft's official support policy states that they support the virtualization of all of Exchange Server's roles, except for the Unified Messaging role. As such, you would think that figuring out which roles are appropriate for virtualization would be relatively simple. Even so, this is a fiercely debated topic, and everyone seems to have their own opinion about the right way to do things.

Since I can't possibly tell you which roles you should virtualize without being ridiculed by the technical editors and receiving a flood of e-mail from readers, I am simply going to explain the advantages and the disadvantages of virtualizing each role, and let you make up your own mind as to whether or not virtualizing the various roles is appropriate for your organization.

The Mailbox Server Role

The mailbox server role is probably the role that receives the most attention in the virtualization debate. Opponents of virtualizing mailbox servers argue that mailbox servers make poor virtualization candidates because they are CPU and I/O intensive, and because the virtualization infrastructure adds to the server's CPU overhead.

While it is true that mailbox servers are I/O intensive, that may not be a deal breaker when it comes to virtualization. Many larger organizations get around the I/O issue by storing the virtual hard drives used by the virtualized mailbox server on a SAN. Smaller organizations may be able to get around the I/O issue by using SCSI pass through storage to host the mailbox database and the transaction logs.

I have seen several different benchmark tests that show that while the abstraction layers used by the virtualization platform do place an increased load on the CPU, CPU utilization only goes up by about 5% (assuming that Hyper-V or VMWare ESX is being used).

The whole point of using virtualization is to make better use of underutilized hardware resources. Therefore, if your mailbox server is already running near capacity, then virtualizing it probably isn't such a good idea. Even in those types of situations though, I have seen organizations implement CCR, and use a virtual machine to host the passive node, while the active node continues to run on a dedicated server.

The Hub Transport Role

The Hub Transport Server role is one of the most commonly virtualized Exchange 2007 roles. Even so, it is important to remember the critical nature of this server role. All messages pass through the hub transport server, and if the server crashes then mail flow stops.

When you are determining whether or not you want to virtualize a hub transport server, it is important to make sure that you have some sort of fault tolerance in place You should also use performance monitoring and capacity planning to ensure that the transport pipeline is not going to become a bottleneck once you virtualize the hub transport server.

The Client Access Server Role

Whether or not the CAS server should be virtualized depends on a number of factors. For instance, many organizations choose to use CAS as an Exchange front end that allows users to access their mailboxes through OWA. If this is how you are using your CAS server, then you need to consider the number of requests that the CAS server is servicing. Some organizations receive so much OWA traffic that they need multiple front end servers just to deal with it all. If your organization receives that much traffic, then you are probably better off not virtualizing your CAS servers.

On the other hand, if you only use CAS because it is a required role, or if your users don't generate an excessive number of OWA requests, then your CAS server might be an ideal candidate for virtualization. The only way to know for sure is to use the Performance Monitor to find out how many of your server's resources are being consumed. Performance monitoring is important because other functions such as legacy protocol proxying (Pop3 / IMAP), Outlook Anywhere (RPC over HTTP), and mobile device support can also place a heavy workload on a CAS server.

The Edge Transport Server Role

The edge transport server is one of the more controversial roles when it comes to virtualization. Many administrators are reluctant to virtualize their edge transport servers because they fear an escape attack. An escape attack is an attack in which a hacker manages to somehow break out of the confines of a virtual machine and take control of the entire server. To the best of my knowledge though, nobody has ever successfully performed an escape attack.

If you are concerned that someone might one day figure out how to perform an escape attack, but you want to virtualize your edge transport server, then my advice would be to carefully consider what other virtual servers you want to include on the server. The edge transport server is designed to sit in the DMZ, so if you are concerned about escape attacks, then why not reserve the physical server for only hosting virtual machines that are intended for use in the DMZ. That way, in the unlikely event that an escape attack ever does occur, you don't have to worry about the physical server containing any data.

The Unified Messaging Server Role

As I stated earlier, the Unified Messaging role is the only Exchange 2007 role that Microsoft does not support in a virtualized environment. Even so, I have known of a couple of very small organizations that have virtualized their Unified Messaging servers, and it seems to work for them. Personally, I have to side with Microsoft on this one and say that just because you can virtualize a unified messaging server, doesn't mean that you should. Unified Messaging tends to be very CPU intensive, and I think that is probably the reason why Microsoft doesn't want you to virtualize it.

Resource Consumption

Capacity planning is an important part of any Exchange Server deployment, but it becomes even more important when you bring virtualization into the picture, because virtualizing implies that your Exchange Server is only going to be able to use a fraction of the server's overall resources, and that the virtualized Exchange Server is going to have to compete with other virtual machines for a finite set of physical server resources. The problem with this is that it can be very tricky to figure out just how much of a server's resources your virtual Exchange Server is actually using.

To show you what I am talking about, I want to show you three screen captures. First, take a look at the screen capture that is shown in Figure A. This screen capture shows a lab machine that is running three virtual servers. One of these virtual servers is a domain controller, another is running OCS 2007, and the third is running an Exchange 2007 mailbox server (this is just a sample configuration, and is by no means a recommendation).

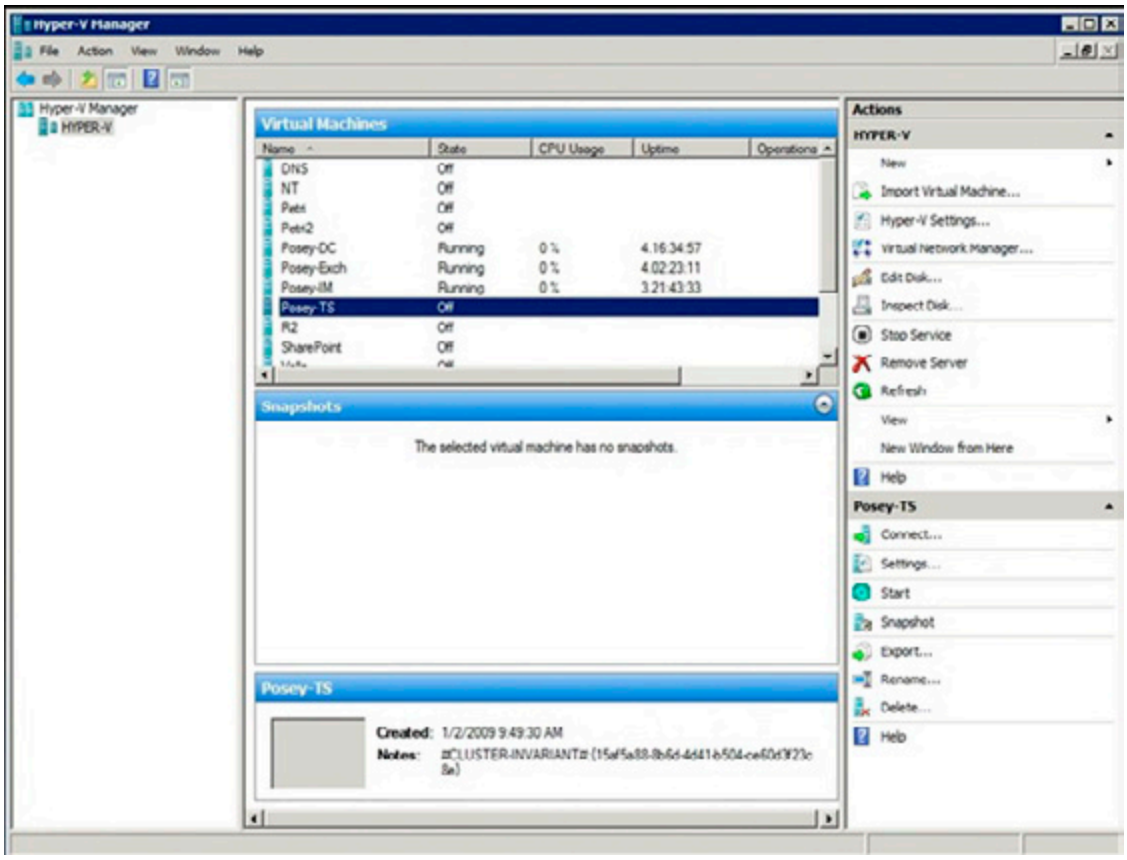

Figure A.

This lab machine is currently running three virtual servers.

If you look at the figure, you will notice that all three machines are running, and have been, for a few days straight. You will also notice that the Hyper-V Manager reports each server as using 0% of the server's CPU resources. Granted, all three of these machines are basically idle, but there is no way that an Exchange 2007 mailbox server is not consuming at least some CPU resources.

The second screen capture that I want to show you is in Figure B. This screen capture was taken from the same server, but this time I put a load on my Exchange 2007 mailbox server. While doing so, I opened up two instances of the Performance Monitor. One instance is running within the host operating system, and the other is running within the virtual Exchange server. Even though these two screen

captures were taken at exactly the same time, they paint a completely different picture of how the server's resources are being used. I will explain why this is the case later on, but you will notice that the host operating system reports much lower resource usage than the guest operating system does, even though there are a couple of other guest operating systems that are also running on the host.

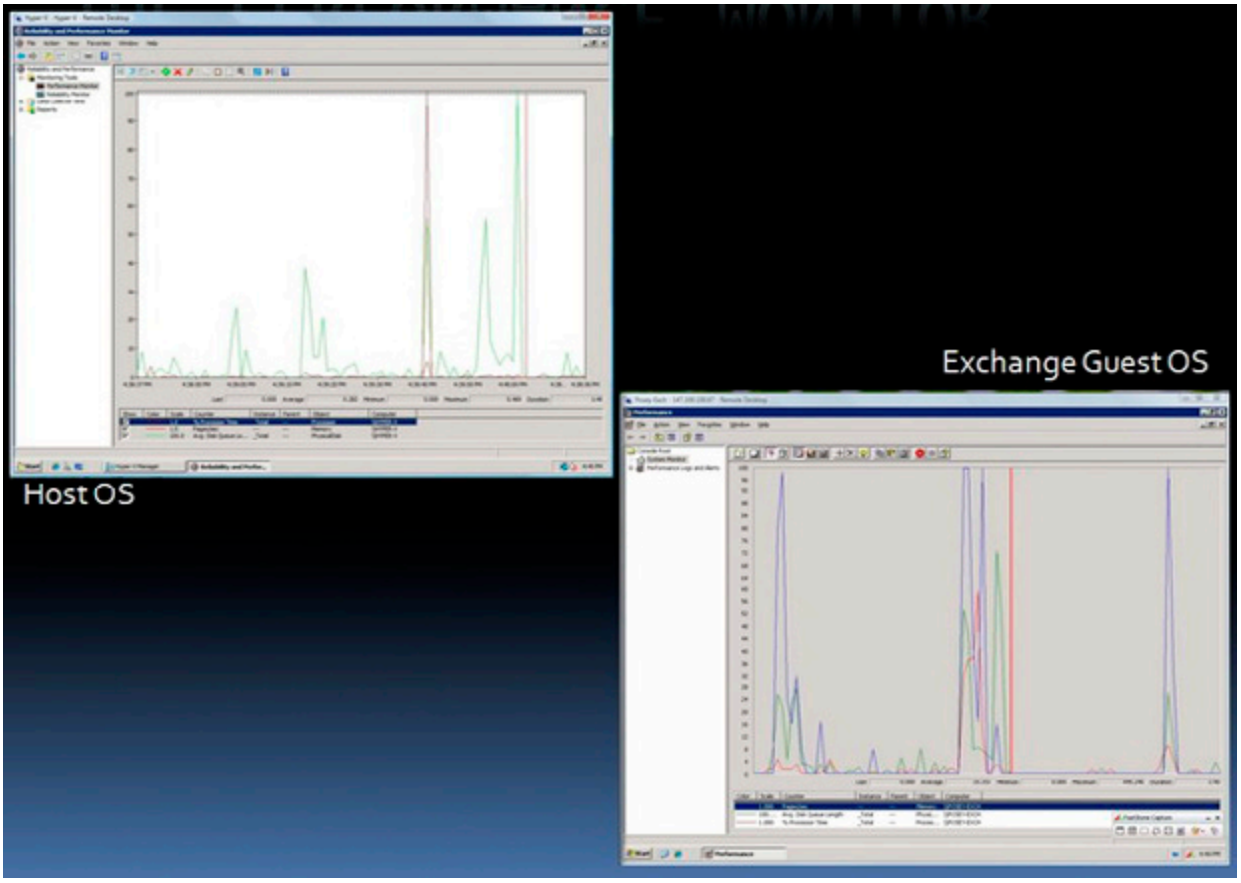

Figure B.

The host operating system and the guest operating systems have completely different ideas about how resources are being used.

The last screen capture that I wanted to show you is the one shown in Figure C. What I wanted to show you in this screen capture is that even the Hyper-V Manager and the Windows Task Manager, which are both running within the host operating system can't agree on how hard the CPU is working. The Hyper-V Manager indicates that five percent of the CPU resources are being used, while the Windows Task Manager implies that three percent of the CPU resources are in use (97% of the CPU resources are free).

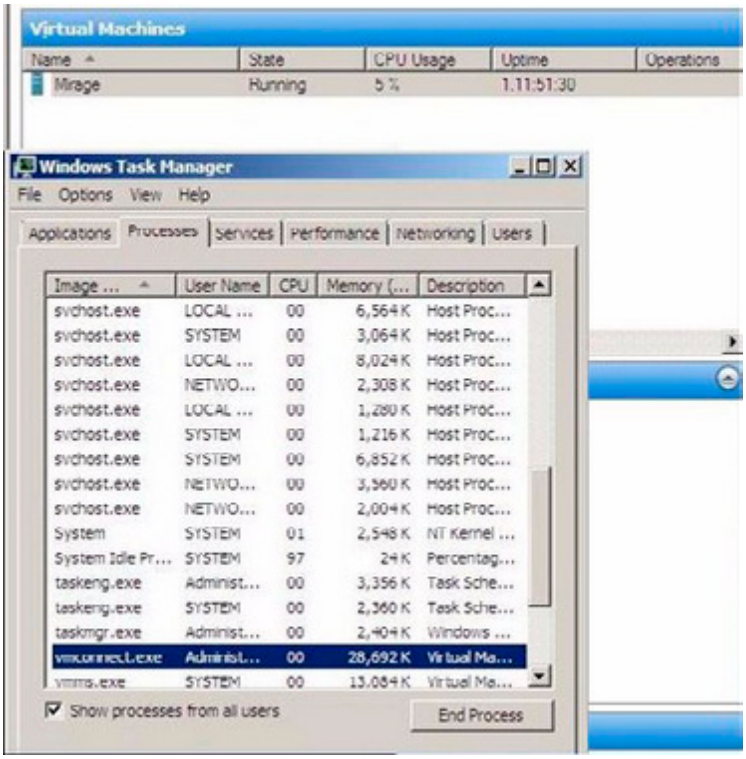

The Anatomy of a Virtual Machine

So why do these discrepancies exist? In order to answer that question, you need to understand a little bit about the anatomy of a virtual machine. Before I get started though, I need to point out that virtual machines are implemented differently depending on which virtualization product is being used. For the sake of this discussion, I am going to be talking about Microsoft's Hyper-V.

As I'm sure you probably know, Hyper-V (like many other virtualization products) is classified as a hypervisor. Hyper-V is what is known as a type 2 hypervisor. A type 2 hypervisor sits on top of a host operating system, as opposed to a type 1 hypervisor, which sits beneath the server's operating system at the bare metal layer.

A type 2 hypervisor isn't to be confused with hosted solution based virtualization products such as Microsoft's Virtual PC or Virtual Server. Such products typically pass all of their hardware requests through the host operating system, which sometimes results in poor overall performance.

In contrast, guest operating systems in a Hyper-V environment do reside on top of the host operating system, but Hyper-V is only minimally dependant on the host operating system (which must be a 64-bit version of Windows Server 2008). The host operating system connects to each virtual machine through a worker process. This process is used for keeping track of the virtual machine's heartbeat, taking snapshots of the virtual machine, emulating hardware, and similar tasks.

Now that I have explained some of the differences between a type 1 and a type 2 hypervisor, I want to talk about how the virtual machines function within Hyper-V. Hyper-V is designed to keep all of the virtual machines isolated from each other. It accomplishes this isolation through the use of partitions. A partition is a logical unit of isolation that is supported by the hypervisor.

Hyper-V uses two different types of partitions. The parent partition (which is sometimes called the root partition) is the lower level of the hypervisor. The parent partition runs the virtualization stack, and it has direct access to the server's hardware.

The other type of partition that is used by Hyper-V is a child partition. Each of the guest operating systems resides in a dedicated child partition. Child partitions do not have direct hardware access, but there is always at least one virtual processor and a dedicated memory area that is set aside for each child partition.

Initially, Hyper-V treats each of the server's processor cores as a virtual processor. Therefore, a server with two quad core processors would have eight virtual processors. It is important to keep in mind though, that Hyper-V does not force a one to one mapping of virtual processors to CPU cores. You can allocate more virtual processors than you have CPU cores, although Microsoft recommends that you do not exceed a two to one ratio.

Hyper-V manages memory differently from the way that it would be managed in a non virtualized environment. Hyper-V uses an Input Output Memory Management Unit (IOMMU) as a mechanism for mapping and managing the memory addresses for each child partition. In a non virtualized environment, low level memory mapping is handled primarily at the hardware level (although the Windows operating system does perform some higher level memory mapping of its own).

Earlier I mentioned that hypervisor based virtualization products tend to be more efficient than hosted solution based virtualization products, which pass all hardware calls through the host operating system. A big part of this efficiency is related to the root partition's ability to communicate directly with the server hardware. As you will recall though, the guest operating systems reside in child partitions, which do not have the ability to talk directly to the server hardware. So what keeps the guest operating systems from suffering from poor performance?

There are several different mechanisms in place that help to improve the child partition's efficiency, but one of the primary things that helps with guest operating performance is something called enlightenment. If you have ever installed a guest operating system in Hyper-V, then you know that one of the first things that you normally do after the operating system has been installed is to install the integration services. Technically, the integration services are not a requirement, and it's a good thing that they aren't. Many non Windows operating systems, as well as most of the older versions of Windows don't support the integration services.

The first thing that most administrators notice after installing a guest operating system is that they can't access the network until the integration services have been installed. This is where enlightenment comes into play. After the integration services have been installed, the guest operating system is said to be enlightened. What this really means is that the guest operating system becomes aware that it is running in a virtualized environment, and as such is able to access something called the VM Bus. This makes it possible for the guest operating system to access hardware such as a network adapter or SCSI drives without having to fall back on an emulation layer. Incidentally, it is possible to access the network from a non enlightened partition. You just have to use an emulated network adapter. In fact, I recently virtualized a server that was running Windows NT 4.0, and it is able to access the network by using the partition's emulation layer.

So let's get back to my original question. Why does the host operating system disagree with the guest operating systems about the amount of resources that are being used? The answer lies in the way that Hyper-V uses partitioning. Each partition is a completely isolated set of logical resources, and the host operating system lacks the ability to look inside of individual partitions to see how resources are truly being consumed.

At first this would seem to be irrelevant. After all, CPU usage is CPU usage, right? Remember though, that child partitions use virtual processors instead of communicating directly with the physical processor. It is the hypervisor (not the host operating system) that is responsible for scheduling virtual processor threads on physical CPU cores.

Ever since the days of Windows NT, it has been possible to determine how much CPU time is being consumed by monitoring the \ Processor(*)\% Processor Time counter in the Performance Monitor. When you bring Hyper-V into the equation though, this counter becomes extremely unreliable (from the standpoint of the system as a whole), as you have already seen.

If you monitor the \% Processor Time counter from within a guest operating system, you are seeing how hard the virtual processors are working, but in a view that is relative to that virtual machine, not the server as a whole. If you watch this counter on the host operating system, the Performance Monitor is not aware of CPU cycles related to the hypervisor.

As you will recall., the screen captures that I showed you earlier generally reflected extremely low CPU utilization for virtual machines. The fact that the host operating system isn't aware of CPU cycles related to hypervisor activity certainly accounts for at least some of that, but there are a couple of other reasons why CPU utilization appears to be so low.

First, Hyper-V allows you to decide how many virtual processors you want to allocate to each virtual server. This can be a limiting factor in and of itself. For instance, if a server has four CPU cores, and you only allocate one virtual processor to a child partition, then that partition can never consume more than 25% of the server's total CPU resources regardless of how the partition's CPU usage is actually reported.

Things can get a little bit strange when you start trying to allocate more virtual processors than the number of physical CPU cores that the server has. For instance, suppose that you have a server with four CPU cores, and you allocate eight virtual processors. In this type of situation, the virtual machines will try to use double the amount of CPU time that is actually available. Since this is impossible, CPU time is allocated to each of the virtual processors in a round robin fashion. When this occurs, CPU utilization is reported as being very low, because the workload is being spread across so many (virtual) processors.

In reality the virtual machine's performance will be worse than it would have been had a fewer number of virtual processors been allocated. Allocating fewer virtual processors tends to cause CPU utilization to be reported as being higher than it would be had more virtual processors been allocated, but performance ultimately improves because there is a significant amount of overhead involved in distributing the workload across, and then allocating CPU time to all those virtual processors.

Conclusion

As you can see, the values that are reported by the Performance Monitor and by other performance measuring mechanisms vary considerably depending on where the measurement was taken. In Part 2, I will show you why this is the case, and how you can more accurately figure out how what system resources a virtual Exchange Server is actually consuming.

Windows Server Virtualisation: Hyper-V, an Introduction

01 June 2009

by Jaap Wesselius

For SQL Server and Exchange Server, Windows Server Virtualization is going to be increasingly important as a way for the administrator to allocate hardware resources in the most efficient way, to offer more robust services, and for deploying services. Jaap Wesellius starts his new series on Hyper-V by explaining what Hyper-V is, how it relates to Windows Server 2008 and how it compares to ESX, Virtual Server and Virtual PC.

Hyper-V Introduction

Microsoft released Hyper-V, its hypervisor based virtualization product, in the summer of 2008; but what's the difference with Virtual Server? And why is Hyper-V a better product than Virtual Server? And what's the difference with VMware ESX for example? In a series of articles I'll try to explain what Hyper-V is, how it relates to other products and try to give some best practices regarding the use of Hyper-V.

Windows Architecture

Before we take a look at the Hyper-V architecture we take a look at the Windows Server 2008 (and basically all Windows NT servers) architecture. When Window Server 2008 is installed on appropriate hardware two modes can be identified:

- **Kernel mode** – this is a protected space where the kernel, the "heart" of Windows Server 2008 is running and where processes run that interact directly with the hardware, for example the device drivers using buffers allocated in kernel mode. When such a process crashes it is very likely that the server will crash as well which will result in a blue screen of death.

- **User mode** – this is a more protected space where applications are running, for example Microsoft Office, or SQL Server, or Exchange Server. When an application in User Mode crashes only the application stops and the server continues running.

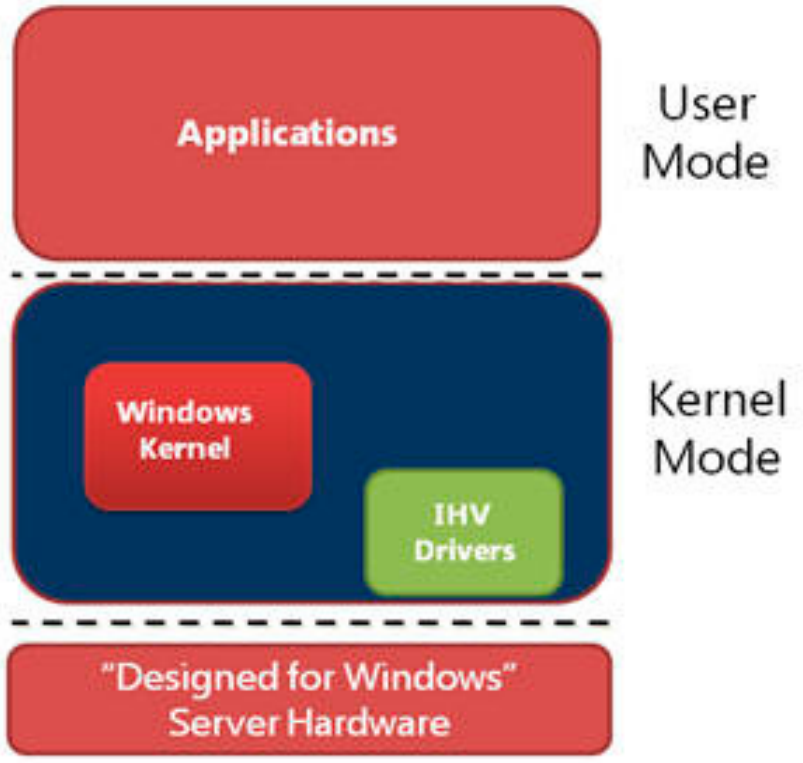

Figure 1. User and Kernel mode running under Windows Server 2008.

When an application needs to access a piece of hardware, for example the hard disk or the network interface the application needs to communicate with the appropriate driver running in Kernel mode. Switching from User mode to Kernel mode is a costly process and consumes a considerable amount of processor cycles. This is known as "mode switching."

Virtual Server and Virtual PC are applications and as such are running in User Mode, the complete environment where the Virtual Machine is running in is emulated. After installing the Virtual Machine additions or when using Hardware Assisted Virtualization some kernel processes are handled directly by the processor. Every piece of hardware the Virtual Machine has to access has to go from User Mode to Kernel Mode and vice versa. The overhead in this scenario is large and will have a large performance impact. The same is true for VMware Server and VMware workstation.

Hyper-V Architecture

Hyper-V is a so called hypervisor. The hypervisor is installed between the hardware and the operating system. Hyper-V is a role in Windows Server 2008 and can only be installed after Windows Server 2008 is installed. When installing the Hyper-V role the hypervisor is "slid" between the hardware and the operating system. Besides the hypervisor a little more is installed as well. The VMBus is installed which is running in kernel mode as well as a Virtual Storage Provider (VSP). Furthermore a WMI provider is installed which is running in User Mode. A VMWorker process is spawn for every Virtual Machine that's started when Hyper-V is running.

> ***Note***
>
> *Hyper-V is only available on Windows Server 2008 X64 edition. Besides X64 capable hardware the server should support hardware virtualization and Data Execution Prevention (DEP) should be enabled on the server. The server's BIOS should support these settings as well.*

After installing the Hyper-V role in Windows Server 2008 the server needs to be rebooted and the server is operational. The original Windows Server 2008 that was installed is turned into a Virtual Machine as well, this one is called the "root" or the "parent partition." It is a very special Virtual Machine since it controls the other Virtual Machines running on the server. I'll get back to this later in this article.

Virtual Machines and the parent partition on Hyper-V are running side-by-side as shown in Figure 2. Virtual Machines are called "child partitions." There are three types of Virtual Machines:

- Hypervisor aware Virtual Machine like Windows Server 2003 and Windows Server 2008.
- Non-hypervisor aware Virtual Machines like Windows Server 2000 and Windows NT4. These Virtual Machines run in an emulated environment.
- Xen enabled Linux kernels (which also support the VMBus architecture). The only one that's available as a standard distribution at this point is SUSE Linux.

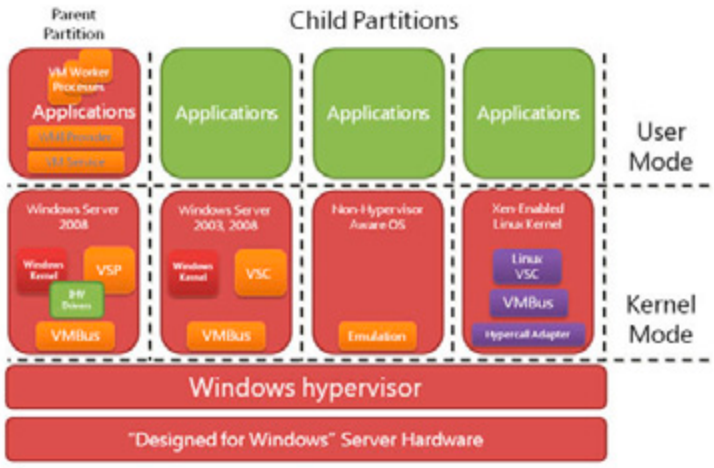

Figure 2. The parent partitions and Virtual Machines in Hyper-V.

Now we're installing a Virtual Machine based on Windows Server 2008. This child partition is running on top of the hypervisor. When the Integration Components are installed the new Virtual Machine can fully utilize the power of Hyper-V. The Integration Components are special Hyper-V drivers, the so called synthetic drivers. Also a Virtual Storage Client (VSC) is installed in the Virtual Machine. These drivers can use the VMBus structure. The VMBus is a point-to-point in-memory bus architecture, running fully in kernel mode. An application running in this Virtual Machine wants to access the network interface or a local disk on the parent partition and makes a request to do so. This request goes from user mode to kernel mode and is sent via the VSC over the VMBus to the VSP. From here the request is sent to the appropriate device. No additional mode switching is needed and this is truly a very fast solution.

A non hypervisor-aware Virtual Machine, for example a Windows NT4 server does not have the Integration Components and a VSC. Everything is emulated, and it is emulated in the VMWorker processes. These processes are running in user mode on the parent partition.

When an application on this Virtual Machine make a request to the local disk the request is sent to the driver running in kernel mode in the Virtual Machine. This is intercepted and sent to the emulator on the parent partition which in turn sends it to the local disk. This means that three additional mode switches are needed. One in the Virtual Machine, from the Virtual Machine to the host partition and on the actual host partition from user mode to kernel mode. This creates additional overhead which results in reduced performance for emulated Virtual Machine. Virtual Server also makes use of a fully emulated environment and thus suffers from the same performance hit.

Virtual Machines running on SUSE Linux and have the Linux Integration Components installed can also fully utilize the new VMBus architecture and thus fully utilize the server's resources. Other Linux clients use a fully emulated Virtual Machine, just like the NT4 example.

Micro-kernelized hypervisor

One difference between ESX and Hyper-V is the type of hypervisor. Microsoft uses a micro-kernelized hypervisor where VMware uses a monolithic hypervisor. So what are the differences between these two?

A micro-kernelized hypervisor is a very thin hypervisor (less then 800 Kb) when an absolute minimum of software in the hypervisor. Drivers, memory management etc. needed for the Virtual Machines are installed in the parent partition. This means that Windows Server 2008 with the appropriate, certified hardware drivers can be used for a Hyper-V server.

A monolithic hypervisor is a hypervisor that contains more software and management interfaces. Network drivers and disk drivers for example are part of the hypervisor and not of the parent partition. This automatically means that only servers that are certified by VMware and have certified drivers can be used for an ESX Server.

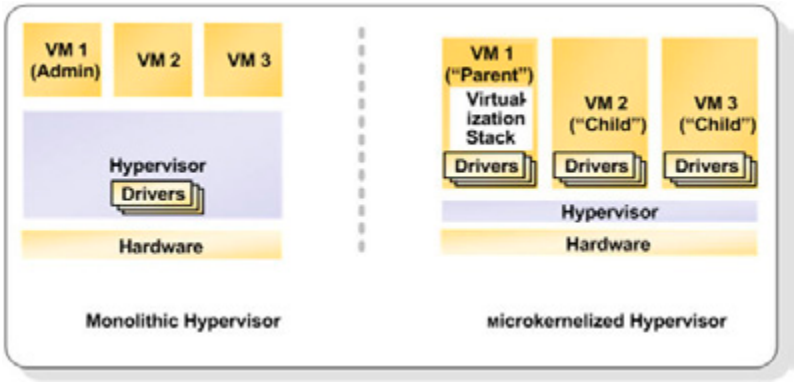

Figure 3. Monolithic versus Micro-kernelized Hypervisor.

Both solution have pros and cons, time will tell which solution is the best one and offers the best performance and scalability.

Security

After installing the Hyper-V role in Windows Server 2008 the original Windows installation automatically turns into a Virtual Machine, the so called parent partition or root. After logging in to the parent partition this just looks like an ordinary Windows Server 2008. But it controls all other Virtual Machines running on the server, so special care needs to be taken.

When the parent partition is compromised with a virus or a Trojan horse not only the parent partition is under somebody else's control, but potentially all Virtual Machines running on this server. The Hyper-V manager is available on this server as well as all WMI interfaces that control the Virtual Machines running on this server. It is a best practice to install no other software on the parent partition and not use it for example for browsing on the Internet. All applications and software should be installed on Virtual Machines and NOT on the parent partition.

A better solution is to use Windows Server 2008 Server Core. This is very minimalistic instance of Windows Server 2008 with few software or services installed. Also the explorer is not present on the Server Core and after logging in to this Server Core only a Command Prompt is shown. Some small GUI's are available though, for example the data-time applet to set the data and time on the server. Managing a Windows Server 2008 Server Core is definitely more difficult than management a "normal" server with a Graphical User Interface (GUI) but once you're used to it and can fully manage it is much safer due to the reduced attack surface.

Microsoft made a couple of design decisions with respect to security. Not using shared memory for example is such a decision. When using shared memory you can over commit memory on your host server. Over committing is assigning more memory to Virtual Machines than there's available on the host server. By sharing memory pages between Virtual Machines it is possible to achieve this. Although this is definitely true it was a security decision made by Microsoft to not use this feature.

Virtual Machines can be compromised as well and this is also a situation you do not want to occur. But when a Virtual Machine is compromised it is not possible to access the hypervisor to take over the host server. It is also not possible to access other Virtual Machines.

This also means that when you have to copy data from one Virtual Machine to another it's just like physical machines. You have to copy this data across the network using file shares. The only option that's possible is to copy plain text between your Parent Partition and a Virtual Machine using the "Copy Text" option in the Hyper-V Manager.

Integration Components

When installing a Virtual Machine initially this is running in an emulated environment. As explained earlier this is not the most efficient way of running a Virtual Machine. After the initial installation you have to install the Integration Components. Open the Hyper-V Manager, select the Virtual Machine, choose Action and select "Insert Integration Services Setup Disk." This will install the Integration Components in your Virtual Machine. When finished reboot the Virtual Machine and it's done.

When installing the Integration Components the synthetic drivers are installed in the Virtual Machine, making it possible to have the Virtual Machine communicate via the VMBus architecture. This will speed up performance dramatically. You can see the Integration Components using the Virtual Machine's device manager:

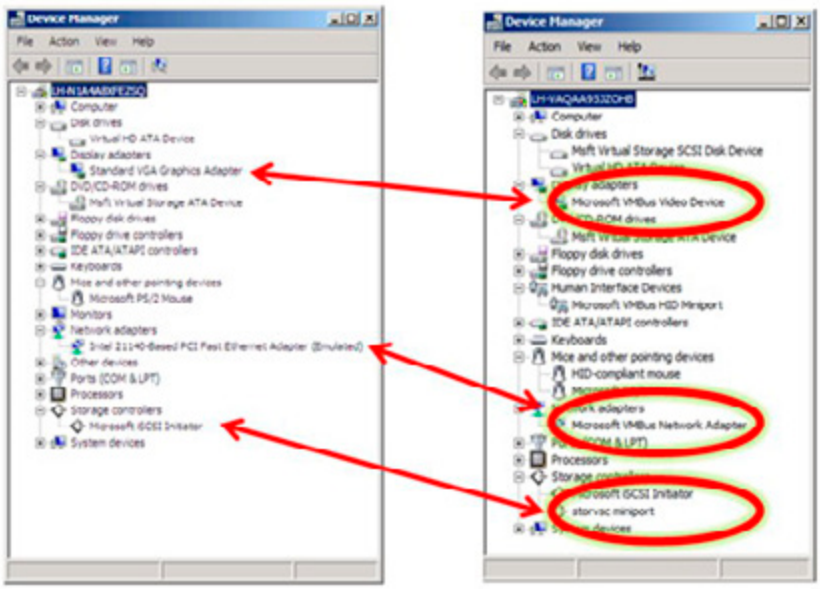

Figure 4. After installing the Integration Components the emulated hardware is replaced by Hyper-V specific hardware.

Besides the synthetic drivers the Integration Components offer more services to Virtual Machines, like time synchronization between the root partition and the Virtual Machine, backup options (volume snapshot) and operating system shutdown from the Hyper-V Manager.

Server Virtualization Validation Program

Microsoft has always been reluctant in supporting virtualized application, especially in the timeframe before Hyper-V. In those days Microsoft only had Virtual Server as virtualization software while VMware was offering ESX Server.

When Hyper-V entered the virtualization market Microsoft had not only to support their own software and application running on Hyper-V, but also their applications running on other virtualization software, from other vendors that is. Microsoft has setup a program where other vendors can have their solutions validated, this program is known as the Server Virtualization Validation Program (SVVP). VMware's ESX Server for example is validated in this program and all recommendations made for running Microsoft applications under Hyper-V also apply for running these applications under ESX Server. When issues are submitted by customers in Microsoft Product Support Services Microsoft does not make a difference between ESX Server and Hyper-V when it comes to troubleshooting. You can find more information regarding the SVVP program on the Microsoft website: http://www.windowsservercatalog.com/svvp.aspx

Conclusion

Microsoft Windows Server 2008 Hyper-V was released in the summer of 2008 and is Microsoft first real hypervisor virtualization solution. It is not an emulated environment like Virtual Server or Virtual PC, but as a hypervisor solution it "sits" between the hardware and the Operating System. With the Integration Components installed you can fully use the functionality offered by Hyper-V. You have to secure the Parent Partition as much as possible to prevent compromising the complete system.

In the next articles I will talk more about the Hyper-V best practices, deploying Virtual Machines, using the System Center Virtual Machine Manager (VMM) 2008 and the "high availability" options and why these aren't really high available in the current release of Hyper-V.

A Beginner's Guide to Virtualizing Exchange Server – Part 2

10 June 2009

by Brien Posey

It isn't easy to measure the consumption of physical resources by servers in virtual machines, since each partition has its' own virtualised view of system resources. Not only that, but there is a subtle difference between virtual processors and physical processors. Brien Posey explains the special performance counters that can be used to get an accurate assessment, and goes on to describe how to test your physical servers to see if they are good candidates for virtualization.

If you read my first article in this series, then you got something of a crash course in Hyper-V's architecture. Of everything that I covered in that article though, there are two main points that you need to keep in mind:

- When it comes to virtualizing Exchange Server, it is critical that you monitor resource consumption so that you can ensure that there are sufficient hardware resources available to effectively service Exchange and any other virtual machines that may be running on the server.
- Each of the various resource monitoring mechanisms that I showed you tells a completely different story regarding how much of the server's physical resources are actually being consumed.

In other words, it is important to find out how much of the server's resources are being consumed, but you are not going to be able to do so in the usual way.

So Why the Discrepancy?

In my previous article, I showed you that when you monitor the amount of CPU time that the server is consuming, you will get different results depending upon where the measurement was taken. I never really said why though. The reason for this is that neither the root nor the child partitions control the APIC timer. In essence, this means that each partition has its own virtualized view of the system's available resources. It isn't that one set of performance counters is right and another set is wrong. All of the counters are correct from their own frame of reference. Remember that the goal of having multiple partitions is to isolate the virtual machines from one another. Therefore, no one single partition (including the root partition) has the full picture of how the system's resources are being used.

Monitoring CPU Utilization

Traditionally, the **\Processor(*)\% Processor Time** performance monitor counter has been used to find out how much CPU time a server is consuming. As you have seen though, this counter does not tell you the whole story when you use it in a Hyper-V environment.

There are ways of measuring CPU utilization in a Hyper-V environment, but before I show you how, it is important that you understand the difference between virtual processors and logical processors. Logical processors map directly to the number of CPU cores that are found in the system. For example, if you have a server that has two quad core CPUs installed, then the server has a total of eight CPU cores. Since logical processors correlate directly to the number of CPU cores, the server also has eight logical processors.

Hyper-V partitions do not directly use logical processors. Instead, they use virtual processors. As I explained in the first part of this series, you can actually allocate more virtual processors to your partitions than the machine has logical processors. Therefore, if you want to know how much of the server's CPU resources are actually being consumed then you need to look specifically at the virtual processor utilization.

If you have ever used the Performance Monitor within the host operating system, then you may have noticed that there are a number of Performance Monitor counters that are specifically related to Hyper-V. One of the most useful of these counters is the **\Hyper-V Hypervisor Logical Processor (_Total)\% Total Run Time** counter. This counter shows you the cumulative CPU consumption across all virtual processors. The output from this counter is presented as a percentage of the server's total CPU resources, just as the \ **Processor(*)\% Processor Time** counter does on a non virtualized system.

Microsoft even gives us some guidelines as to what is considered to be an acceptable level of utilization. According to Microsoft, if this counter is reflecting a value of 60% or less then the server's CPU utilization is healthy. If the counter reflects a value of 60% to 89% then the server needs to be monitored, because it is in danger of exhausting available CPU resources. A value of 90% or above is considered to be a critical state in which the server's performance will suffer.

That is how you measure the server's CPU consumption as a whole. There are a couple of other Performance Monitor counters that you need to know about though; the \Hyper-V Hypervisor Logical Processor(_Total)\% Total Run Time and \Hyper-V Hypervisor Virtual Processor(_Total)\% Total Run Time counters.

The \Hyper-V Hypervisor Logical Processor(_Total)\% Total Run Time counter shows you how hard the CPU is working from a logical processor standpoint, while the \Hyper-V Hypervisor Virtual Processor(_Total)\% Total Run Time counter looks at the workload from a virtual processor standpoint. Generally speaking, these two counters should be balanced so that they display roughly about the same value.

If you find that logical CPU utilization is high, but that virtual CPU utilization is low, then it is often a good indication that there are more virtual processors allocated to the server's partitions than the machine has logical processors. Although you can create more virtual CPUs than the number of logical CPUs that the server contains, you should try to use a one to one ratio when possible.

If you find that you have allocated an excessive number of virtual processors then Microsoft recommends using the \Hyper-V Hypervisor Virtual Processor(*)\%Guest Run Time counter to determine which virtual processors are consuming the most CPU resources. You can then begin removing lesser used virtual processors from guest machines in an effort to try to reach a one to one virtual to logical processor ratio.

On the flip side, you may occasionally run into situations in which the logical processor utilization is low, but the virtual processor utilization is high. When this happens, you should consider allocating additional virtual processors to the partitions that need them most. Of course you have to make sure that the operating system that is running within that partition will support the extra processors before you actually begin allocating any additional virtual processors. If you find that you have CPU resources to spare, but the guest operating systems do not support adding additional processors, then you might consider adding additional virtual machines to the physical server to make better use of the server's resources.

Planning for Exchange

The techniques that I have shown you work well for figuring out how much of a server's resources are being consumed, but they may not do you a lot of good if you are trying to figure out whether or not a server has the resources to host a virtualized Exchange Server. Fortunately, there are a couple of tricks that you can use to gauge a server's capabilities before you try to virtualize your Exchange Server.

The first thing that you need to keep in mind is that from a hardware requirement standpoint, Microsoft doesn't make any differentiation between a physical server and a virtual server. For instance, suppose that the Exchange Server role that you plan to deploy requires 2 GB of disk space. That 2 GB requirement remains the same whether you are going to be deploying Exchange on a physical server or on a virtual server.

Of course adhering to the hardware requirements for Exchange Server will only get you so far. You have probably noticed that I have spent a whole lot of time in this series discussing how you can figure out how much of a server's CPU time is available. The reason for that is that

Hyper-V makes it really easy to create a virtual server that adheres to specific hardware specifications. For instance if you need a virtual machine with 2 GB of RAM, three hard drives that are 100 GB each, and two network adapters, then you can easily create such a machine by using the Settings option in the Hyper-V Manager.

CPU resources are a little bit more difficult to allocate though. It's easy to tell Hyper-V how many virtual processors you want to assign to a virtual machine, and you even have the option of allocating CPU resources as a percentage of the server's total CPU resources. What you can't do however, is to tell Hyper-V to give your virtual machine the equivalent of a 2.8 GHz quad core processor. In other words, you can tell Hyper-V how many virtual processors to assign to a virtual machine, or you can tell it to use a percentage of the overall CPU resources, but there isn't a way of requesting a specific amount of processing power. That's why CPU monitoring and capacity planning is so important.

Of course this still leaves the question of how you can determine whether or not your host server is up to the job of running a virtualized instance of Exchange Server. If your Exchange Servers are currently running on physical hardware then the assessment process is easier than you might think.

The Microsoft Assessment and Planning Toolkit

Microsoft offers a free utility called the Microsoft Assessment and Planning Toolkit . This utility can be used for a lot of different purposes, but one of the things that you can use it for is to test your physical servers to see if they are good candidates for virtualization.

Unfortunately, this utility isn't Exchange Server aware, so it isn't going to tell you that you can't virtualize one of your Exchange Servers because it isn't running a configuration that is supported in a virtual environment. What it will do, is collect performance data from your physical servers and then use that performance data to determine whether or not the server is a good candidate for virtualization based on its resource consumption, and on the resources that are provided by your host server.

I recommend running the Assessment and Planning Toolkit on a computer that's running Windows Vista, rather than running it directly on one of your servers. That way, you can minimize the utility's performance impact on the servers that you are going to be benchmarking. However, the machine that you run the utility on needs to have Microsoft Office installed. You won't even be able to install the utility unless you install Microsoft Office first.

Gathering Performance Data

The first step in analyzing your physical servers is to create a simple text file containing the NetBIOS names of the physical servers that you want to analyze. Just use Notepad to create the list, and put each server name on its own line.

Now that you have created a list of servers to analyze, open the Microsoft Assessment and Planning Solution Accelerator. When the console opens, click on the Select a Database link. You should now see a dialog box asking you if you would like to select an existing database, or if you would like to create a new database. Since this is the first time that we have used the Assessment and Planning Toolkit, choose the option to create a new database. Provide the database with a name, and click OK.

Next, select your newly created database from the console's Assessment pane, and then click the Prepare Recommendations for Server Consolidation Using Windows Server 2008 Hyper-V or Virtual Server 2005 R2. You will now see a message explaining that the wizard that you have chosen to use requires performance data. The message contains a link that you can use to capture performance data for your computers. Go ahead and click this link.

At this point, you will be prompted to supply the name and path of the text file that you created earlier. After doing so, click Next and you will be prompted to provide a set of WMI credentials for the computers that you have specified. You can use the same credentials for each computer on the list, or you can specify separate credentials for each machine if necessary. When you are done, click Next.

You should now see a screen that asks you when the benchmark tests should complete. By default the tests are set to run for an hour, but you can specify a different period of time if you wish. Keep in mind though, that if you set the testing period to be too short, then the wizard may not have enough data to make a good recommendation. Click Next, and you will see a summary of the settings that you have chosen. Take a moment to review these settings, and then click Finish. The performance monitoring process will now begin.

Analyzing the Performance Data

Now that we have completed the data collection process, it is time to analyze the results. To do so, click on the Prepare Recommendations for Server Consolidation Using Windows Server 2008 Hyper-V or Virtual Server 2005 R2 link. When you do, Windows will launch the Server Virtualization and Consolidation Wizard.

The first thing that you will need to do when the wizard starts is to select the virtualization product that you plan on using. The wizard allows you to select either Hyper-V or Virtual Server 2005 R2.

Click Next, and you will be asked to provide some details regarding the host machine's CPUs. The wizard asks some fairly detailed questions regarding things like the sizes of the level 2 and level 3 caches, and the bus speed. Granted, answering these questions isn't rocket science, but it may require you to look up your server's specs. If you don't have a way of getting the specific details for your server's CPUs, then you can just select the make and model of the CPUs that the server is using, and the wizard will fill in the details for you. They might not be completely accurate in every situation, but they should at least be close enough to make the wizard's results reasonably accurate.

After you have filled in your CPU information, click Next, and you will be prompted to enter some details about your host server's disk configuration. This information is fairly easy to fill in. You just need to know the type and speed of the disks that will be used, and what type of array configuration (if any) will be used. When you are done, click Next.

The following screen prompts you to enter the number of network adapters that are installed in the host server, and the speed of those adapters. This screen is also the place where you tell the wizard how much memory is installed in your host server.

Click Next, and you will be asked if you would like to set a limit as to the number of virtual machines that the server can host. Personally, I recommend setting a limit, although the actual limit that you set is going to depend on what you feel comfortable with. The reason why I recommend setting a limit is because the recommendations that the wizard is going to make are based on an hour's worth of performance monitoring. Your peak usage may cause significantly higher resource consumption than what was recorded during the monitoring period. Besides that though, it is smart to leave some room for future growth.

Click Next, and you will be asked to provide a text file containing a list of computer names. You should use the same text file that you used when you gathered performance data on your servers.

Click Next one more time, and you will see a screen that displays all of the settings that you have entered. Take a moment to verify that you have entered this information correctly, and then click Finish. The Microsoft Assessment and Planning Solution Accelerator will now create the requested reports.

Viewing the Results

When the report generation process eventually completes, click Close, and you will be returned to the main console screen. You can access the reports by clicking the *View Saved Reports and Proposals* link. When you do, you will find that the wizard has created two documents. The first document is an Excel spreadsheet that contains the raw performance data from each of the servers that you are considering as virtualization candidates. The second document is a Microsoft Word document that contains detailed recommendations for virtualizing your proposed servers. These recommendations are based on the performance data that was collected and the information that you provided about your host server.

One thing that is important to understand is that the wizard does not perform any types of benchmarks on your host server. Therefore, the recommendations that the wizard makes are only as good as the information that you provide it with.

Conclusion

There are obviously many more virtual server resources that you can monitor other than just the CPU. Keep in mind though, that once you understand the basic concepts involved in monitoring CPU resources, many of those same concepts can be applied to monitoring other types of resources.

Increasing the Availability of Virtualized Applications and Services

22 October 2009

by Nirmal Sharma

By using a virtualized clustering computing environment with failover, you can improve server availability without using as many physical computers. A group of independent computers can work together to increase the availability of virtualised applications and services. If one of the cluster nodes fails, another node takes over to provide the service without disrupting the service. Nirmal Sharma explains the failover process under Hyper-V and how to improve the performance of a failover.

This article explains the internal process behind the Hyper-V Virtual Machine Resource DLL and the functions used to interact with cluster components to improve the failover process for virtual machines.

Most of the article talks about Hyper-V Resource DLL. It doesn't really show how to cluster Virtual Machines or how to configure Quick Migration in Hyper-V for Virtual Machines. Instead the article focuses more on the Hyper-V Resource DLL for Virtual Machines and the Failover Process for Virtual Machines running on Hyper-V Server.

Terms

Before we move ahead, let's define some important terms that we will be using.

Cluster Service
The Cluster Service is the main component of the Clustering Software which handles the communication between Resource Monitor and its managers. All the clustering managers run under the Cluster Service.

Resource Monitor
The Resource Monitor is part of the Clustering Software. This runs under the Cluster Service (Clussvc.exe) to handle the communications between the Resource DLL and the Clustering Software.

Resource DLL
The Resource DLL ships with cluster-aware applications. The functions executed by the Clustering Software are supported by the Resource DLL. The main function of the Resource DLL is to report the status of the application resources to the Clustering Software and execute the functions from its library as and when needed.

Cluster Configuration Database
The Cluster Configuration Database is a registry hive that contains the state of the cluster. It is located at HKLM\Cluster at registry.

Resources
A resource is an entity that can provide a service to a client and can be taken offline and brought online by the Clustering Software. A resource must have its associated Resource DLL so that the Resource Monitor can communicate with the resources using this DLL. The Virtual Machines running on Hyper-V can be configured as a Resource in the Cluster. The Resource DLL for Virtual Machines is VMCLUSRES.DLL

Windows Clustering

Microsoft introduced its first version of clustering software in Windows NT 4.0 Enterprise Edition. Microsoft has significantly improved the clustering software in Windows 2000, Windows Server 2003 and Windows Server 2008. There are two types of clustering technologies: Server Cluster (formerly known as MSCS) and Network Load Balancing Cluster (NLB). MSCS or Server Cluster is basically used for High Availability. NLB is of course used to load balance the TCP/IP traffic. The MSCS or Server Cluster capability is also known as Failover Clustering. The support for Virtual Machines running on Hyper-V in a cluster is available only with Failover Clustering.

Virtual Machines and High Availability

Support for Clustering Virtual Machines was introduced in Windows Server 2008 running the Hyper-V Role and has been continued in the versions that followed.

Windows Clustering includes many components such as Cluster Service, Resource Monitors, Node Manager, Membership Manager, Event Log Processor, Failover Manager, and Cluster Database Manager. The whole purpose of Failover clustering is to provide high availability of application resources. Clustering doesn't get involved in deciding how much CPU and Memory should be utilized by an application.

An application running in the clustering environment must be cluster-aware. A cluster-aware application supports the functions executed by the cluster service or its components as shown in Figure 1.1. There is no way for Cluster Service to know about the availability of resources of an application in the cluster unless the application is cluster-aware. For example, if a node holding the application resources fails, the Cluster Service running on that node must be notified in order to start the failover process for the application's resources. Cluster Service does this by receiving the responses from the Resource Monitor. The Resource Monitor tracks the Virtual Machines with the help of Resource DLLs provided by Hyper-V Role.

You cannot cluster Virtual Machines running on Virtual Server. The Virtual Machines running on Virtual Server do not provide any Resource DLL which can be used with the clustering software to make them highly available. On the other hand, the Virtual Machines running on Hyper-V are fully cluster-aware Virtual Machines, supporting/responding to all functions executed by the cluster service. The Resource DLL of Hyper-V Virtual Machines, which supports all the functions, is **VMCLUSRES.DLL.** Hyper-V provides only one DLL for its Virtual Machines in the cluster. There are not any other DLLs provided by the Hyper-V Role. We will discuss that DLL in detail in this article.

Tip

A Resource DLL is a separate application component that is specifically written to support cluster functions (for example, Open, Terminate, Online, Offline, Retry and so on).

The Clustering Software Resource Monitor tracks the Hyper-V Virtual Machines availability through VMCLUSRES.DLL by performing two checks: IsAlive and LooksAlive. Implementing these tests is application specific and hence why cluster-aware applications are expected to provide their resource DLL. The Cluster Server doesn't need to know about application-specific functions. It just executes the functions provided by the Resource DLLs. Hyper-V implements many other functions in its Resource DLL. The functions are shown in **Figure 1.1.** These functions are Hyper-V virtual Machine-specific and not related to clustering in any way.

Tip

The two basic checks (**IsAlive** *and* **LooksAlive**) *are supported by every Resource DLL or cluster-aware application.*

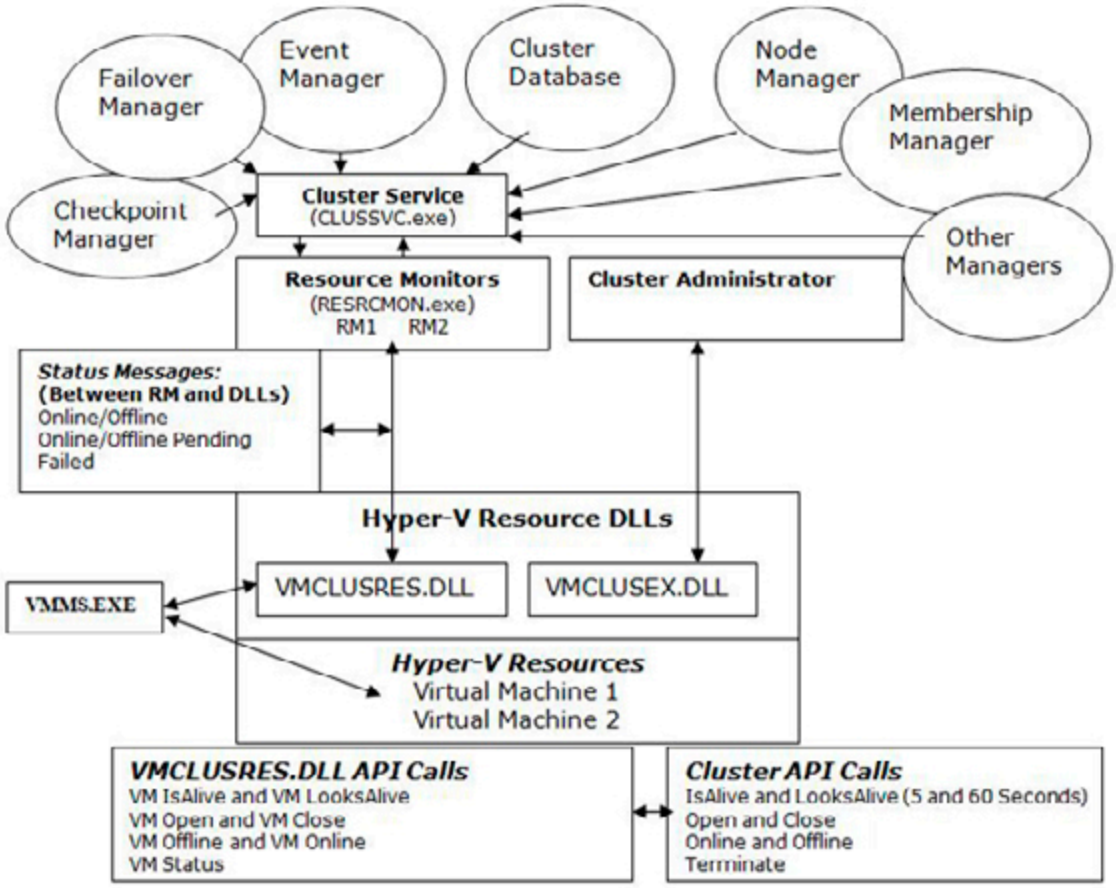

FIGURE 1.1–Cluster Components and Hyper-V Cluster Resource DLL.

In **Figure 1.1** you can see the DLL VMCLUSRES.DLL is installed when the Hyper-V Role is enabled initially. Before you can cluster Virtual Machines running on Hyper-V, you need to install the Failover Clustering Software on Windows Server 2008 or 2008 R 64-bit edition. After installation is completed, you click on "Services and Applications" in Failover Cluster Management and then select the "Virtual Machines" as Cluster resource.

Tip

If you don't see "Virtual Machines" then try running the following commands. This DLL must be registered before you can cluster Virtual Machines.

Regsvr32.exe /u VMCLUSRES.DLL
Regsvr32.exe VMCLUSRES.DLL
The above command reregisters the VMCLUSRES.DLL with the Failover Clustering Software.

The next DLL is VMCLUSEX.DLL. This DLL works as a proxy between the Cluster Administrator and the Hyper-V Manager. The main function of this DLL is to provide interfaces to configure and control Virtual Machines configuration parameters and screens. If this DLL is missing or corrupted you can't access Virtual Machines. VMCLUSEX.DLL doesn't implement any cluster-specific control functions. As an example, when you right click on a Virtual Machine resource using the Failover Cluster Manager, you will get "Bring this Virtual Machine Online" option to start the Virtual Machine. The same will be reflected in Hyper-V Manager. You will see the Virtual Machine starting in the Hyper-V Manager also.

VMMS.EXE which is the main process of Hyper-V needs to know the status of Virtual Machines running on the Hyper-V Server. The Resource DLL is written to update the status of the Virtual Machines in a cluster to VMMS.EXE. VMMS.EXE, in turn, shows the status of each Virtual Machine in Hyper-V Manager.

VMCLUSRES.DLL which sits between Resource Monitor and Virtual Machines plays an important role in the failover process. Without this DLL Hyper-V cannot function as a cluster-aware application.

Tip

A malicious code running in your system may corrupt the DLL files.

- Re-run the Hyper-V Setup (disabling and enabling the role).
- Copy VMCLUSRES.DLL from a working computer.

Figure 1.1, above, also shows the functions defined in VMCLUSRES.DLL. The Hyper-V Virtual Machine-Specific functions are mapped with the cluster-specific functions. For example, Cluster's **IsAlive** and **LooksAlive** functions are mapped with VM **IsAlive** and VM **LooksAlive** respectively. However, there are no static mappings defined within VMCLUSRES.DLL. VMCLUSRES.DLL knows which function to execute. The same way, other Virtual Machines functions are also mapped to related cluster functions as shown in Figure 1.1.

VM **IsAlive** and VM **LooksAlive** functions are executed by VMCLUSRES.DLL at a predefined interval. Most of the monitoring task is done by performing a VM **IsAlive** query. VM **IsAlive** is implemented in such a way that it performs all the checks for Hyper-V Virtual Machines. It checks to make sure all the:

- Virtual Machines in cluster are online.
- Virtual Machines are configured with correct dependencies.
- The registry entries for Virtual Machines resources are configured correctly.

VM **LooksAlive** is used to perform a thorough check on the Virtual Machines in the cluster. This check might take some time as it includes checking the configuration of Virtual Machine, Virtual Machine Configuration file location (XML), VHD location, etc. It might take some time for **LooksAlive** to perform these checks and report back the status to the Resource Monitor. To avoid the delays in reporting, the Resource Monitor cluster component depends on the results reported by **IsAlive** which is configured to execute every 5 seconds by default. **IsAlive** only checks the status of Virtual Machine in the Cluster (e.g. Online or Failed). Based upon that, the action is taken by the Resource Monitor. Think of a situation where only **LooksAlive** is used to get the status of Virtual Machines in the Cluster. This may result in slightly more downtime of the Virtual Machines as **LooksAlive** calls are executed every 60 seconds! Now, you could ask why not decrease

the interval of **LooksAlive**. Well, if you do so, you would see performance issue on the cluster. Please note that the Resource Monitor component of Clustering Software executes **IsAlive** and **LooksAlive** queries against the whole Cluster Group. It is the responsibility of the Resource DLL (VMCLUSRES.DLL) to execute VM **IsAlive** and VM **LooksAlive** against its Virtual Machine resources. By default, the **IsAlive** check is performed every 5 seconds and **LooksAlive** check is performed every 60 seconds as shown in Figure 1.2 below.

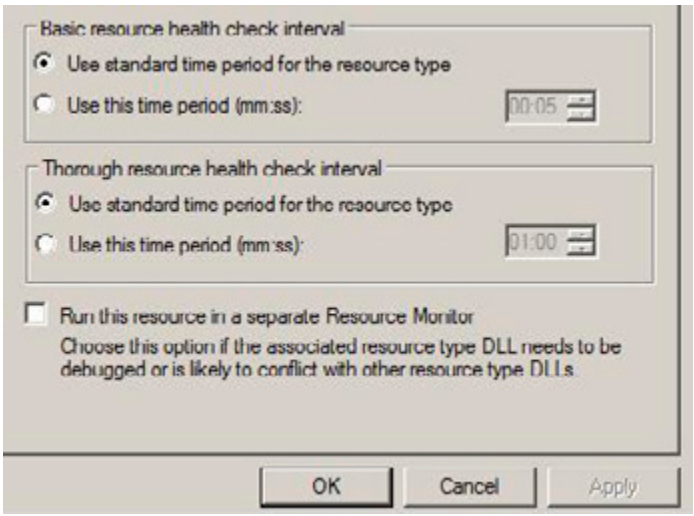

FIGURE 1.2: IsAlive and LooksAlive Interval of Virtual Machine Resource.

The default interval can be changed per Virtual Machines to improve failover response time as shown above in **Figure 1.2.**

In previous versions of Windows Clustering, it was not possible to define the **IsAlive** and **LooksAlive** interval per Resource. Now, starting with Windows Server 2008 cluster, it is possible to define the **IsAlive** and **LooksAlive** intervals per resource.

When you setup a cluster for the first time, the Cluster Service running on the node takes a snapshot of the cluster configuration and saves it in HKLM\Cluster key. This Key contains the cluster configuration such as the resource name, their GUID, node holding the resources and status. This is generally called cluster configuration database. As an example, for Virtual Machines it includes the following:

Resource Name	GUID		Node Name	Status	PersistentState
Virtual Machine 1	{GUID1}		Node1	Online	1
Virtual Machine 2	{GUID2}		Node1	Online	1
Virtual Machine 3	{GUID3}		Node1	Offline	0

The **PersistentState** keeps the status of the Resources or Virtual Machines in the Cluster. The above shown Status column is just for your reference. The **PersistentState** 1 means Online and 0 means Offline. The "Status" column is not stored as a registry entry.

This is also shown in the Cluster Registry hive:

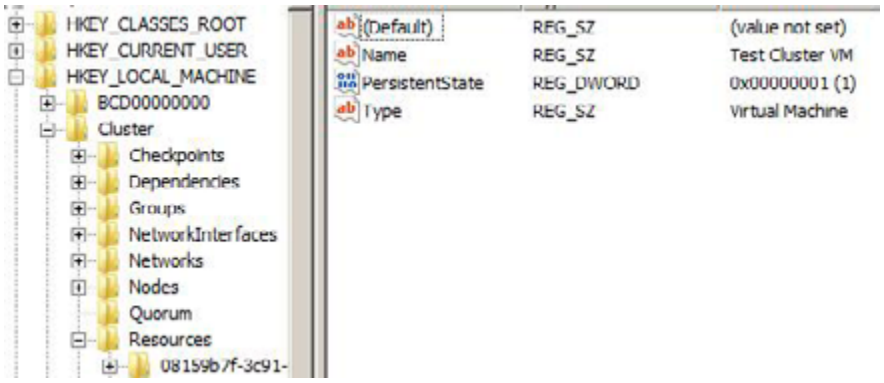

FIGURE 1.3: PersistentState Entry in the Cluster Registry for Virtual Machine.

As you can see in Figure 1.3, the **PersistentState** registry entry value of Virtual Machine "Test Cluster VM" is 1 which indicates that the Virtual Machine is Online in the cluster.

Before the Resource Monitor executes any cluster function against the Virtual Machines or Cluster Groups, it looks at the cluster configuration database to check the status of all resources and their GUIDs. For example, let say we have a cluster group named "HyperV VMs." All the Virtual Machines of Hyper-V reside in this group. When **IsAlive** interval expires (5 seconds by default), the Resource Monitor executes the **IsAlive** call against the "Hyper-V VMs" Cluster Group. It hands over the Resource GUID and Status to the Hyper-V Virtual Machines Resource DLL (VMCLUSRES.DLL). VMCLUSRES.DLL in turn executes the VM **IsAlive** call to check the Virtual Machines availability. Please note that VMCLUSRES.DLL doesn't really know about the status of Virtual Machines. It is the Resource Monitor who supplies this information to VMCLUSRES.DLL.

Next we look at VM Open, VM Close, VM Online and VM Offline. These functions are called whenever Virtual Machines are moved across Hyper-V Servers or taken offline/online or when there is the need to call them. For example, you might want to take a Virtual Machine offline for maintenance purposes on a Hyper-V node. In that case, the Resource Monitor executes the Offline function and in turn VMCLUSRES.DLL executes the VM Offline function to take the Virtual Machine offline. The same will be updated to the VMMS.EXE process in background so that it is aware of the Virtual Machine status. We will discuss these functions later in this article. As a whole, these functions are executed by the Cluster Service and supported by the Hyper-V Resource DLL. That's why Hyper-V Virtualization are known as pure cluster-aware Virtualization Software!

The Resource Monitor determines the state of Virtual Machines by checking the **PersistentState** value at the registry. This value could be either 1 or 0. 1 is for Online and 0 is for Offline. For example, if you stop a Virtual Machine on a cluster node, the value 0 is set for that service or resource at the registry. If you stop the Virtual Machine using command line or Hyper-V Manager, the value is still updated in the Cluster Configuration Database. It is because Resource DLL of Hyper-V and VMMS.EXE always talk to each other to get the status of Virtual Machines and update accordingly in the Cluster Configuration Database. When you stop a Virtual Machine using a command line or WMI Script, you are actually interacting with VMMS.EXE service which, in turn, executes the Stop command on behalf of you. The status of Virtual Machine is updated in the Cluster Configuraiton Database. This may not work for other applications in the cluster. As an example, Exchange Server. Operations occurring out of the cluster for Exchange Server resources are not reflected at the cluster configuration database. In this case, the **IsAlive** query may not function correctly. The value supplied by the resource monitor will indicate that the Resources are running. Thus **IsAlive** will not take any action against the stopped Cluster Resources. The value is updated in the Cluster Configuration Database only when the **LooksAlive** is executed which performs a thorough check for the resources. The thorough check includes checking the Exchange Services.

How does Hyper-V Virtual Machine Resource DLL help in the failover process?

The status messages shown above Figure 1.1 are generated through **IsAlive** calls. When the **IsAlive** interval expires, the Resource Monitor executes the Cluster **IsAlive** calls. The Hyper-V Cluster Resource DLL in turn executes VM **IsAlive** against all Virtual Machine Resources. The messages returned by these calls include one of the following:

- Online/Offline.
- Online/Offline Pending.
- Failed.

The above status messages are passed back to the Resource Monitor. In turn this reports the need to take any action to the Cluster Service.

As shown in Figure 1.1, the Resource Monitor sits between the Hyper-V Resource DLL and the Cluster Service. Any calls made to Hyper-V Virtual Machine Resources have to take place at VMCLUSRES.DLL first. For example, if the Cluster Service needs to check the availability of Hyper-V Virtual Machine resources, it will make a call to the Resource Monitor; in turn this will ask VMCLUSRES.DLL to check the status of the Hyper-V Virtual Machine Resources and report back. If the Resource Monitor doesn't receive any response from VMCLUSRES.DLL or it cannot detect the Virtual Machine availability, it will pass the status back to Cluster Service. Cluster Service then passes this status message to related Managers as shown in above figure. Managers take the action as per the status passed by lower layer components. The status message could indicate a failure of Virtual Machine resources or could indicate a simple status message. These messages and cluster actions are discussed later in this article with an example.

In addition, if functions executed by the Resource Monitor do not exist in the Resource DLL, the request is simply discarded and no operation is carried out.

Hyper-V Server doesn't really utilize its own mechanism to failover the Virtual Machines on the surviving node. Instead Resource DLLs are written to "support" the failover process. The following figure shows a simple failover process:

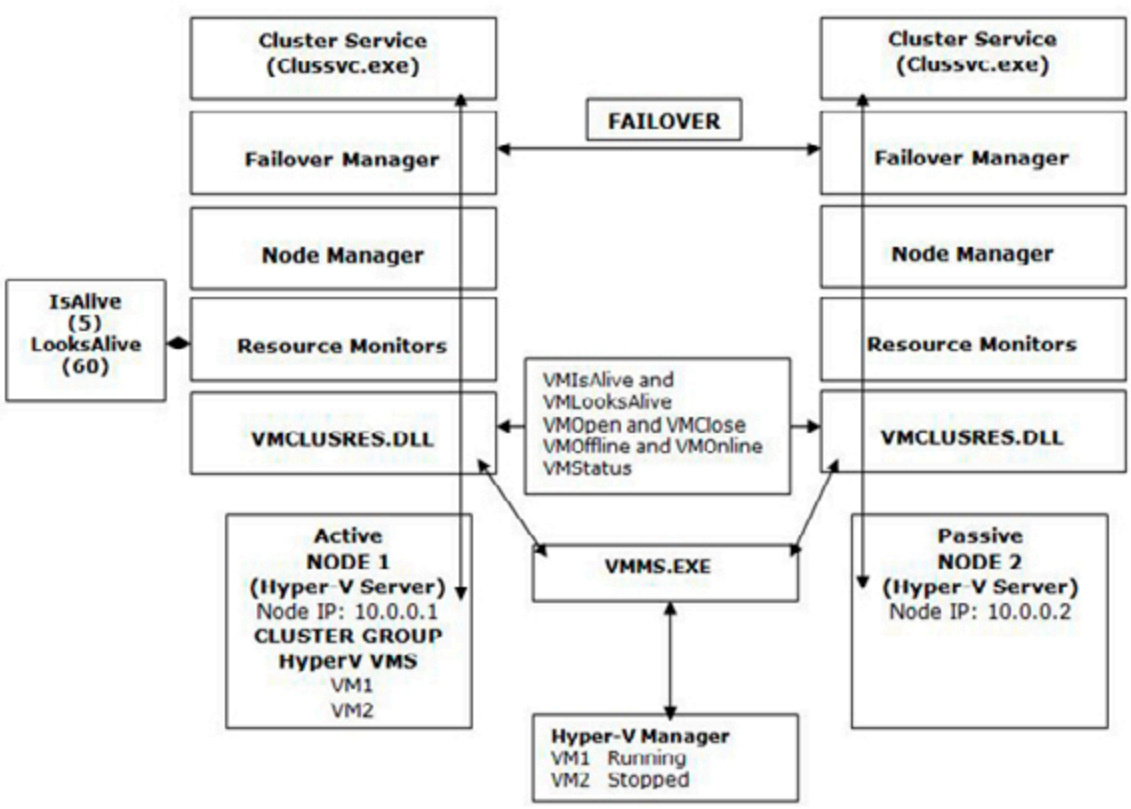

FIGURE 1.4 – VMCLUSRES.DLL and Status Messages in Hyper-V Virtual Machines Failover Process.

- After **IsAlive** interval expires (by default every 5 seconds), Cluster Service asks the Resource Monitor to report the status of Virtual Machines.
- Resource Monitor checks the status of Virtual Machine Resources in Cluster configuration database (HKLM\Cluster). It provides VMCLUSRES.DLL with the Virtual Machine Resources GUID and their current status (PersistenState.
- VMCLUSRES.DLL executes its own function (VM **IsAlive**) after it receives a signal from the Resource Monitor to perform a check on the Virtual Machines. It checks and reports back the status to Resource Monitor. VMCLUSRES.DLL will report the following status messages:

```
Online/Offline

Online/Offline Pending

Failed/Stopped
```

- After the Resource Monitor receives the status, it compares the status messages received from VMCLUSRES.DLL with the one stored in the Cluster configuration database. It then takes the action as per the status reported by the VMCLUSRES.DLL as listed below:

If comparison is successful, no action is taken. For example, status message received in step 2 is "Online" and VM **IsAlive** query also reports the same status.

If comparison is unsuccessful, the following actions are taken:

- If status message received in step 2 is "Online" and VM **IsAlive** query reports "Offline," the Resource Monitor executes an "Online" function. VMCLUSRES.DLL receives this message and executes VM Online function to bring the Virtual Machine online. This status message is also reported to the VMMS.EXE process.

  Tip

  The Resource Monitor doesn't take any action for Online/Offline status messages because an Administrator might have stopped the resource for maintenance purposes, but the same should also be reflected in the Cluster configuration database before **IsAlive** *is called. The Resource Monitor only takes action when the comparison is not successful as stated above.*

  Furthermore, there shouldn't be any inconsistencies in the Cluster configuration database. If there were any, these wouldn't last longer than 5 seconds since **IsAlive** *calls always update the status at the Cluster configuration database.*

- The mechanism isn't really straight forward. There could be one more message returned by VMCLUSRES.DLL that is "Failed." In this case the Resource Monitor sends a message (Restart) back to VMCLUSRES.DLL to restart the Virtual Machine resource in the cluster. VMCLUSRES.DLL in turn executes the "VM Online" function to bring the failed Virtual Machines online.

  Tip

  VMCLUSRES.DLL doesn't actually implement a separate Restart function. Instead it always uses its own implemented VM Online function. If a resource doesn't come online within the specified interval or after a few attempts, the resource is considered to be failed and then the Failover process starts. The same is notified to the VMMS.EXE as it needs to keep the status of all the Virtual Machines running in the Cluster.

- After the Virtual Machine resource has failed, the message is passed back to the Resource Monitor. The Cluster Service receives this message from the Resource Monitor and starts the failover process with the help of the Failover Manager. The Failover Manager on each node will communicate with the Failover Manager on another selected cluster node to improve the failover process. Before Failover Manager on the node where the Virtual Machine resource has failed communicates with another Failover Manager, it needs to get the list of nodes available in Cluster. This is where the Node Manager comes into picture. It supplies the list of nodes available in the cluster and the first available node at the top of the list will be selected for failover.

- Once the list of nodes has been obtained by the source Failover Manager, it will talk to Failover Manager on the target node. The Failover Manager on the target node supplies the list of Virtual Machines Resources along with GUID and **PersistentState** to Resource Monitor. Since this is a failover process, the Resource Monitor knows what to do next. It lists all the Virtual Machines with its flag (Online or Offline) and instructs the Resource DLL of Hyper-V to execute the VM Online function from its library.

- The Resource DLL, in turn, executes the VM Online function to bring the resources online on the target node. The same is updated to the VMMS.EXE process of Hyper-V.

- If the Virtual Machine is started successfully within a few attempts, the failover process doesn't occur.

Thus if there is no Resource DLL for Hyper-V Virtual Machines, the failover process could take a longer time to move the resources from one node to another surviving node. Because Hyper-V Resource DLL is competent enough to handle the cluster functions executed by the Clustering Software, it doesn't need to wait to decide which action to take. As stated above, the cluster-aware functions are mapped with Hyper-V Resource DLL-specific functions, so it is easier for Hyper-V Resource DLL to execute these functions as soon as they are executed from the Resource Monitor.

In figure 1.4 you see VMMS.EXE and Hyper-V Manager. Every function executed by the VMCLUSRES.DLL is also notified to VMMS.EXE. VMMS.EXE, in turn, refreshes the status of its VMs on the Hyper-V Server. This is required in order to know the exact status of a VM running on the Hyper-V Server. As an example, an Administrator could open the Hyper-V Manager to get the status of all the Virtual

Machines on the Hyper-V Server. If a Virtual Machine has failed and this is not communicated to VMMS.EXE, then there could be confusion, since the Failover Cluster Manager would report one status and the Hyper-V Manager would report a different status.

> ***Tip***
>
> *IsAlive is executed every 5 seconds for a Virtual Machine in the cluster. You could decrease this value to 1 or 2 to speed up the failover process.*

Conclusion

To summarize, Virtual Machines running on Virtual Server are not cluster-aware because they do not provide any Resource DLL. Virtual Machines running on Hyper-V are cluster-aware because they provide a Resource DLL as they ship along with a cluster Resource DLL.

We saw how the Cluster Service doesn't talk to VMCLUSRES.DLL directly. In fact, it uses its Resource Monitor. The status messages passed by the Hyper-V Resource DLL are received by the Resource Monitor to perform any appropriate action.

Finally we also saw how the Hyper-V Resource DLL plays an important role for its Virtual Machines in the cluster. Resource DLLs allow Hyper-V Virtual Machines to be fully cluster-aware VMs. The functions executed by the Resource Monitor on behalf of the Cluster Service are supported by the Hyper-V Resource DLL. This makes the failover process faster.

Microsoft Hyper-V Networking and Configuration - Part 1

11 December 2009

by Nirmal Sharma

> ***In the first of a series of articles on Hyper-V Networking, Nirmal explains Hyper-V networking and VLAN tagging, and shows how to set up a Virtual Network switch. Once this is done, the Hyper-V Virtual Network Switches can be used to connect unlimited no. of Virtual Machines.***

Most of the article talks about Hyper-V Networking. It doesn't really elaborate on basics of Networking. Instead the article focuses more on the Hyper-V Networking and VLAN Tagging with examples.

The first article in this series explains the following topics:

Virtual Networking Overview

- Hyper-V Virtual Network Switch Overview.
- Microsoft Hyper-V Virtual Network Switch Types.
- Microsoft Hyper-V Virtual Network Maximum configuration.
- What happens when you create a Virtual Network Switch?

Terms Used Throughout This Article:

Parent Partition
A Windows Server 2008 running Hyper-V Role is called the Parent Partition or Root Partition. The Operating System (Windows Server 2008) running on the Root is the "Management Operating System." Parent Partition is responsible to create Child Partition and also controls the communications between all the Virtual Machines.

Child Partition
A Virtual Machine running on Hyper-V Server is called the Child Partition. The Parent Partition creates the Child Partition.

Virtual Switch or Virtual Network Switch
A Virtual Switch is a software component of Virtualization Software. Virtual Machines are connected to Virtual Switch in order to allow communications between Virtual Machines. A Virtual Switch, just like Physical Switch, does more than communications.

VLAN
A VLAN is a Virtual LAN. A VLAN is a method of creating independent logical networks. VLAN is a broadcast domain created by the physical or virtual switches.

VLAN ID
A unique number called the VLAN ID identifies the each VLAN. Each VLAN is separated by assigning unique VLAN ID.

VLAN Trunk and **Access Modes**
There are two modes a particular port on a Virtual or Physical Switch can operate in; Access Mode and Trunk Mode. Access Mode is used for end devices that will not require access to multiple VLANs. Trunk Mode is used for passing multiple VLANs to other network devices that need to have communication to multiple VLANs on one occasion.

Integration Services Component
A Hyper-V component used to enhance the performance of Virtual Machines running on Hyper-V. The Integration Services component is similar to the VM Additions of Virtual PC but it's more than that in the functionality.

VMBUS
VMBUS is a logical inter-partition communication channel. VMBUS allows Virtual Machines to access hardware resources faster. VMBUS is available only when you install the Integration Service Components on the Virtual Machines running on Hyper-V.

Virtual Networking Overview

Virtualization has been in the market for a long time. Every vendor has to design the Virtualization software in such a way that it dictates the physical environment. This includes physical networking as well.

Virtual Networking operates at Layer 2. Layer 2 cannot perform IP Routing which is basically done by the Layer 3 switches. Microsoft Hyper-V and VMWare both implement Virtual Networking. Virtual Networking is designed to meet the requirement to move from the physical to the virtual environment. Microsoft Hyper-V Virtual Networking has been designed as part of its Virtualization Software; Hyper-V. Microsoft entered into the Virtualization market recently. The way that the underlying components of Microsoft Hyper-V Networking behaves are completely different from VMWare. This series of articles does not compare Virtual Networking within Hyper-V and VMWare: Instead, it focuses only on Microsoft Hyper-V Virtual Networking technology.

In the Virtual World, there are no fundamental limitations. It is the responsibility of the vendor to improve the way that Virtual networking works in the implementation of the Virtual component. Microsoft Hyper-V has been designed to improve Virtual Networking by introducing new Networking Technologies. I'll elaborate on this later on in this series of articles.

Hyper-V Virtual Network Switch Overview

Virtual Network Switch implementation in Hyper-V provides the following functionality:

Virtual Network Switch operates at Layer 2.

- Switch maintains a table called MAC Table. This MAC Table contains the MAC Addresses of Virtual Machines connected and the Virtual Machine names.
- Hyper-V Virtual Network Switch has a learning algorithm in which it learns the MAC address of a Virtual Machine. This MAC Address, once learned, is stored in the MAC Table of the switch.
- Unlimited Virtual Machines can be connected to a Virtual Network Switch.

Microsoft Hyper-V Virtual Network Switch Types

Microsoft Hyper-V implements three types of Virtual Switches or Networking Types as shown in figure 1.1.

Hyper-V Networking: Three types of Virtual Networks:

Type	Parent OS	VMs on Same HV	VMs on Remote HV	LAN	Remark
External	x	x	x	x	Conn. Lost Temporarily
Internal	x	x			
Private		x			
Dedicated		x		x	

FIGURE 1.1 – Hyper-V Virtual Network Types.

As you can see in Figure 1.1, there are four networking types or Virtual Switches in Hyper-V. There is one more type called Dedicated. This is either not visible or not available in Hyper-V Virtual Network Manager. You see only three types. We will discuss the Dedicated type later in this article series.

Before I get into more details, let me explain you the default configuration of Hyper-V Networking. In a default Hyper-V implementation:

There are no virtual Network Switches created.

- Virtual Machines created on Hyper-V Networking are not associated with any of the Virtual Network Switch types shown in Figure 1.1.
- There is no Network Adaptors configured in Virtual Machines (depending on the guest Operating Systems).

The default Hyper-V Implementation looks like as shown below in Figure 1.2:

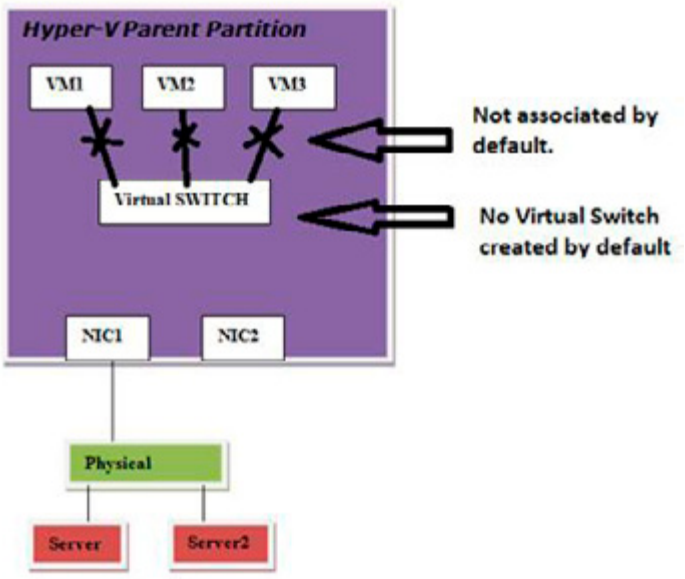

FIGURE 1.2 – Hyper-V Default Implementation and Network Configuration.

As you can see in Figure 1.2, there are three VMs created on Hyper-V Server; VM1, VM2 and VM3. By default, these VMs are not associated with the Virtual Network Switch and cannot have any communication with each other. If any of the VMs running on Hyper-V (VM1, VM2

and VM3) need to have communication with Server1 on physical LAN then they can't because Virtual Network Switch is not created by default.

Coming back to Virtual Network Switch types; you create Virtual Switches using the Virtual Network Manager found in the Action Pane on the right-hand side of the Hyper-V Manager.

As shown in figure 1.1, the "**External**" Virtual Network Switch allows you to have communication with Parent Partition of Hyper-V, Virtual Machines running on same Hyper-V Server, Virtual Machines running on Remote Hyper-V Server, and physical servers on the LAN. The External Network Switch requires that you have at least one Physical NIC (e.g. not associated with any other Virtual Network Switch). You can have one External Virtual Network Switch per Physical NIC.

One thing you notice is the remark column for "External" Virtual Network Switch. It says "Conn. Lost Temporarily," that means, connection is lost temporarily if you create an External Virtual Network Switch. Why so? The External Virtual Network Switch is mapped to a physical NIC on the Hyper-V Server. This is basically a binding of Virtual Network Services to a Physical NIC.

The "**Internal**" Virtual Network Switch allows you to have communication between Parent Partition of Hyper-V and the Virtual Machines running on the same Hyper-V Server. You cannot have communication with any other VMs which are associated with a different Virtual Network Switch or physical servers. Internal Virtual Network Switch does not require the availability of Physical NIC as communication happens internally or on the same Hyper-V Server. You can create Internal Virtual Network Switch without the Physical NIC also.

The "**Private**" Virtual Network Switch allows you to have communication between only the Virtual Machines running on the same Hyper-V Server. Communications are allowed only between the Virtual Machines which are connected to that Internal Virtual Network Switch.

Hyper-V Virtual Networking Maximum Configuration

Support For	MAXIMUM	Remark
Virtual NICs Per Virtual Machine	12 NICs	4 Legacy and 8 VMBus NICs
VLAN	Unlimited	
Virtual Machines Per VLAN	Unlimited	
External Network Virtual Switch Per Hyper-V Server	1 Per Physical NIC	
Internal Network Virtual Switch Per Hyper-V Server	Unlimited	
Private Network Virtual Switch Per Hyper-V Server	Unlimited	
Virtual Machines Per Virtual Network Switch	Unlimited	
Wireless		No Support for wireless
VLAN ID Tagging	External, Internal	
VLAN ID Tagging On Virtual Machines	One Per Virtual Machine	

FIGURE 1.3 – Virtual Network Types and Maximum Configuration.

As you can see in figure 1.3, Hyper-V Virtual Machine supports two types of Networking Cards; Legacy and VMBus NICs. The support for Legacy Network Card is included for Guest Operating Systems which are not supported by Hyper-V. There can be 4 Legacy Network Adaptors. You install a Legacy Network Adapter from the property of Virtual Machine and then selecting "Add New Hardware." Legacy Network Adapters use Device Emulation architecture to have communication with Parent Partition and to access Network resources. Please check the MSDN Article that lists all the Guest Operating Systems which are supported on Hyper-V: Guest operating systems that are supported on a Hyper-V virtual machine.

VMBus Network Card Adapters are available only when you install the Integration Services Component. The Integration Services component of Hyper-V leverages the VMBus architecture for best networking performance. There can be a maximum of 8 VMBus Network Cards.

A Virtual Machine running on Hyper-V can support a maximum of 12 Network Cards (4 Legacy and 8 VMBUS NICs).

VLAN Support is also included in Virtual Machines. A Virtual Machine running on Hyper-V can be configured with a VLAN ID. You can create unlimited number of VLANs in Hyper-V. You can have unlimited number of Virtual Machines per VLAN.

As shown in Figure 1.3, you can have one External Virtual Network Switch per Physical NIC on Hyper-V Server. This type of Network Virtual Network Switch allows Virtual Machines to have communication with LAN Servers also. The External Virtual Network Switch is mapped to Physical NIC in order to allow communication with Physical Devices.

There is no limitation for Internal and Private Network Virtual Switch. The reason is that Internal and Private operate internally and the communication is restricted to the Virtual Machines running on the Hyper-V. So you can create unlimited number of Private and Internal Virtual Network Switches.

Unlimited Virtual Machines can be connected to a Virtual Network Switch. This can be External, Private or Internal Network Virtual Switch.

In figure 1.3, you'll also notice the support for VLAN IDs on Virtual Switches. VLAN IDs can be specified on each Virtual Network Virtual Switch except Private Network Virtual Switch. These VLAN IDs can be used to create VLANs.

Similarly, VLAN ID support is also available for Virtual Machines. You can have one VLAN ID per Virtual Machine.

Note

There is no support for Direct Wireless for Virtual Machines running on Hyper-V. Instead you need to create a bridge between Wireless NIC and a Virtual Machine NIC.

What happens when you create a Virtual Network Switch?

You create Virtual Network Switch by selecting "Virtual Network Manager" located on the right pane of the Hyper-V Manager and then select the Network Virtual Switch Type you want to create. You need to select "Virtual Network Manager" to create any Virtual Network Switch Type. However, creation of External Virtual Network Switch is different from Private and Internal Network Switches.

When you create External Virtual Network Switch, you need to select a Physical NIC to which Virtual Switch will be mapped. This process involves modification of Physical NIC components.

Before creating External Network Virtual Switch, the network connection folder has only one network connection. The property of the Physical Network Connection looks like as shown below: (I assume you have only one Physical NIC attached to the Hyper-V Server).

FIGURE 1.4 – Property of Physical NIC.

When you create the External Network Virtual Switch using Virtual Network Manager, Hyper-V Components or VMMS.EXE process makes the following changes:

- Unbind the following services, protocols, clients from the Physical NIC:
  - Client for Microsoft Networks
  - File and Print Sharing for Microsoft Networks
  - TCP/IP Protocol IPv4
  - TCP/IP Protocol IPv6
  - Any other Service, client or protocol
- Bind the "**Microsoft Virtual Network Switch Protocol**"
- Create a new network connection in the Network Connections folder with the name you had specified when creating the External Virtual Network Switch. Let's say the name while creating the Virtual Switch you gave is "EXT NET Switch."
- Bind the following Services, protocols, clients to the External Virtual Network Switch (EXT NET Switch):
  - Client for Microsoft Networks
  - File and Print Sharing for Microsoft Networks
  - TCP/IP Protocol IPv4
  - TCP/IP Protocol IPv6
- Unbind the following protocol from the External Virtual Network Switch; EXT NET Switch:
  - "Microsoft Virtual Network Switch Protocol"

When you open the Network Connections folder you will see two Network Connections created; one Local Area Connection for physical NIC and other one is "EXT NET Switch" for External Virtual Network Switch as shown in figure 1.5:

FIGURE 1.5 – Network Connections with External Virtual Network Switch.

The bindings will also be changed. When you look at the property of Physical NIC and virtual Network Switch, the resulting image will look like as shown in Figure 1.6.

Property of EXT NET Virtual Switch.

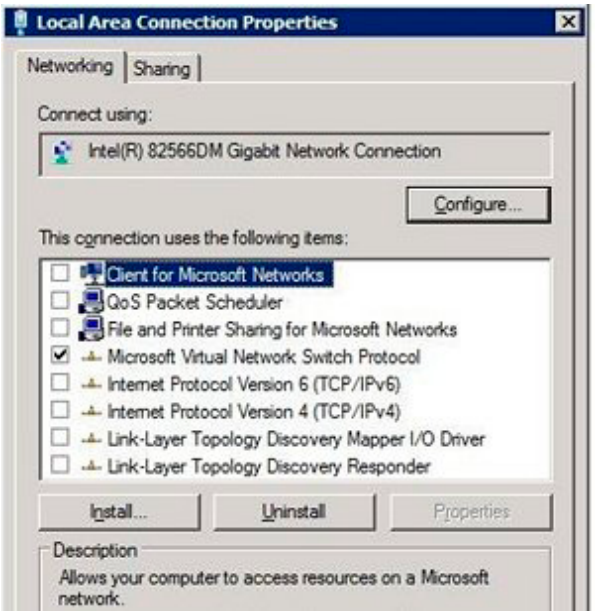

Property of Local Area Connection.

FIGURE 1.6 – Property of Virtual Network Switch and Physical Connection.

In the third step, the Virtual Network Manager unloads and loads the Network Driver for the Physical NIC from the memory. That's the reason the network connection is lost temporarily when you create External Virtual Network Switch. Any connected device to Hyper-V Server has to retry to reconnect. As an example; if you had connected to a Virtual Machine from a remote computer using VMConnect.exe, the connection is lost and you need to reconnect using VMConnect.exe.

A warning is shown to the user about temporary loss of network connection when you create External Virtual Network Switch.

The creation of Private and Internal Network Virtual Switches is same as shown above. However, you don't have to have the physical NIC to create these switches.

Conclusion

In this article we saw how Hyper-V Networking is different and the networking types introduced in the RTM version of Hyper-V. There are three networking types available; External, Internal and Private. We also saw the default configuration of Hyper-V Networking in which no network communication is possible unless a Virtual Network Switch is created. The Hyper-V Virtual Network Switches can be used to connect unlimited no. of Virtual Machines.

In the next series of this article, we will primarily focus on the following topics:

Hyper-V Networking and Packet Flow.

- Hyper-V Networking Examples (including VLAN Configuration).
- Hyper-V Networking using SCVMM.
- Configuring Hyper-V Networking using SCVMM.
- Three Ways to configure VLAN Trunking for Hyper-V Virtual Machines.

Unified Messaging

An Introduction to Unified Messaging

26 March 2009

by Brien Posey

For many years, everyone involved in the administration of offices has wished that there was a way of combining the voicemail, fax, the PBX with the Exchange Email system, particularly for highly mobile users. . With Exchange Server 2007, this has now become a reality. Even though it forces the Exchange Administrator to become familiar with telephony, the benefits are such that it is well worth the effort.

One of the least understood Exchange 2007 features has got to be Unified Messaging. Unified Messaging isn't really all that difficult to install, but it does require some specialized hardware and a degree of telephony knowledge, and that seems to scare some administrators away. I'll start out by talking about what Unified Messaging is, and about some of its features. I will also talk about some common misconceptions associated with Unified Messaging. Later on I will talk about the benefits that Unified Messaging can bring to the enterprise, and about where I think that Unified Messaging is headed in the future.

What is Unified Messaging?

Unified Messaging is an Exchange Server role that was introduced in Exchange Server 2007. The basic philosophy behind Unified Messaging is that users communicate in a variety of different ways. Some users prefer to send E-mail messages, while others prefer using the telephone. In fact, I tend to think that this is a big part of the reason why smart phones are so popular. They allow people to receive E-mail messages and telephone calls on the same device.

In a corporate environment though, a user typically has two separate mailboxes; one for E-mail messages, and another one for voice mail messages. Furthermore, voice mail has traditionally been tied to the telephone. Although it is common for voice mail to be remotely accessible, users often find themselves scribbling down names, numbers, or messages on pieces of paper, which often get lost.

Microsoft designed Exchange 2007 so that the Inbox allows users to store E-mail messages, voice mail messages, and faxes all in the same place. This frees the user from having to look for messages in multiple locations. It also gives users a way to make voice messages searchable, as I will show you later on.

If you look at Figure A, you can see an example of what a user's Inbox looks like after Unified Messaging has been enabled. Notice in the figure that the message at the top of the list contains a telephone icon, which reflects the fact that it is a voice mail message. In this case, the From field displays the phone number that the call was from, and the subject line also tells us that this is a voice message, and who it is from.

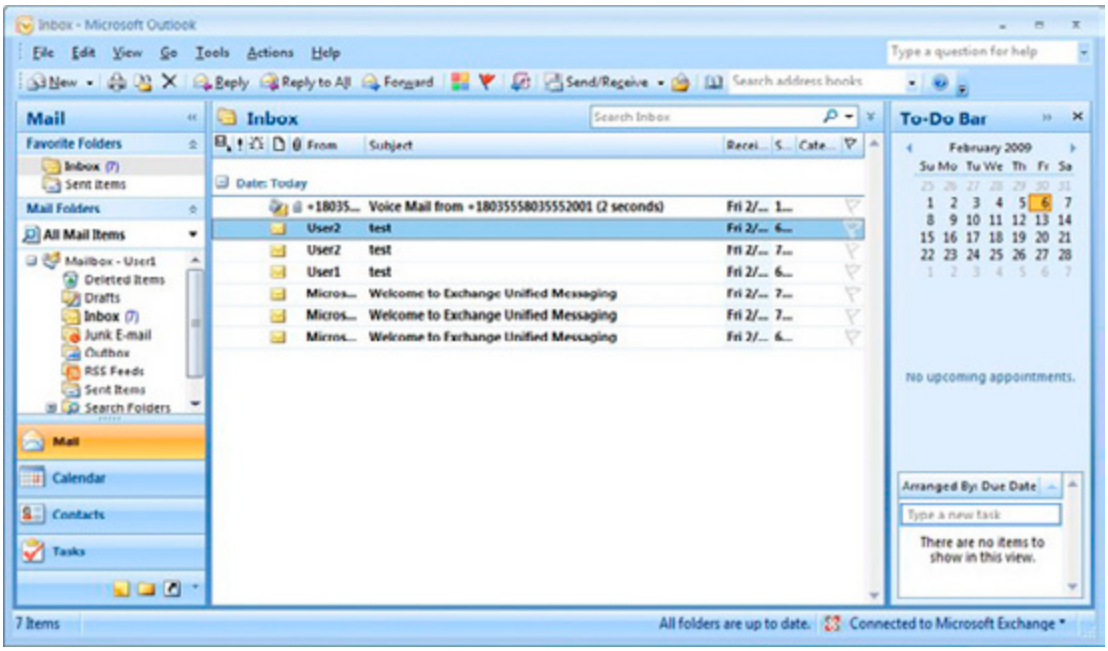

Figure A: Voice messages are stored in the user's Inbox.

If you look at Figure B, you can see what it looks like when I open this voice message. As you can see in the figure, Outlook displays an audio plug-in just beneath the message's Subject line. This allows you to play the message directly through Outlook.

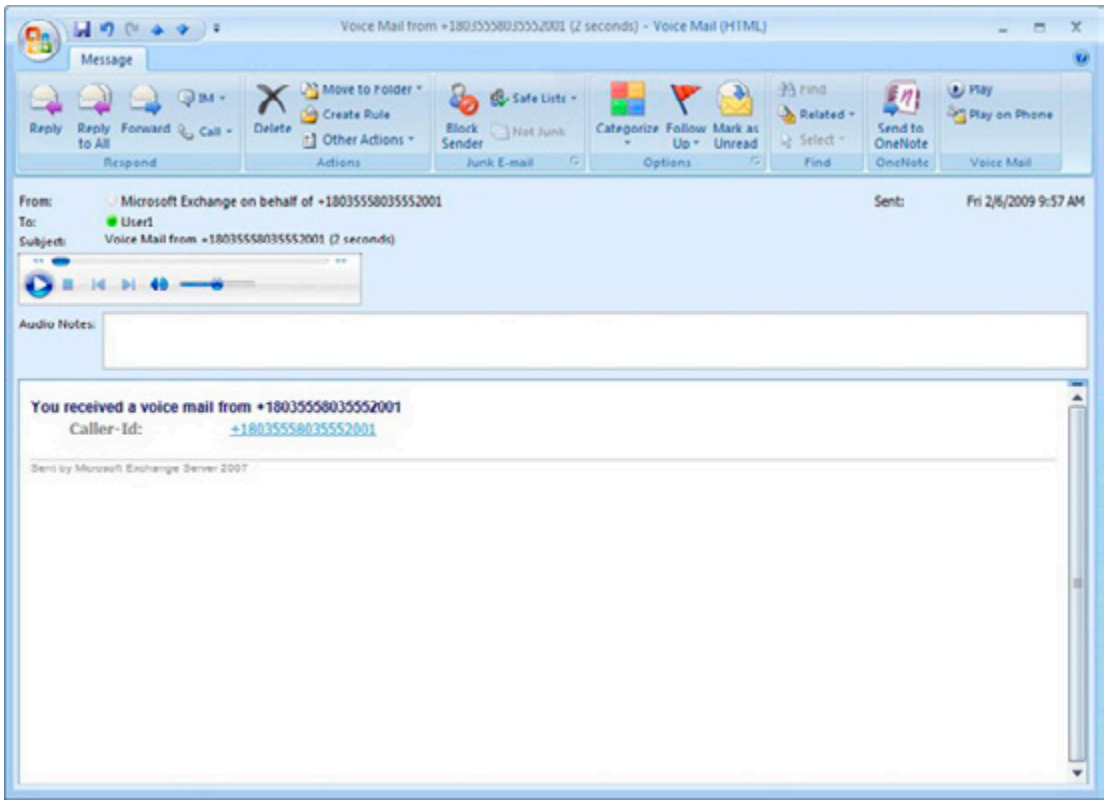

Figure B: This is what a voice message looks like.

Remember earlier when I said that users often take notes about voice messages on paper, and a lot of times those notes end up getting lost? Well, if you look just beneath the audio plug-in, there is an Audio Notes field. Users can take notes about the message in this field. These

notes are saved when the user closes the message. That way, when a user performs a search on their Inbox, Outlook searches both E-mail messages, and the notes that have been attached to voice mail messages.

Notice in the figure that this message contains a Reply button, just like an E-mail message does. If you click the Reply button, you can respond to the voice message with an E-mail message. Of course this only works if the message was left by a user who has Unified Messaging enabled, because Outlook must look up the phone number in the Active Directory, and figure out which user the number belongs to so that it can send a message to them.

Another thing that I want to point out is the Call icon, located within the Respond section in Figure B. Clicking this icon allows you to return the call rather than requiring you to respond to the voice message through E-mail.

One last thing that I want to show you about this message is the Play on Phone icon. Normally, when you play a voice message, the voice message is played directly through Outlook, on your computer's speakers. The problem with this is that in a crowded office, privacy can sometimes be a concern. If you don't have a set of headphones to plug into your computer, you can use the Play on Phone icon to redirect the message to a telephone. That way, you can listen to the message without the entire office over hearing it.

When you click this icon, Outlook asks you to enter the phone number of the phone that you want to play the message on. Outlook then dials the number, and allows the user to listen to the message on the telephone.

Although Unified Messaging is great for storing voice messages and faxes alongside E-mail messages, this in only one of the two major Unified Messaging features. The other major feature is Outlook Voice Access (OVA). Just as Outlook Web Access (OWA) allows users to access their mailboxes through a Web browser, OVA allows users to access their mailboxes over the telephone.

When you initially configure Unified messaging, you're given the opportunity to associate a phone number with a dial plan, as shown in Figure C. This phone number is known as a Subscriber Access Number. It is the number that users can call to access Outlook Voice Access.

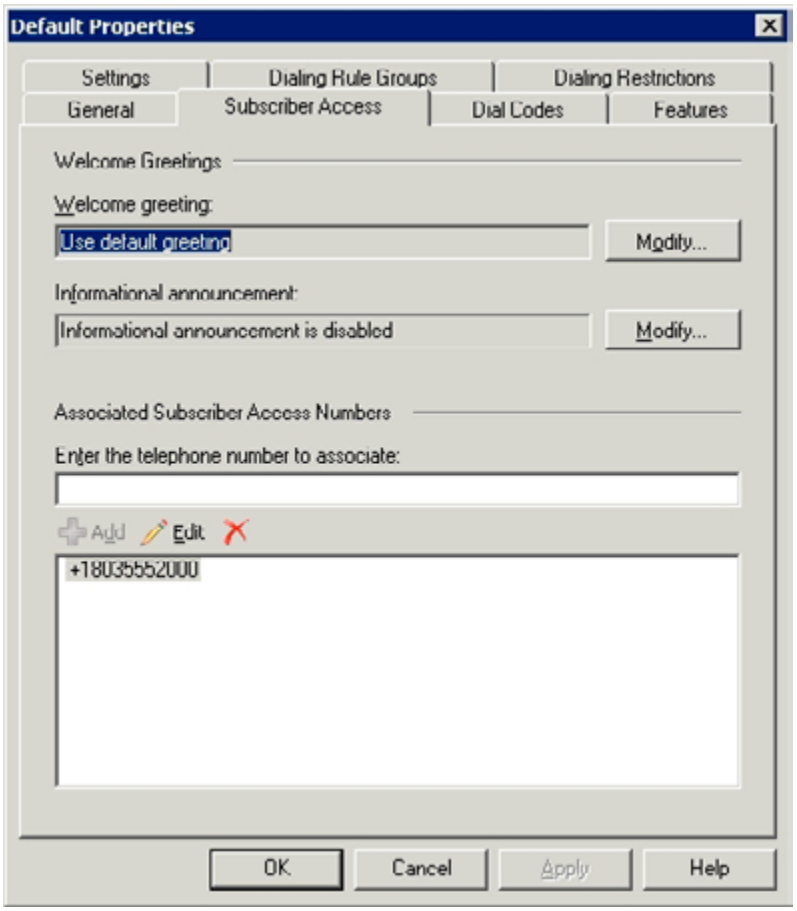

Figure C: A subscriber Access Number provides access to OVA.

After Unified Messaging is set up, you must enable the various mailboxes for unified messaging. When you do, you are required to either provide a PIN for the mailbox, or to tell Exchange Server to automatically generate a PIN. This PIN is used when the user logs in to OVA.

After a mailbox has been enabled for Unified Messaging, Exchange Server sends a Welcome to Exchange Unified Messaging message to the mailbox. This message confirms the user's extension number (which they should already know), and provides them with their PIN. Once the user has this information, they can use Outlook Voice Access to interact with their mailbox over the telephone. You can see an example of the Welcome to Exchange Unified Messaging message in the figure below.

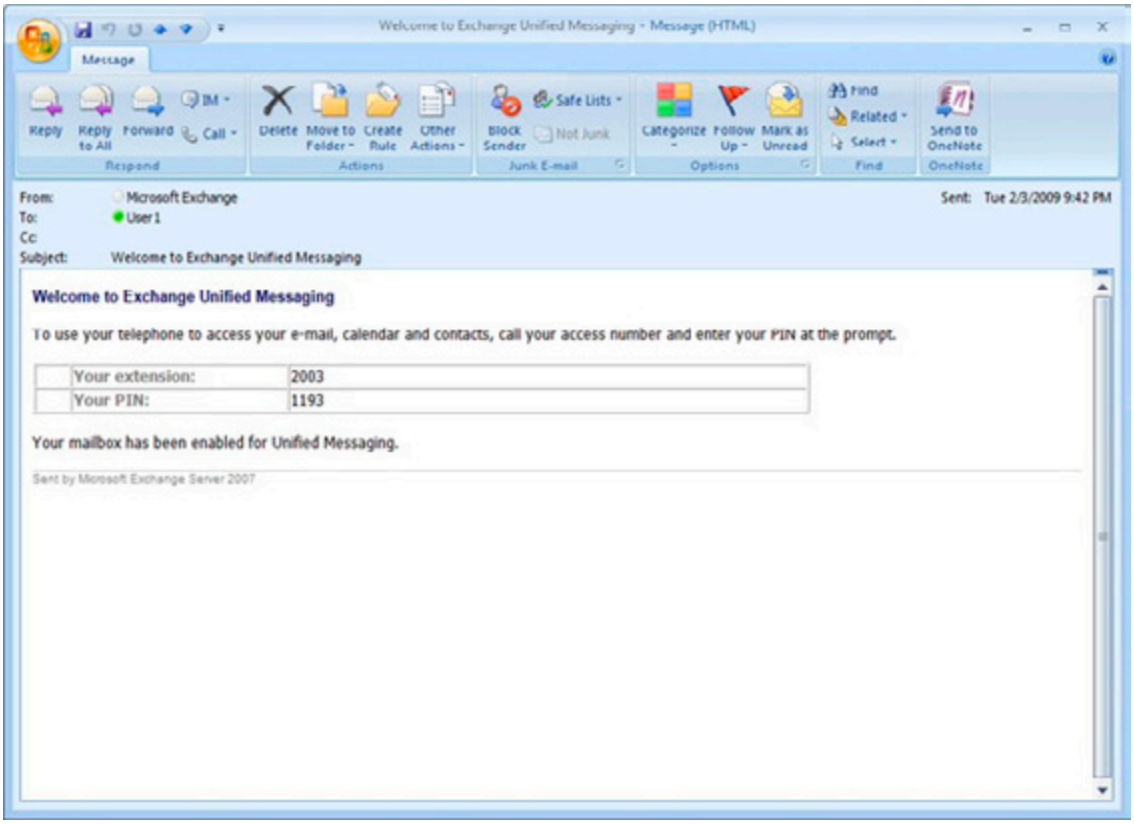

Figure D: Users are automatically sent a message containing their extension number and PIN.

If a user wants to use OVA, they simply call the subscriber Access number. When they do, Exchange Server will prompt them to enter their extension followed by their PIN. The user enters this information by using the touch tones on their telephone. This authentication method is used in place of the traditional username and password.

Once the user has logged in to OVA, they are free to listen to their e-mail messages or to verbally compose new e-mail messages. They also have the ability to interact with their calendar. Users can schedule or reschedule meetings, mark a timeslot as busy, or do just about anything else that they could do to their calendar through Outlook.

The nice thing about OVA is that users can interact with it verbally, or they can use the touchtone keypad on their phone. It has been my experience that the voice recognition software that OVA uses is quite good. In fact, I have even been able to use OVA to verbally check my e-mail in my car while I am driving.

Misconceptions about Unified Messaging

Probably the biggest misconception about Unified Messaging is that it is a full-blown unified communications solution. In actuality though, nothing could be further from the truth. It important to remember that Unified Messaging is only one part of Microsoft's unified communications strategy. Microsoft offers a completely separate product called Office Communications Server 2007 (OCS) that is designed to be a full-featured unified communications product.

OCS 2007 is the successor to Microsoft's Live Communications Server product line. It allows users to place and receive voice and video phone calls using Microsoft Office Communicator. Furthermore, OCS 2007 can be configured as a mediation server, which can act as a bridge to the organization's IP enabled PBX system. This allows external phone calls to be routed to the desktop. Of course users also have the ability to place external calls as well.

The one thing that OCS 2007 really seems to be missing is a voicemail feature. As you might have already figured out though, OCS 2007 and Exchange Server 2007 can be configured to work together so that voicemail capabilities and OVA can be integrated into the organization's unified communications solution.

One thing to keep in mind though is that the process of making OCS 2007 and Exchange Server 2007 work together is not completely seamless. Exchange 2007 was completed before OCS 2007. Therefore, Microsoft does not offer full OCS 2007 support in Exchange 2007 unless you install SP1. Even then, the setup process requires you to run an obscure script that tells Exchange Server 2007 to treat the OCS Enterprise Pool as a VoIP gateway. The script also sets the necessary permissions for the two products to work together.

The Benefits of Unified Messaging

Unified Messaging is the most resource intensive of all of the Exchange 2007 server roles, and using it requires additional hardware that bridges the gap between the corporate network and the PBX system. Often times, the resource intensive nature of the Unified Messaging role means that companies that want to use Unified Messaging will end up having to deploy multiple Unified Messaging servers. Even if it does not initially seem that Unified Messaging is going to be all that resource intensive, it is important to remember that voice mail is "real time" in nature. Something as simple as a dropped packet can introduce jitter and cause part of a message to be lost.

Organizations may also find themselves having to beef up their mailbox servers because of the additional load to voice messages place on the storage subsystem. Obviously, implementing Unified Messaging is going to be an expensive endeavor in most cases because of all of the additional hardware requirements. It is therefore essential to consider how Unified Messaging is going to be a benefit to the organization.

In my opinion, the biggest benefit that Unified Messaging provides is Outlook Voice Access. In this day and age, users often need access to their mailboxes at a moment's notice. Occasionally, users may find themselves in a situation in which they do not have access to an e-mail client, or in which using a traditional e-mail client is impractical. Telephones however, are universal. This makes it possible for users to check, and respond to messages from almost anywhere in the world.

I also think that having voice messages and e-mail messages stored in the same mailbox provides a great benefit to the organization. Using consolidated storage may not initially seem like a big deal. However, it may allow remote users to work more efficiently. It also allows them to instantly be alerted to new voice messages. Traditionally, users would have to dial into a voicemail system in order to find out if they have any new messages waiting or not, although some systems support paging, message forwarding, and other alerting features.

I also like the idea that a consolidated mailbox allows users to have one place to manage all of their communications with clients. Since voice notes are searchable, it becomes easy for a user to enter a client's name into the Outlook Search box and see every conversation that they have ever had with the client, regardless of whether it was over voicemail or e-mail.

One last benefit that I want to talk about is faxing. Fax machines are quickly becoming obsolete, but they do still have their place. This is especially true in situations in which a signature is required on a document. Unified Messaging not only give each user fax capabilities, but it allows faxes to be routed directly to the recipient rather than ending up in a communal fax server in which privacy may be an issue.

The Future of Unified Messaging

In the future, I expect to see a much closer relationship between Exchange Server and Office Communications Server (OCS). This is evidenced in the way that Microsoft has begun aligning the version numbers of these products.

A few years ago, Microsoft decided to renumber their internal Exchange Server release numbers to match the version numbers that were being used by the Microsoft Office team. Both the Exchange team and the Office team used 12 as the version number for their products that were released in 2007. What is interesting is that although Exchange 2007 and OCS don't work together quite as seamlessly as Exchange and Office do, the OCS team also adopted 12 as their most recent version number.

The OCS team has recently released a new version of OCS which is called Microsoft Office Communications Server 2007 R2. The OCS team used 13 as the version number for this product, so the Exchange team and the Microsoft Office team chose 14 as the version number for their next release. This shows that the Exchange team is committed to making Exchange and OCS a "better together" experience.

At the present time, Microsoft hasn't publicly released any information regarding how Unified Messaging will evolve in the next Exchange Server release, and in future releases beyond Exchange 14. Even so, there are some clues as to what might become of Unified Messaging, and Microsoft's Unified Communications infrastructure as a whole.

Microsoft has firmly committed itself to its Software + Services initiative. The basic idea behind this initiative is that Microsoft plans to use the power of locally installed software in conjunction with distributed Web services to create a much more collaborative computing experience in which data and information services are available anywhere, regardless of a user's location and the type of device that they are using.

Although Microsoft's complete Software + Service vision is still several years away from being fulfilled, Exchange 14 is slated to be the first version of Exchange that can be operated as either software or as a service. Exchange 14 is designed to be installed on local servers, but Microsoft has also made it clear that they intend to offer hosted Exchange services that are powered by Exchange 14 as well.

Microsoft's Software + Services initiative has already started showing up in some of the company's other product lines as well. If you want to see what the future of Unified Messaging and Unified Communications may be like, then I think that the product that is the most worth paying attention to right now is Xbox 360.

The reason why I say this is that Xbox 360 fits into Microsoft's Software + Service initiative in that it can act as a standalone unit for playing games and movies, but it also offers collaborative services through a service known as Xbox Live.

Xbox Live offers a glimpse into the future of Unified Communications in that it allows gamers to talk to each other through a VOIP connection while they are playing games, or outside of a game. A feature called Xbox Live Vision also offers video conferencing capabilities.

OCS offers VOIP calling capabilities today through a feature called Enterprise Voice, and it also offers video conferencing capabilities. Although these features are designed to be easy to use, they tend to be difficult for an administrator to initially setup. In fact, before an administrator can deliver all of the capabilities offered by Unified Messaging and Unified communications, they need to have expertise in Exchange Server, OCS, traditional networking, telephony, and Windows. The deployment is so complicated in fact, that I am in the process of writing an entire book on the subject. In contrast, Xbox Live users can access these same basic capabilities today, but with almost none of the headaches that are associated with the initial setup process.

I look for Microsoft to turn Unified Messaging and Unified Communications into extensible services in the Exchange 15 time frame. My guess is that by doing so, they will make unified communications and unified messaging much easier to implement.

At the same time though, I don't expect these products to totally become hosted services. If Microsoft redesigned all of their server products so that they were only available as hosted services, or so that using the products as hosted services was the most attractive licensing model, there would almost certainly be a tremendous backlash from the IT community.

Imagine for instance what would happen if Exchange were only available as a hosted service. Administering Exchange would largely be confines to tasks such as setting up mailboxes or resetting passwords. If that happened, most of today's Exchange administrators would be out of work. The reason why I say that there would be a backlash is because most administrators aren't going to recommend that

a company implement a solution that they know is going to put them out of a job. Instead, they would probably look for a competing product that offered the promise of job security. I think that Microsoft probably realizes this and although they are committed to delivering on their software + services initiative, they will still make future versions of Exchange and OCS flexible enough to meet varying demands of their customers and complicated enough to appease Exchange admins.

Conclusion

Unfortunately, I don't have a crystal ball so I can't tell you what the future holds for Unified Messaging or for unified communications in general. Even so, I can tell you that both Unified Messaging and OCS seem to be gaining popularity, so it is highly unlikely that they are going to go away in the future.

Organizations have been using PBX systems for many years as a way of giving them the freedom to configure their internal phone systems without having to involve the telephone company each time that they want to make a change. I think that as future versions of these products are released, we will eventually start seeing unified communications products such as Exchange and OCS start to replace traditional PBX systems in the enterprise.

Moving to Office Communications Server 2007 R2

05 May 2009

by Desmond Lee

Office Communications Server 2007 R2 has some exciting new features, which make its deployment well worth considering. Desmond focuses on using the side-by-side method to migrate existing OCS RTM servers and clients over to R2, which you can likewise adopt to create an entirely new R2 installation.

Like many organizations, you may already be operating a Unified Communications infrastructure based on Live Communications Server 2005 SP1 or Office Communications Server 2007. OCS 2007 R2 edition has come with some major improvements and exciting new productivity features that compel companies to look into deploying this new version to stay one step ahead. Whether you are beginning with bare soil or have an existing OCS setup, this article lays out the key considerations to help you move as early and as effortlessly as possible to the world of Unified Communications with OCS 2007 R2.

Migration Strategies

Those of you who are starting out from scratch may like to read my LAST ARTICLE to kick start your OCS deployment with R2. Although it is written from a virtualization perspective, the principles and instructions in that last article apply equally to a production environment. Nevertheless, look out for pointers in this article to help ease your implementation woes and challenges. At the time I write this, running OCS in a virtualized environment in production is not supported regardless of the underlying virtualization platform i.e. Microsoft Hyper-V, VMware Esx Server or the recently announced vSphere.

Update *(13 May 2009): Microsoft released a blog post on Office Communications Server 2007 R2 Virtualization. You can read and find out more about this topic* HERE.

If you already have a production setup based on either Live Communications Server 2005 SP1 (LCS) or OCS 2007 RTM (some refer it to R1), resist the urge to simply pop in the R2 DVD to run setup. There are significant architectural changes in the product with updated AD environment and back-end database requirements. They must be fulfilled for a successful rollout whether building the environment brand new or migrating to R2 in phases with coexistence of previously supported versions. Because there is no direct upgrade path to R2 from OCS 2003, Office Communications Server 2003 users will have to migrate to at least LCS 2005 SP1 first.

You can take one of two alternative migration strategies on the road towards implementing R2; these are "uninstall/reinstall" and "side-by-side."

"Side-By-Side"

To keep operations going with minimum disruptions and inconvenience to end-users, the side-by-side migration is typically the preferred approach. This is at the expense of having to invest in new hardware for the R2 infrastructure with varying degrees of complexity to maintain proper coexistence. In addition, new names for the servers, pools and IP addresses have to be assigned. Since OCS 2007 R2 at its core is exclusively 64-bit, it can only run on x64 versions of Windows Server 2003 SP2 or higher and Windows Server 2008.

"Uninstall/Reinstall"

The uninstall/reinstall method allows you to preserve system settings and repurpose compatible hardware wherever possible. However, you will have to be prepared for longer periods of downtime to deal with major tasks such as backup/restore of existing user, configuration data and server rebuilds.

To help you quickly get up to speed with migrating the core infrastructure to R2, I have attempted to distill the mass of information down to what I believe are the most important and essential steps. This can be taken as a guide to migrate using the uninstall/reinstall or side-by-side methods. The latter is the focus of this article where migration starts from the internal network before moving out to the perimeter network. This is commonly known as an "inside-out" migration strategy. During the migration process, the R2 deployment operates independently yet can communicate with existing LCS or OCS RTM deployments in the same AD framework.

Pre-Migration Tasks

Quick side-by-side Migration Checklist

- Raise Active Directory domain and forest functional level.
- Move default location of OCS global settings.
- Apply hotfixes to OCS RTM server roles and clients (co-existence).
- Prepare back-end database server and x64 server for R2 installation.
- Run preparation steps for AD Schema, forest and domain.
- Prepare DNS records and digital server certificates.
- Deploy R2 server roles and services and install R2 administrative tools.
- Move users to home on R2 Front End servers.
- Inside-out migration for remaining internal server roles out to DMZ.
- Standardize and rollout R2 client applications and supporting components.
- Decommission OCS RTM servers.

The very first step is to raise both the domain and forest functional level to at least the Windows Server 2003 mode required by R2 to function (see TechNet Library HERE). To do this, login to an active domain controller holding the Schema Master role with an appropriate administrative account in Active Directory.

This task should be performed at the earliest possible moment so that the changes have sufficient time to propagate to all domain controllers throughout the AD forest. The domain controllers themselves can be running Windows Server 2003 SP1 or higher and Windows Server 2008 (Standard or higher editions) with a minimum of one Global Catalog server available. The processor architecture does not matter and can either be x86 or x64.

By default, both LCS and the RTM versions of OCS store the global settings in the root domain's System Container within AD. LCS/OCS servers must be able to access these settings anywhere in the forest. In an environment already installed with LCS 2005 SP1 or OCS 2007 RTM, the ability to change the location for storing the global settings in the System Container to the Configuration Partition or vice versa is disabled (see Figure 1).

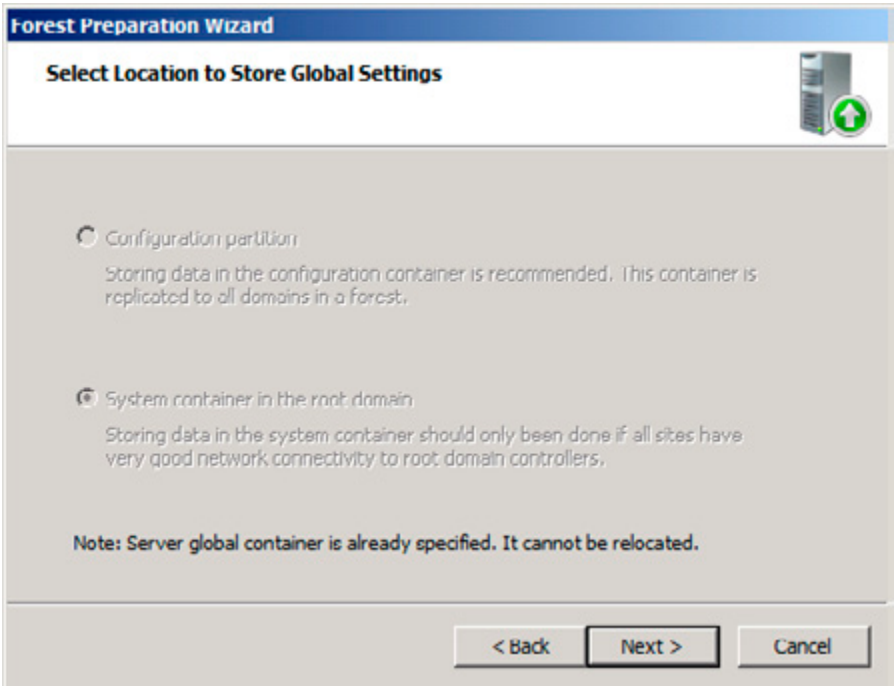

Figure 1.

You can download the Microsoft Office Communications Server 2007 Global Settings Migration Tool and follow the instructions in the TechNet library to relocate the OCS global settings. You must use the OCS RTM tools to move the global settings *before* launching the step to prepare the AD schema. Ascertain that you first stop all LCS and OCS services ahead of time and re-start them after the move. Otherwise, you will not be able to make this change once the schema has been upgraded to R2. The OCS global settings themselves are actually created during the Forest Prep step. Like raising an AD domain or forest functional level, this operation is a one-way street. Presently, no official tools are known to exist that can move this back or to enable relocation of the global settings after a R2 schema extension is complete. A clean installation of R2 will by default store its global settings in the Configuration Partition.

Why should I care if I have only a single domain, single forest AD structure? Basically, a number of AD sites are usually created to optimize network traffic in a decentralized AD topology involving multi-site deployments. If sub-domains are present or added at a later stage, OCS servers and UC-enabled users deployed in the child domains will continue to reach out to a domain controller and Global Catalog server in the root domain to retrieve data for the global settings. Failure to move the OCS global settings could cause long startup times for the core OCS services, synchronization and update errors, long replication delays, inter-domain instant messaging/presence issues or user login problems. This becomes even more critical if you are considering deploying Enterprise Voice (VoIP) for mobile users across disparate geographical locations with anywhere access via Microsoft Office Communicator client (MOC) or a supported browser in R2.

You therefore have to move the OCS global settings to the Configuration Partition from the pre-R2 default location. Also known as a Naming Context (NC), this is one of 3 stores in AD – Schema, Configuration and Domain (or System Container in OCS speak) – where only the first 2 NCs are replicated to every domain controller in the forest. In so doing, the global settings can be directly accessible on local DCs and GCs in the respective domains; all without the need to traverse over slow, costly and often unreliable WAN links. You can access the various global settings by navigating to the Forest - <FQDN> node / Properties then click on the Global Properties, Voice Properties or Conferencing Attendant Properties menu options using the R2 administrative snap-in.

Infrastructure Readiness

To ensure that R2 servers can operate problem free in an existing OCS RTM environment, you should download and apply the list of updates on the RTM server roles (Table 1). Check that you have updated documentation of your current environment (including user account inventory) and make a full system backup before you go ahead with any changes. Since it usually takes an organization much more time to test and deliver patches to its fleet of Windows client machines, I suggest that you also prepare the user base running the MOC client well ahead of the first batch of R2 users going live. Until all clients are migrated and homed on R2 server pools, different MOC versions will negotiate and continue to communicate using the most common feature sets available to them. Eventually, you should develop a plan to rollout the MOC R2 client to all your users after the entire end-to-end infrastructure has been upgraded to the R2 equivalents.

Hotfix/Update	All OCS Server Editions & roles	Mediation Server role	MOC 2007	Notes
KB956389 (Nov 2008)	X			Supersedes KB952783 (Aug 2008). Fixes issues described in KB958560 (Nov 2008) and KB958561 (Nov 2008)
KB956831 (Nov 2008)	X			Supersedes KB946763 (Jan 2008). Fixes issues described in KB957489 (Oct 2008), KB957595 (Oct 2008) and KB957648 (Oct 2008).
KB956829 (Oct 2008)		X		Supersedes KB952780 (Aug 2008). Fixes issues described in KB957490 (Oct 2008) and KB957648 (Oct 2008)
KB961552 (Mar 2009)			X	Supersedes KB957465 (Dec 2008) with fixes for a number of issues.
KB953582 (Oct 2008)	X	X	X	Fixes registry extensions permission errors with program installation in Vista and Windows Server 2008. Must apply before installing the R2 administrative tools.
KB953990 (Jan 2009)	X	X	X	Applies to machines with .NET Framework 2.0 SP1 on Windows (x86 and x64).

Table 1: OCS and MOC RTM Updates.

In the event that you are installing an R2 Enterprise pool, a dedicated Windows server (or active/passive cluster) running Microsoft SQL Server 2005 SP2 or SQL Server 2008 must already be setup and configured on another physical server. Both the operating system and database application can either be 32- or 64- bit editions. A new database instance will be automatically created when you install the R2 Enterprise pool. The reason behind this is due to the back-end database schema change in R2.

This same box can host multiple SQL databases and instances supporting the unique needs of the Enterprise pool, Archiving or Monitoring Server R2 server roles. Nonetheless, it is not a recommended practice to collocate the pool's back-end database and archiving database for performance reasons. Note also that collocating any R2 server roles on the same back-end SQL Server remains unsupported. For R2 Standard Edition setup, the SQL Server 2005 Express Edition SP2 database is automatically created on the local machine to store user and configuration data as part of the setup process.

Let the Show Begin

With the pre-migration tasks out of the way and the back-end infrastructure in place, you are almost ready to run the R2 setup executable. As stated before, you need to provision a machine with a supported x64 edition of Windows and join it to the domain so as to conduct a fresh R2 install.

Now if you are targeting an AD environment with domain controllers that run only x86 editions of Windows, you can still execute Step 1 to 7 under "Prepare Active Directory" on the main R2 setup screen from a member server with good IP connectivity to the Schema Master. Clearly, the member server must be operating on an x64 edition of Windows as the R2 installation Wizard relies on the 64-bit version of setupEE.exe or setupSE.exe. If not, you will have to explicitly extract and deploy the 32-bit version of LCSCmd.exe from the DVD media to prepare AD.

To avoid the hassle of switching between consoles, I recommend that you execute all setup steps for R2, including those for preparing AD, on a remote Windows Server 2008 server that will eventually be installed with one of the R2 server roles. This must be the full graphical version of Windows Server 2008 (not Server Core). The optional Remote Server Administration Tool (RSAT) feature is needed to accomplish this on Windows Server 2008. For AD environments with multiple domains, you must execute the Domain Prep step for each domain where R2 servers will be deployed. Even if the latter is not planned for a particular domain, this procedure must still be done if you intend to UC-enable any user account in that domain. This step to prepare an AD domain mirrors the requirements of Exchange Server 2007. Again, set aside adequate time to allow AD replication to run its course in the forest.

As soon as AD preparation steps for the schema, forest and domain are completed, you can commence with the actual deployment of an R2 Standard or Enterprise Edition server pool. Mixing RTM and R2 servers in the same pool is not supported; consequently you can rule out the rolling upgrade option where servers are upgraded one by one in the pool. After all, an in-place upgrade is not possible between disparate processor architectures i.e. OCS RTM x86 to R2 x64.

Even though the Enterprise expanded topology continues to be supported, Microsoft recommends targeting new R2 Front End server pools as well as Edge Server roles and services to be based on the consolidated topology. Such configurations simplify deployment and reduce operational overhead since all server roles are collocated on a single machine. More computers with analogous, collocated server roles can be added to the hardware load-balanced pool to scale out and scale up the infrastructure afterwards, thereby improving overall system availability.

> ***Application Server***
>
> *The new Application Server is an integrated component of an R2 front-end server comprising of the Conferencing Attendant, Conferencing Announcement Service, Response Group Service and Outside Voice Control. Each of these UC applications can be activated individually. The Application Server component always runs as a separate process within the front-end server and cannot be segregated for deployment on another machine in the pool.*

Other than making sure that the prerequisites for hard- and soft- ware are met, you should be ready to call up the passwords for the RTCService and RTCComponent service accounts that were created from the previous OCS installation when prompted during setup. For R2 Standard Edition or Enterprise consolidated topology setup on Windows Server 2008 to succeed, you need to separately install the Web Server (IIS) server role then select the IIS 6 Management Compatibility and security role services (Figure 2).

On top of that, you should arrange for the required number of digital server certificates configured with the correct Fully Qualified Domain Name (FQDN) of the distinct R2 server roles. You will need this at the "Configure Certificate" setup step. An OCS infrastructure relies heavily on digital certificates to enforce server and client authentication as well as securing all communications channels with encryption (signaling and media).

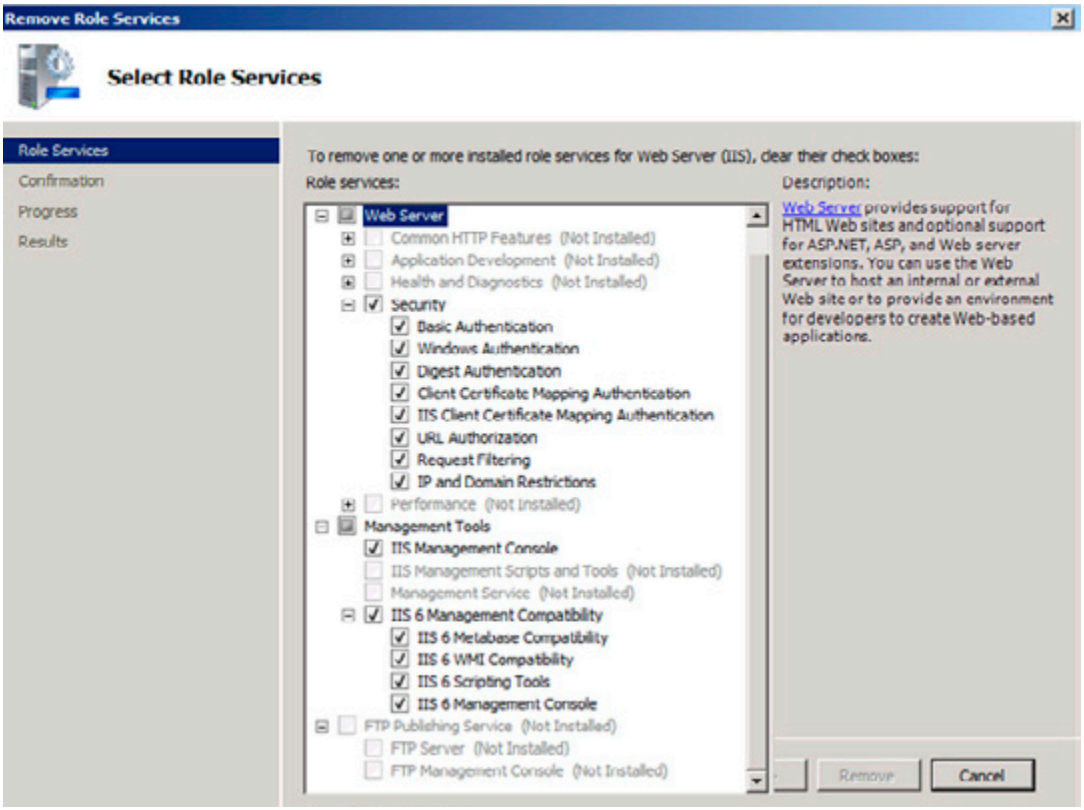

Figure 2.

Nowadays, you can procure what is known as "Unified Communications (UC) Certificates" from public root Certificate Authorities (CA) such as Entrust. This type of certificate is designed to be compatible with Exchange Server 2007 and OCS 2007. It supports multiple, unique domains and hostnames that are inherently more secure than wildcard certificates (the latter is not supported in any OCS versions). Also, added support for X.509v3 Subject Alternative Name (SAN) extensions makes deploying the same SSL certificate for multiple virtual hosts, applications or different SIP domains on the same physical box possible. It is worthwhile to note that certain R2 server roles require SAN certificate to be deployed. To make certificate management easier and further contain cost, consider setting up an internal Windows Server 2008 Enterprise root and sub-ordinate CAs integrated with AD and turn on the automatic web-enrollment feature. For external, public facing R2 server roles, you should continue to procure and deploy certificates from trusted public root CAs to avoid unnecessary errors from unknown or unrecognized certificates that are generated from private, internal CAs.

You should be able to execute the remaining setup steps to arrive at a functional side-by-side R2 environment comparable to an entirely new installation. I recommend that you run the R2 validation tests at the end and verify the core functionality with a small test bed of RTM and R2 users. Once you feel comfortable and confident enough, you can plan for the definite move of production users over to the new home pool.

Administration

Unlike their predecessors, you need to manually install the updated set of R2 administrative tools using the Office Communications Server 2007 R2 Deployment Wizard. They essentially come in 4 different flavors – a new Office Communications Server 2007 R2 Administrative snap-in, management components for Active Directory Users and Computers (dsa.msc), a snap-in extension for the Computer Management console plus the LCSCmd.exe command-line tool. Furthermore, other standalone administrative tools must be installed to administer new R2 applications such as the Group Chat Administration Tool and Response Group Service.

While R2 is available in a 64-bit edition only, the OCS 2007 R2 Administrative Tools are shipped in both x86 and x64 versions. Supported platforms include Windows Server 2003 SP2, Windows Server 2008 (x86, x64) and Vista Business or Enterprise with SP1 (x86, x64). Beware of the fact that you need to apply the relevant patches for the latter 2 platforms beforehand.

Installation of the R2 and OCS RTM administrative tools on the same box is not supported by Microsoft. Moreover, they can only be used to manage user settings on the respective OCS server versions, are not upward or backward compatible and cannot be mixed i.e. R2 administrative tools for R2 server pools only. An exception to this rule is the use of the Move User Wizard in the R2 or ADUC snap-in to transfer UC-enabled users homed on an OCS RTM server to a new R2 server pool. In fact, this is the easiest and recommended method to migrate a single user or in bulk over to the R2 platform through a graphical interface. Obviously, you can automate this with VBScript, PowerShell or your favorite scripting language.

Conclusion

You have learnt that there are two migration strategies you can consider in deploying OCS 2007 R2 in your organization. In this article, I focused on using the side-by-side methodology to migrate existing OCS RTM servers and clients over to R2. This approach can similarly be adopted for an entirely new R2 installation since it operates independently from existing OCS RTM or LCS 2005 SP1 in the environment.

After raising the Active Directory domain and forest functional levels to the required mode and laying the groundwork by relocating the OCS global settings, you apply the appropriate hotfixes to ensure peaceful co-existence of OCS RTM and R2. Once the relevant digital server certificates are prepared, we launched into the actual R2 setup process to deploy server roles and services. A successful deployment enables users to be moved to an R2 home server and migration to continue in phases with the remaining internal servers roles. I have also showed you that it is necessary to manually install the new and updated R2 administrative tools while keeping the previous versions to administer pre-R2 user settings.

In Part 2 of this article, I'll look into recommendations to standardize the client desktop to support the different R2 applications, other infrastructure concerns and migration of the remaining internal server roles before moving out to the Edge servers in the screened network.

References

Supported Migration Paths and Coexistence Scenarios
http://technet.microsoft.com/en-us/library/dd425356(office.13).aspx .

Migration From Office Communications Server 2007
http://technet.microsoft.com/en-us/library/dd572505(office.13).aspx .

Microsoft Office Communications Server 2007 R2 Resource Kit (ISBN: 9780735626355)
http://www.microsoft.com/learning/en/us/Books/13113.aspx .

Microsoft Office Communications Server 2007 R2 Documentation
http://www.microsoft.com/downloads/details.aspx?displaylang=en&FamilyID=e9f86f96-aa09-4dca-9088-f64b4f01c703.

PowerShell

Managing Exchange 2007 Mailbox Quotas with Windows PowerShell

24 November 2008

by Ben Lye

The use of PowerShell with Exchange Server 2007 can do a great deal to ease the task of managing mailbox quotas. You can, for example, use scripts to customize quota messages, retrieve mailbox sizes, and set the quotas. Ben shows you how, in the 2nd installment of his Top Tips for SysAdmins.

Most Exchange organizations use mailbox quotas to help users manage the size of their mailbox and to keep storage utilization in check. Managing quotas can be a burden on Exchange administrators, the help desk, and users, but with a few tweaks things can be made easier.

Exchange Mailbox Quota Types

Mailbox quotas can be set at the database or mailbox level. Quotas set at the mailbox level will take precedence over database quotas. Additionally, quotas can be a hard limit, preventing sending and/or receiving e-mail, or just a warning threshold which triggers a warning message. It's important to choose the right type of quotas to set – do you want users to be warned about their mailbox size, or do you want to prevent them from sending or receiving e-mail if they exceed their quota? The three mailbox quota types are:

Issue Warning Quota – this is not a hard limit, but a warning threshold. When it has been exceeded the user will get a warning message about their mailbox size, but will still be able to send and receive e-mail.

Prohibit Send Quota – this is a hard limit, and once a mailbox size exceeds it the user will no longer be able to send e-mail, but will still be able to receive e-mail.

Prohibit Send and Receive Quota – this is also a hard limit, and once it is exceeded the user will no longer be able to send or receive e-mail messages. Incoming e-mail will be returned to the sender.

Customizing Quota Messages

To make things easier for users, and to reduce the number of calls to the helpdesk, the messages which Exchange sends when a user is exceeding a quota limit can be customized to include more useful or practical information than the standard message gives. Custom quota messages also support HTML, so you can include formatting or other HTML, such as link to a self-help knowledgebase article offering tips on reducing mailbox size. In previous versions of Exchange modifying the quota messages was difficult and usually required programming skills, or custom DLL files, but with Exchange 2007 it can all be done with the PowerShell command line.

While there are only three mailbox quota types, there are four quota message types as a warning quota can be set with or without a hard limit quota. Custom messages can be set for all four quota message types:

WarningMailboxUnlimitedSize – this message type is sent to mailboxes which have no size limit when the warning quota has been exceeded.

WarningMailbox – this message type is sent to mailboxes which have a size limit (such as a Prohibit Send or Prohibit Send and Receive quota) when the warning quota has been exceeded.

ProhibitSendMailbox – This message type is sent when the Prohibit Send storage quota is exceeded.

ProhibitSendReceiveMailBox – This message type is sent when the Prohibit Send and Receive storage quota is exceeded.

The **New-SystemMessage** cmdlet is used to set a custom quota message:

```
New-SystemMessage -QuotaMessageType WarningMailbox -Language EN -Text "My custom quota message."
```

This command will set an HTML custom message which includes a hyperlink:

```
New-SystemMessage -QuotaMessageType WarningMailbox -Language EN -Text "<p>You are approaching the
maximum size of your mailbox.</p><p>Please see this article for details on how to reduce your
mailbox size: <a href=http://intranet.example.com/kb/mailboxsize.html> http://intranet.example.
com/kb/mailboxsize.html</a></p>"
```

Once a custom quota message has been set it can be viewed using the **Get-SystemMessage** cmdlet:

```
Get-SystemMessage -Identity EN\WarningMailbox | Format-List
```

It can be modified using the **Set-SystemMessage** cmdlet:

```
Set-SystemMessage -Identity EN\WarningMailbox -Text "My modified custom quota message."
```

And it can be removed using the **Remove-SystemMessage** cmdlet:

```
Remove-SystemMessage -Identity EN\WarningMailbox
```

Quota messages are sent according to a schedule defined on each mailbox database. The default schedule is daily, between 1am and 1.15am. The schedule can be altered using the **Set-MailboxDatabase** cmdlet.

This command will set the database "Mailbox Database" on server "EXCHANEG01" to send quota notifications on Sundays and Wednesdays between 7am and 8am:

```
Set-MailboxDatabase -Identity "EXCHANGE01\Mailbox Database" -QuotaNotificationSchedule
"Sun.7:00-Sun.8:00,""Wed.7:00-Wed.8:00"
```

You must be delegated the Exchange Organization Administrator role to use the *-SystemMessage cmdlets. You must be delegated the Exchange Server Administrator role and be a member of the local Administrators group for the target server to use the Set-MailboxDatabase cmdlet.

Retrieving Mailbox Sizes and Quotas

Mailbox sizes are retrieved using the **Get-MailboxStatistics** cmdlet.

```
Get-MailboxStatistics juser | fl TotalItemSize
```

Mailbox Quotas are retrieved using the **Get-Mailbox** cmdlet.

```
Get-Mailbox juser | fl *Quota
```

This simple script combines Get-MailboxStatistics with Get-Mailbox to show mailbox size and prohibit send quota in one command:

```
# Get-MailboxQuota.ps1
# Script for showing mailbox size and quota

# Exit the script if username is not found
If ($args[0] -eq $null) {
   Write-Host "Error: No user specified" -ForegroundColor "Red"
   break
}

# Get the username from the command line argument
$username = $args[0]

# Get the mailbox, break if it's not found
$mb = Get-Mailbox $username -ErrorAction Stop

# Get the mailbox statistics
$mbstats = Get-MailboxStatistics $username

# If the mailbox is using the database quotas then read them, otherwise read them from the mailbox
If ($mb.UseDatabaseQuotaDefaults -eq $true) {
   $quota = (Get-MailboxDatabase -Identity $mb.Database).ProhibitSendQuota.Value.ToMB()
} else {
   $quota = $mb.ProhibitSendQuota.Value.ToMB()
}

# Get the mailbox size and convert it from bytes to megabytes
$size = $mbstats.TotalItemSize.Value.ToMB()

# Write the output
Write-Host "Mailbox:  " $mb.DisplayName
Write-Host "Size (MB): " $size
Write-Host "Quota (MB):" $quota
Write-Host "Percent:  " ($size/$quota*100)
Write-Host
```

You must be delegated the Exchange View-Only Administrator role to use the Get-Mailbox and Get-MailboxStatistics cmdlets.

Setting Mailbox Quotas

Mailbox quotas can be set in two places – directly on the mailbox or on the mailbox database. By default Exchange 2007 sets a quota of 2000MB on all new mailbox databases, and all mailboxes in the database inherit this value.

The **Set-MailboxDatabase** cmdlet is used to set default quotas on a mailbox database using the PowerShell command line.

This command will set the default warning quota on the database "Mailbox Database" on server EXCHANGE01 to 975MB, and the limit at which users will no longer be able to send mail to 1000MB:

```
Set-MailboxDatabase "EXCHANGE01\Mailbox Database" -IssueWarningQuota 975MB -ProhibitSendQuota
1000MB
```

The **Set-Mailbox** cmdlet is used to set quotas on individual mailboxes.

This command will set the warning quota for user juser to 1475MB, and the limit at which the user will no longer be able to send mail to 1500MB. It will also configure the mailbox not to use the database default quotas:

```
Set-Mailbox juser -IssueWarningQuota 1475MB -ProhibitSendQuota 1500MB –UseDatabaseQuotaDefaults
$false
```

Quota increase requests will be fairly common for most organizations which use mailbox quotas. Quota increases are usually governed by an IT policy, and increases are usually in fixed amounts. This PowerShell script will automatically increment the quota size of a specified mailbox by a given amount. This script, or something like it, can be used to decrease the administrative overhead of mailbox quotas. The script reads the current quota from the database or from the mailbox, shows what the existing quota is and what the new quota will be, then prompts for confirmation before setting the new quotas. It then displays confirmation that the new values have been set. If the current quota is not a multiple of the increment specified it will be rounded up to the next increment, rather than having an increment added, which ensures that quotas are always a multiple of the desired increment value.

```
# Increase-MailboxQuota.ps1
# Script for incrementing mailbox quotas

# Amount to increase prohibit send quota by in megabytes
$QuotaIncrement = 250

# Amount to subtract from prohibit send quota to set warning quota
$WarningDifference = 25

# Get the username from the arguments
$username = $args[0]
# Prompt if no location was passed
if (-not $username) {
   $username = Read-Host "Username"
}

# Get the mailbox
$Mailbox = Get-Mailbox -Identity $username -ErrorAction Silentlycontinue

# Error if the mailbox wasn't found
If (-not $mailbox) {
   Write-Host "User not found" -Foregroundcolor:Red
   break
}
```

```
# Get the mailbox information and size
$DisplayName = $Mailbox.DisplayName
$Database = $Mailbox.Database
$UsingDBQuotas = $Mailbox.UseDatabaseQuotaDefaults
$MailboxSize = (Get-MailboxStatistics -Identity $Mailbox.Name).TotalItemSize.value.ToMB()


# Get the current quota values
if ($UsingDBQuotas -eq $True)
   {
      # Database quotas are being used so read them from the DB
      $Database = Get-MailboxDatabase -Identity $Database
      $ProhibitSendQuota = $Database.ProhibitSendQuota.value.ToMB()
      $IssueWarningQuota = $Database.IssueWarningQuota.value.ToMB()
   }
else
   {
      # Mailbox quotas are being used so read them from the mailbox
      $ProhibitSendQuota = $Mailbox.ProhibitSendQuota.value.ToMB()
      $IssueWarningQuota = $Mailbox.IssueWarningQuota.value.ToMB()
   }

# Calculate the new prohibit send quota
If (($ProhibitSendQuota % $QuotaIncrement) -eq 0) {
   # Existing quota is a multiple of $QuotaIncrement so increase it by $QuotaIncrement
   $NewProhibitSendQuota = $ProhibitSendQuota + $QuotaIncrement
} Else {
   # Existing quota is not a multiple of $QuotaIncrement so round it up to nearest multiple of
$QuotaIncrement
   $NewProhibitSendQuota = $ProhibitSendQuota + ($QuotaIncrement - ($ProhibitSendQuota %
$QuotaIncrement))
}

# Calculate the new warning value
$NewIssueWarningQuota = $NewProhibitSendQuota - $WarningDifference

# Show what we're going to do
Write-Host ""
Write-Host "Full Name:       ," $DisplayName
Write-Host "Database:        ," $Database
Write-Host "Using Default Quota: ," $UsingDBQuotas
Write-Host ""
Write-Host "Mailbox Size (MB):  ," $MailboxSize, "MB"
Write-Host ""
Write-Host "Current Quota:     ," $ProhibitSendQuota, "MB"
Write-Host "Current Warning:    ," $IssueWarningQuota, "MB"
Write-Host ""
Write-Host "New Quota:       ," $NewProhibitSendQuota, "MB"
Write-Host "New Warning:      ," $NewIssueWarningQuota, "MB"
Write-Host ""
$Continue = Read-Host "Continue [Y/N]?"

# Ask if we want to continue
Switch ($Continue) {
   "Y" {$Continue = $True}
   "y" {$Continue = $True}
```

```
}

# Stop here if not continuing
If ($Continue -ne $True) {
   break
}

# Set the new values on the mailbox
$NewProhibitSendQuota = [STRING]$NewProhibitSendQuota + "MB"
$NewIssueWarningQuota = [STRING]$NewIssueWarningQuota + "MB"
Set-Mailbox -Identity $Mailbox -UseDatabaseQuotaDefaults $False -ProhibitSendQuota
$NewProhibitSendQuota -IssueWarningQuota $NewIssueWarningQuota

# Update the mailbox quota information
$Mailbox = Get-Mailbox $Mailbox
$ProhibitSendQuota = $Mailbox.ProhibitSendQuota.value.ToMB()
$IssueWarningQuota = $Mailbox.IssueWarningQuota.value.ToMB()

# Write some output to confirm the new values
Write-Host ""
Write-Host "Updated Quota:     ," $ProhibitSendQuota, "MB"
Write-Host "Updated Warning:    ," $IssueWarningQuota, "MB"
Write-Host ""
```

You must be delegated the Exchange Recipient Administrator role to use the Set-Mailbox cmdlet

Configuring the Mailbox Information Cache Refresh Interval

Exchange quota information is stored in Active Directory, and by default is cached by Exchange for up to two hours. This means that it could take up to two hours for a quota change to take effect. The recommended interval for Exchange to refresh quota information is 20 minutes, which can be set by adding three registry values.

> ***Note***
>
> *Setting the cache refresh intervals too low can adversely affect the performance of Exchange. Incorrectly editing the registry can cause serious problems that may require you to reinstall your operating system. Problems resulting from editing the registry incorrectly may not be able to be resolved. Before editing the registry, back up any valuable data. You need to be a local administrator on the Exchange server in order to edit the registry.*

- Start the registry editor on your Exchange 2007 Mailbox server
- Locate the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\MSExchangeIS\ParametersSystem key.

- Create the "Reread Logon Quotas Interval" value
  - Right-click ParametersSystem, select New, and then select DWORD value.
  - Name the new DWORD value "Reread Logon Quotas Interval."
  - Right-click Reread Logon Quotas Interval, and then click Modify.
  - Enter a decimal value of 1200 seconds (20 minutes)
- Create the "Mailbox Cache Age Limit" value
  - Right-click ParametersSystem, select New, and then select DWORD value.
  - Name the new DWORD value "Mailbox Cache Age Limit."
  - Right-click Mailbox Cache Age Limit, and then click Modify.
  - Enter a decimal value of 20 (20 minutes)
- Locate the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\MSExchange ADAccess key.
- Create the "CacheTTLUser" value
  - Right-click MSExchange ADAccess, select New, and then select Key.
  - Name the new key Instance0.
  - Right-click Instance0, select New, and then select DWORD value.
  - Name the new DWORD value "CacheTTLUser."
  - Right-click CacheTTLUser, and then click Modify.
  - Enter a decimal value of 300 (5 minutes)

Alternatively, copy this text file and paste it into a file called MailboxCache.reg, then import it into the registry of each of your Exchange 2007 Mailbox servers

```
Windows Registry Editor Version 5.00
[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\MSExchangeIS\ParametersSystem]
"Reread Logon Quotas Interval"
=dword:000004b0
"Mailbox Cache Age Limit"
=dword:00000014
[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\MSExchange ADAccess\Instance0]
"CacheTTLUser"
=dword:0000012cc
```

The Exchange Information Store service needs to be restarted for the change to become effective. More information about these changes can be found on the Microsoft TechNet website.

So You Thought PowerShell Was Only For Exchange 2007

08 May 2009

by Jonathan Medd

PowerShell makes a lot of sense as a means of gathering information, and automating large sections of your administration tasks. It is not just for the latest version of Exchange, though. Powershell can be used on previous versions too. Jonathan Medd gets you started!

Introduction

When Microsoft released Exchange 2007 they built the Exchange Management Console on top of Windows PowerShell 1.0; when you execute commands from the console, underneath it uses PowerShell cmdlets to carry out the requested actions. Exchange 2007 SP1 ships with around 400 native PowerShell cmdlets to efficiently configure, manage and automate your messaging environment; everything from mailbox and database management, through junk mail settings and public folders. Whilst this has been a fantastic move forward for those organisations who have either migrated from previous Exchange versions or other messaging platforms or even a completely fresh Exchange 2007 installation not every Exchange administrator is lucky enough to be in that boat and may well still be managing an Exchange 2003 (or earlier) environment.

The aim of this article is to demonstrate how you can use standard PowerShell techniques to query information stored in Event Logs, WMI and Active Directory to make management of your pre-Exchange 2007 environment more efficient than with the standard GUI based tools provided.

What Do I Need To Get Started

Installation of Exchange 2007 has a dependency on Windows PowerShell 1.0 being installed on the server so that you are able to execute commands from the Exchange Management Console. For the purposes of what we are going to look at it is not recommended that you install PowerShell on your Exchange 2003 server, rather install it on your management workstation. (It would be unlikely to cause an issue on your Exchange 2003 server, but it's not good practise to install unnecessary software there).

On your management workstation you require the following:

- Windows XP (or above)
- .NET Framework 2.0 or above

PowerShell downloadable executable, there are different versions for XP, Vista and x86 or x64 – available here http://www.microsoft.com/windowsserver2003/technologies/management/powershell/download.mspx. .

For the WMI queries we are going to run you will need to provide an account with local admin rights on the Exchange Server since that is a requirement for remote WMI access.

Querying Event Logs Using WMI

Later on we'll come onto using some of the Exchange specific WMI classes which are available to us for extracting information from an Exchange server, but for now I will show you how you can use WMI to query the Event Logs on any Windows based server.

Windows Management Instrumentation (WMI) is the infrastructure for management data and operations on Windows-based operating systems, think of it as a database of the OS. You may have previously used other tools for querying WMI such as WMIC or VBScript and consequently not had the easiest experience with WMI – I recently heard WMI described as "Voodoo'! It doesn't need to be like this either with the previously mentioned tools or best of all with PowerShell which makes WMI your friend.

Using the PowerShell cmdlet Get-WmiObject you can query WMI on local or remote computers and easily obtain valuable results. For instance with a very simple example you could use **Get-WmiObject Win32_ComputerSystem** to return information about the local computer.

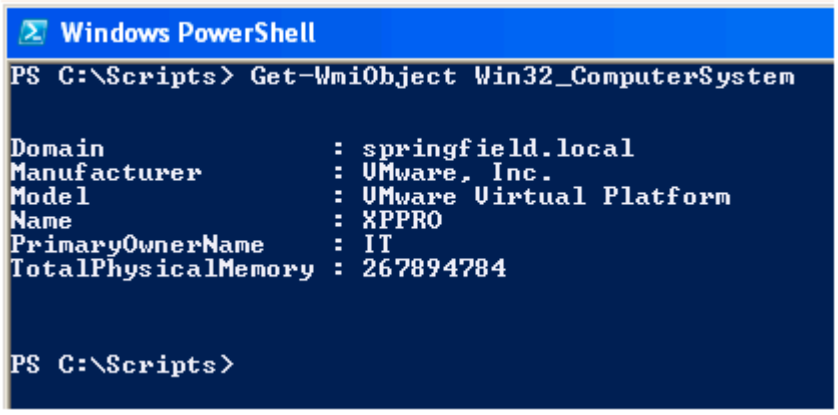

PowerShell by default will return a standard set of properties for the WMI class you have queried as above, however there are typically many more properties hidden away underneath which can be exposed. You can pipe the results of your WMI query to the **Format-List** cmdlet to discover what properties are available to you.

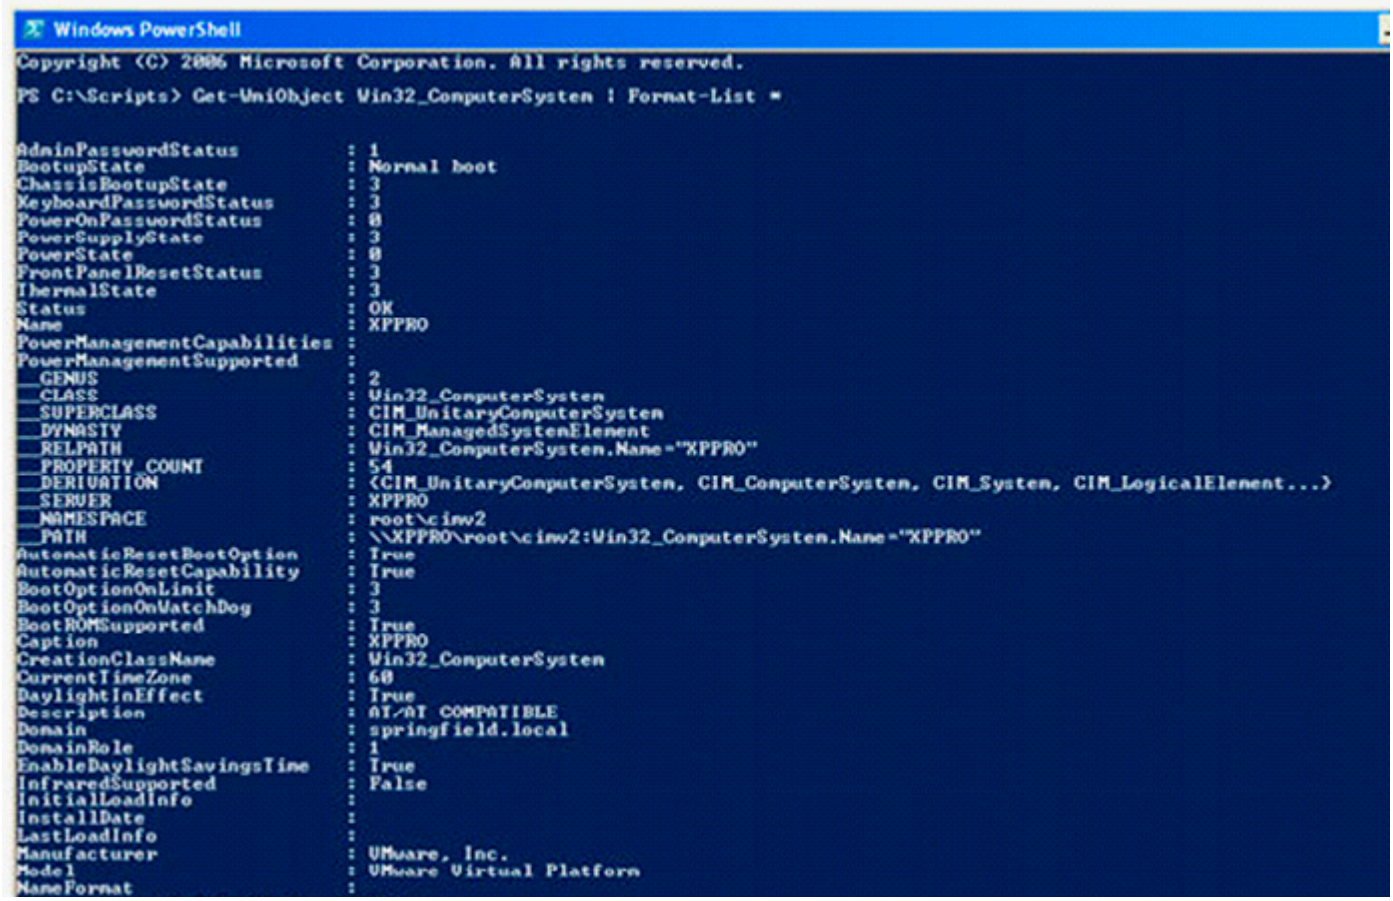

Now you know what's available you can stipulate the particular properties you wish to display. PowerShell will then only output the results for the information you are specifically interested in.

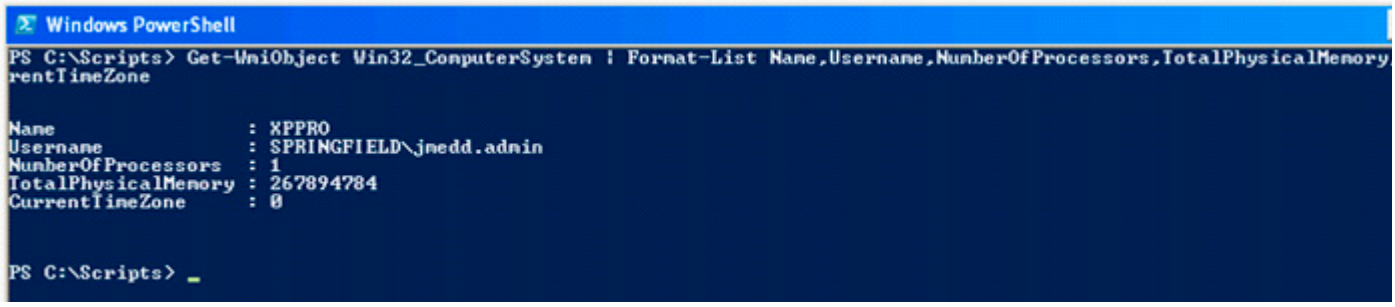

Moving on, we can now use some slightly more advanced techniques to query the event logs on remote machines. If you are an experienced Exchange administrator it is likely that at some point you will have taken your mailbox databases offline so that they can be defragged to reclaim the white space inside them. In a large environment you may have hundreds of databases and looking for candidates for defragging could be a time consuming process, typically you would use an event log entry similar to that below from the Application log which shows how much free space is in a database after the online defragmentation process has completed. You would not want to manually browse through these records on multiple Exchange servers over multiple mailbox stores.

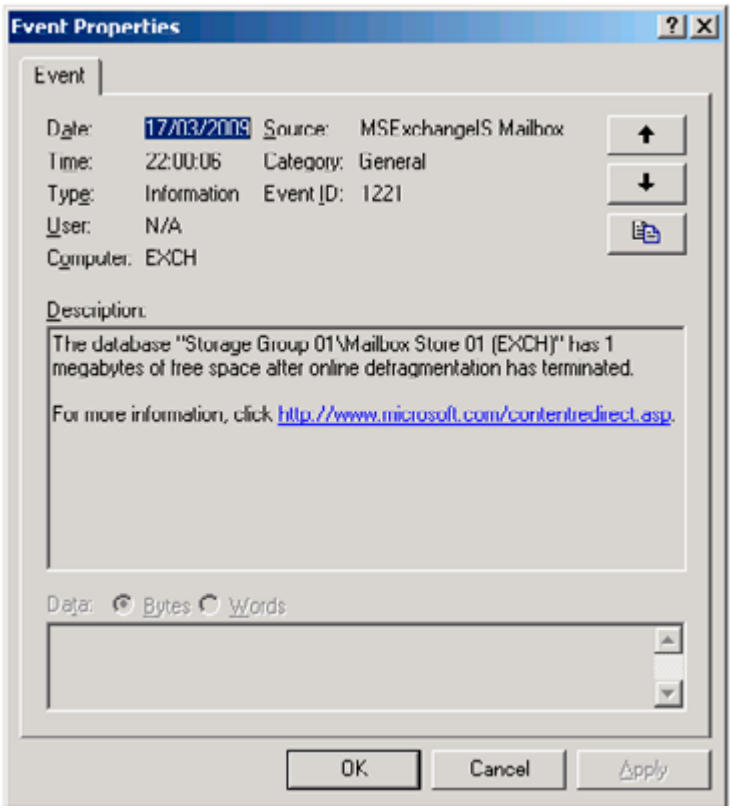

Instead, it is far better to use PowerShell to query those entries in the event logs and return to you only the information you need. For instance, give you only those databases which have over 2GB free space in them. The following command will carry out most of this work for you:

```
Get-WmiObject
-computer ExchangeServerName-query("Select * from Win32_NTLogEvent Where Logfile='Application'
and Eventcode = '1221' and TimeWritten >='"+$WmidtQueryDT+"'")
```

(Do not worry too much about **$WmidtQueryDT**, it's supplied for you in the full script file, **Get-FreeSpace.ps1**, where we use some .NET code to get a date into a format that WMI likes)

Step-by-step we use the:

- Get-WmiObject cmdlet to query the remote server ExchangeServerName by using the –computer parameter.
- Use the query parameter to specify a SQL style query.
- Select everything from the Win32_NTLogEvent WMI class.
- Filter it on the Application Log.
- Filter it on Event ID 1221.
- Filter it on the last 24 hours – this assumes that you run online maintenance on your databases daily.

This will return the information we need. In the script file we then run some extra code to read those log files and filter it down so that only those whose Message Text includes a value over 2GB and export the results to a CSV file ready for use.

Exchange 2003 WMI classes

When you install Exchange 2003 some additional WMI classes will be added specific to Exchange information which you are able to query. These are documented on MSDN (http://msdn.microsoft.com/en-us/library/ms876479(EXCHG.65).aspx) where you can find full details of what they have to offer and some VBScript examples of how you might use them. (Note: these classes have been removed from Exchange 2007 and above.) Fortunately for you I'm going to show you how to use them in PowerShell instead. The best one to start with is the Exchange_Mailbox class.

So far I have not mentioned WMI namespaces. In the previous examples the WMI classes we have been using are all part of the default WMI namespace Root\CIMV2. Since it's the default, PowerShell assumes that's the namespace you wish to use if you don't specify one. The Exchange 2003 classes do not live in this namespace; they belong to the root\MicrosoftExchangeV2 namespace, so when using the Get-WmiObject cmdlet with these classes we must use the namespace parameter. A simple command to query an Exchange server about its mailboxes is the following:

```
Get-WMIObject
-namespaceroot\MicrosoftExchangeV2-classExchange_Mailbox-computer ExchangeServerName
```

This will return information about mailboxes across all Mailbox Stores on the specified Exchange server. By using a couple of additional standard PowerShell cmdlets we can produce more readable output. Firstly we can use the **Sort-Object** cmdlet and specify the property **MailboxDisplayName** so that the mailboxes are returned in alphabetical order. Then we can use **Format-Table** to specify which properties are returned.

```
Get-WMIObject
-namespaceroot\MicrosoftExchangeV2-classExchange_Mailbox-computer ExchangeServerName| Sort-
ObjectMailboxDisplayName| Format-TableMailboxDisplayName,Servername,StorageGroupName,StoreName,
Size-auto
```

The following example shows the kind of results this would produce.

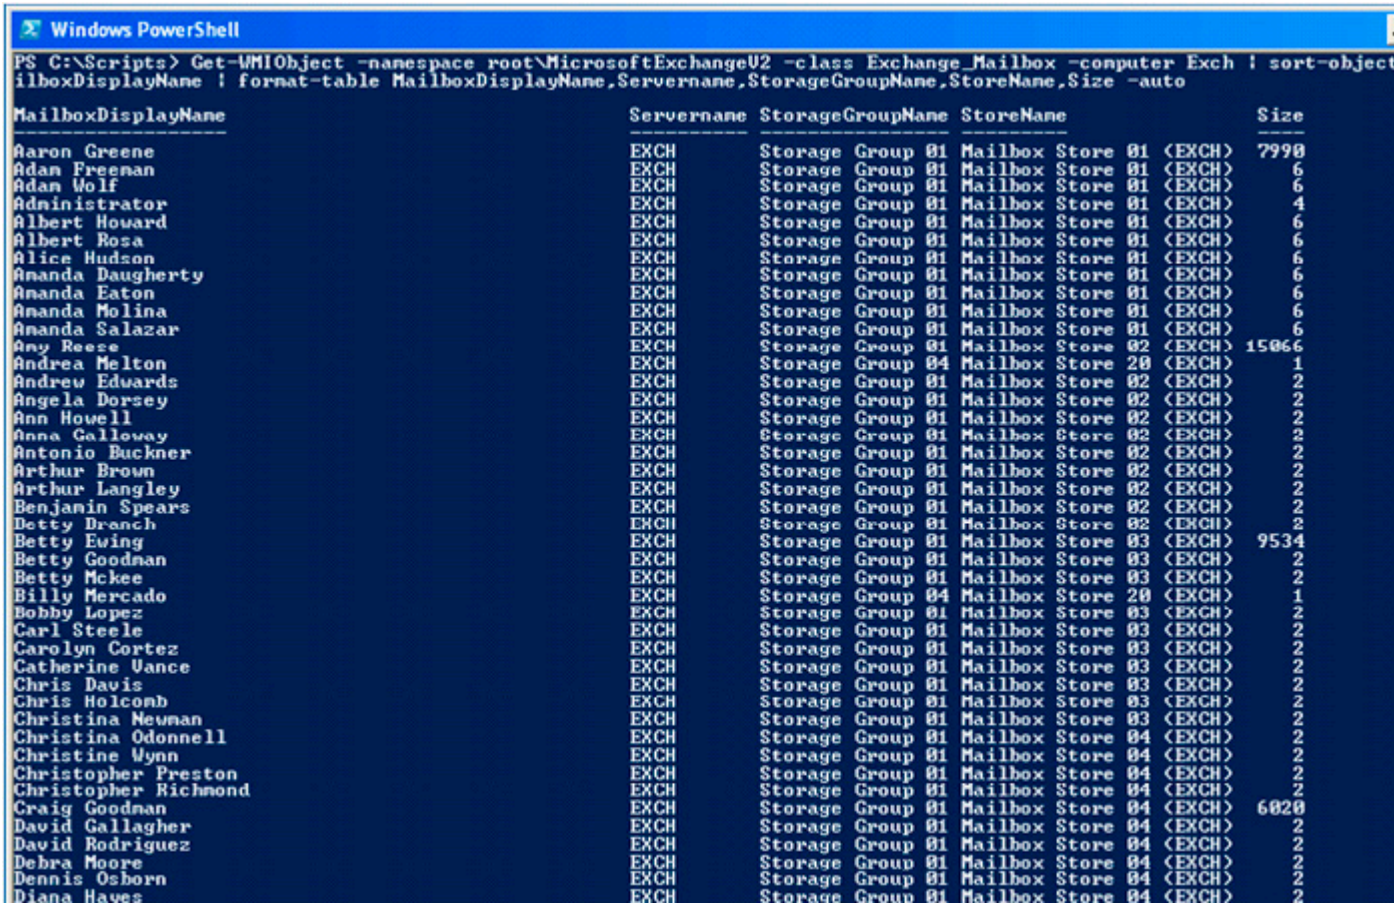

A common scenario for Exchange management is the demand from users for large mailboxes and consequently lots of storage. Your manager may well want to know who is using up all the extra storage he only just paid for last month that you told him you needed. Using the GUI Exchange tool it is not easy to aggregate this information across your server farm if you have multiple databases over many Exchange servers. However, we can take the previous example, make some modifications to it and hey presto you will have the report for your manager.

This time instead of sorting by **MailboxDisplayName** we sort by the **Size** of the mailbox, then we use the **Select-Object** cmdlet and use the –First parameter to return only a subset of the results – in this case 10. This will give us the top 10 largest mailboxes (you can obviously adjust the number of results returned to suit your needs). Since we return the Exchange Servername, Storage Group and Mailbox Store it also makes it simple to track these mailboxes down.

```
Get-WMIObject
-namespaceroot\MicrosoftExchangeV2-classExchange_Mailbox-computer ExchangeServerName| Sort-
ObjectSize-Descending| Select-Object-First10 | Format-TableMailboxDisplayName,Servername,Storag
eGroupName,StoreName,Size-auto
```

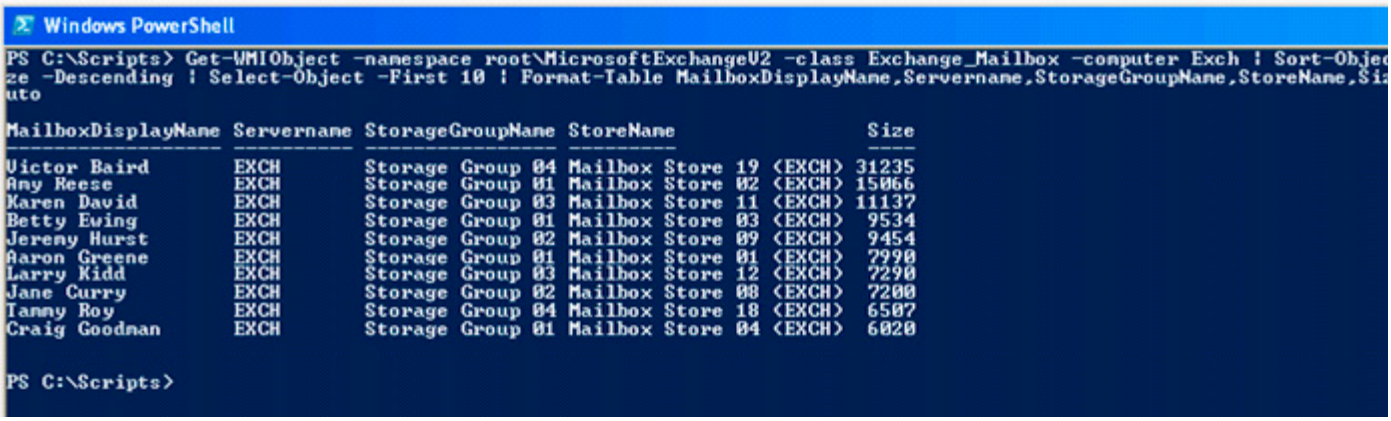

It's not just mailbox data we can query though with **Get-WmiObject**, another useful Exchange WMI class is the **Exchange_Logon** class. Not surprisingly this will tell us about who is logged into Exchange and other useful information like what client they are using, for instance you might wish to report on which of your users are logged in using Outlook Web Access.

This time we use the **–filter** parameter of **Get-WmiObject** to return a restricted amount of data, in this case where the **ClientVersion** reported by the **Exchange_Logon** class matches "HTTP" ; we also exclude entries where the **LoggedOnUser** is NT Authority\System which is something we are obviously not particularly interested in. We may well get multiple results per user returned via the Exchange_ Logon class for the same client connection so when sorting the data with **Sort-Object** we use the **–unique** parameter so that we only get one relevant result in the output.

```
Get-WMIObject
-namespaceroot\MicrosoftExchangev2-classExchange_Logon-ComputerNameExchangeServerName-
filter"ClientVersion = 'HTTP' and LoggedOnUserAccount != 'NT AUTHORITY\\SYSTEM'"| Sort-
ObjectMailboxDisplayname-unique| Format-TableLoggedOnUserAccount,MailboxDisplayName,Servername,
StorageGroupName,StoreName,ClientVersion-auto
```

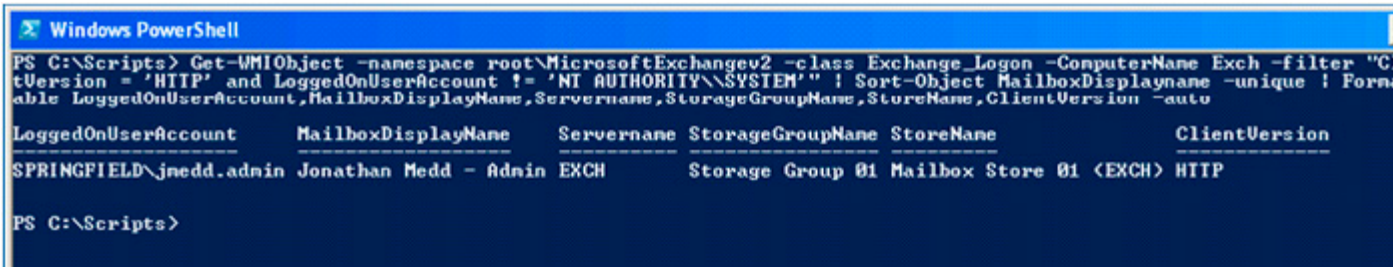

(In the demo environment used for this article I was not able to simulate multiple logons, but you can imagine the kind of output you would receive in a larger environment.)

As an Exchange administrator another common best practise is to keep mailboxes spread evenly over Mailbox Stores and Storage Groups as best is possible. To do that though you need the information easily to hand of the number of mailboxes across these areas; again the out of the box Exchange GUI tools do not help you with this task, particularly across large environments. PowerShell can again come to your rescue and very simply provide this information.

We use the **Exchange_Mailbox** class again, but this time instead of returning names, locations and sizes of mailboxes we count the results and return that data. We store in a text file the names of the Storage Groups we are interested in (as below) and use the **Get-Content** cmdlet to store these names into a variable.

```
$StorageGroups
=Get-ContentC:\Scripts\StorageGroups.txt
```

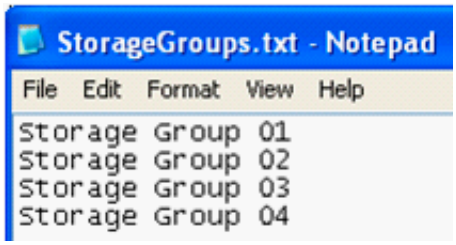

We then use a **foreach** statement to loop through each of these Storage Groups and count the number of mailboxes in each one. You will see below that the WMI query is filtered on the **StorageGroupName** property and the entire command is encapsulated in brackets with .count added at the end to give us the total in that query. Finally **Write-Host** is used to display some text and the results to the screen.

```
foreach
($SGin$StorageGroups)
{
$Total=(Get-Wmiobject-namespaceroot\MicrosoftExchangeV2-classExchange_Mailbox-computer
ExchangeServerName-filter"StorageGroupName = '$SG'").count;Write-Host$SG":"$Total"users"}
```

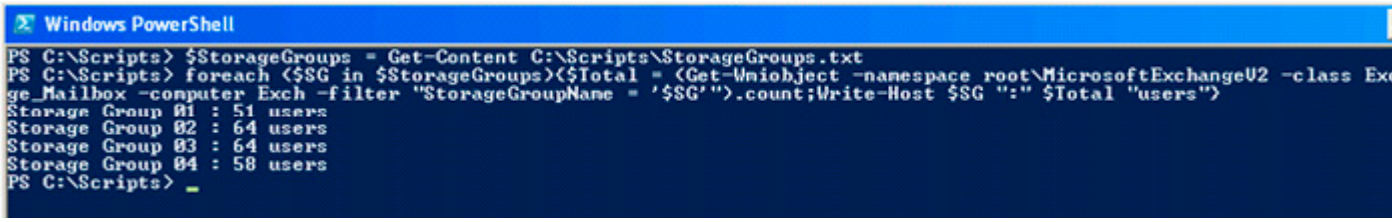

It doesn't take a genius to extend this to be more granular and look at Mailbox Stores instead of Storage Groups.

Provide your list of Mailbox Stores to PowerShell with

```
$MailboxStores
=Get-Contentc:\Scripts\MailboxStores.txt
```

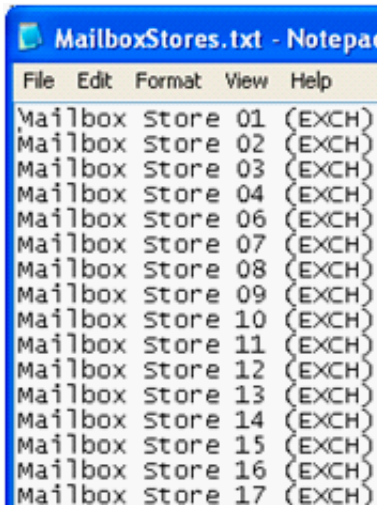

Then use a **foreach** statement to loop through each of the stores and return the number of users in each one.

```
foreach
($Storein$MailboxStores)
{
$Total=(Get-WmiObject-namespaceroot\MicrosoftExchangeV2-classExchange_Mailbox-computer
ExchangeServerName-filter"Storename = '$Store'").count;Write-Host$Store":"$Total"users"}
```

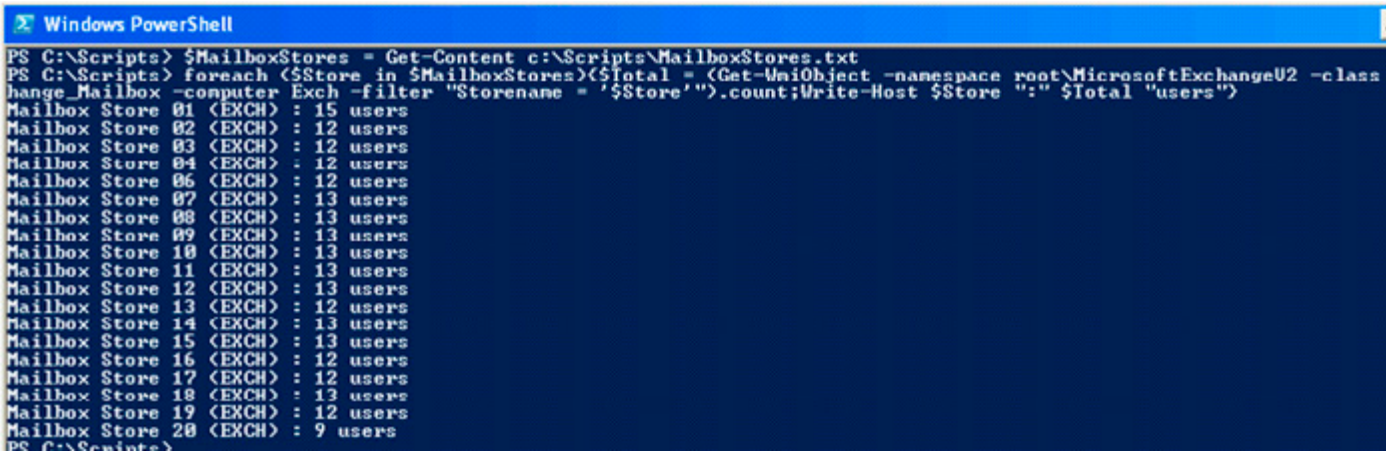

Querying Active Directory to Determine Exchange Information in a Network

Another way we can use PowerShell to garner information about our Exchange environments is to query Active Directory which stores a lot of Exchange configuration information. One of the ways to do this is to use a couple of classes within the .NET DirectoryServices Namespace which is a .NET wrapper for ADSI (http://msdn.microsoft.com/en-us/library/system.directoryservices(VS.71).aspx), in particular the DirectoryEntry and DirectorySearcher classes.

I have heard stories of administrators being put off getting into PowerShell because they've been told that you need to be a .NET programmer to use it – this is so not true. You don't need to know any .NET to effectively use PowerShell, however PowerShell does have access to .NET so it can be to your advantage and potentially open a few doors if you just explore some of the basics.

In this example we are going to query Active Directory to find out the names of the Exchange servers in our environment. First of all we create a DirectoryServices.DirectoryEntry object for the current Active Directory domain.

```
$root
=New-ObjectSystem.DirectoryServices.DirectoryEntry("LDAP://RootDSE")
```

Then we use ADSI to get a reference to the configuration naming context within AD. All Exchange information in AD (except for per-recipient information) is in the configuration naming context.

```
$configpartition
=[adsi]("LDAP:// CN=Microsoft Exchange,CN=Services," +$root. configurationNamingContext)
```

Create a directory search object:

```
$searcher
=New-ObjectSystem.DirectoryServices.DirectorySearcher($configpartition)
```

Filter the search on the Exchange Server objectclass:

```
$searcher
.filter ='(objectclass=msExchExchangeServer)'
```

Use the FindAll method to execute the search:

```
$ExchServer
=$searcher.FindAll()
```

Finally, return the names of all the results:

```
$ExchServer| foreach{$_.properties.name}
```

In my demo environment there is only one Exchange server so the results are not particularly exciting, however you can imagine how useful this could be in a large deployment of Exchange servers.

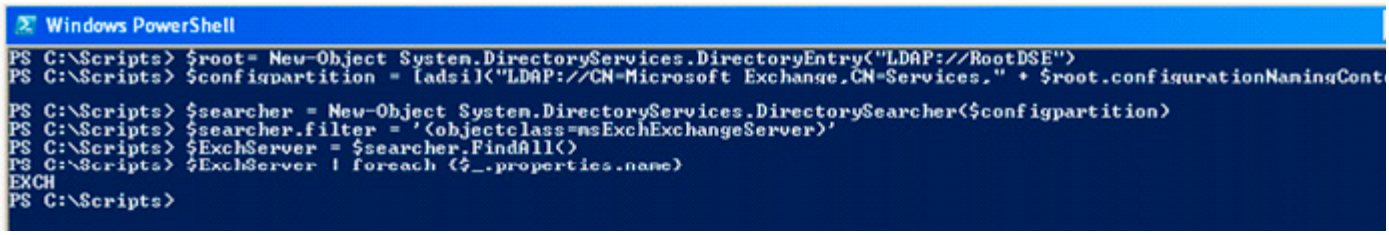

To save typing all of those lines out each time you want to run the code you could turn it into a filter like this:

```
filterGetExchangeServers{
$root
=New-ObjectSystem.DirectoryServices.DirectoryEntry("LDAP://RootDSE")
$configpartition
=[adsi]("LDAP:// CN=Microsoft Exchange,CN=Services," +$root. configurationNamingContext)
$searcher
=New-ObjectSystem.DirectoryServices.DirectorySearcher($configpartition)
$searcher
.filter ='(objectclass=msExchExchangeServer)'
$ExchServer
=$searcher.FindAll()
$ExchServer
| foreach{$_.properties.name}
}
```

and then you can either include it in a script or if it's something you might run regularly then include it in your PowerShell profile (http://www.microsoft.com/technet/scriptcenter/topics/winpsh/manual/profile.mspx).

Note

A filter is similar to a function in PowerShell, however it enables you to take advantage of the pipeline.

What we have now though is something potentially very powerful. By putting the GetExchangeServers filter together with one of our earlier WMI queries we can now run that query against all of our Exchange Servers without even having to specify their names!

To find the top 10 largest mailboxes on each Exchange server run the GetExchangeServers filter and pipe it to the Get-WmiObject command.

```
GetExchangeServers |
ForEach-Object{Get-WmiObject-namespaceroot\MicrosoftExchangeV2-classExchange_Mailbox-computer
$_| Sort-ObjectSize-Descending| Select-Object-First10 | Format-TableMailboxDisplayName,Servernam
e,StorageGroupName,StoreName,Size-auto}
```

Other objectclasses you could also query in Active Directory include:

- Storage Groups: `msExchStorageGroup`.
- Mailbox Databases: `msExchPrivateMDB`.
- Public Folder Databases: `msExchPublicMDB`.
- Simply replace the value for the `objectclass` in the `$searcher.filter` line.

Exchange 2003 Powerpack for PowerGUI

You might be reading this article and thinking "Well this scripting stuff is all very well, but it's really not for me, I'm a GUI kind of administrator and what I'd really like is someone just to package up all these scripts for me and let me run them from the click of a button."

If that's you then you are in luck because Quest has created a tool to make that dream a reality by creating a fabulous free tool called PowerGUI (http://powergui.org). Not only do you get a brilliant scripting editor (yes for free) the package also includes a tool which lets you create custom management consoles. So say you don't like the console which ships with Exchange or it has something missing and you don't want to wait for the next product release, you can create your own console by packaging up PowerShell scripts into PowerPacks – or even better download one which somebody else has already created for you.

Simply download and install PowerGUI, then head over to the PowerPacks part of the site (http://powergui.org/kbcategory.jspa?categoryID=21) and you will find many pre-made PowerPacksfor a variety of different products including Exchange, SQL, VMware, Citrix, etc ready to download.

The PowerPacks are completely open so you can modify / add / delete any part of them to make it more suitable for your own environment. However, as of version 1.7 it is also now possible to lock down and brand a PowerPack, for instance if you wanted to provide one for helpdesk use.

I made a PowerPack for Exchange 2003 (http://powergui.org/entry.jspa?externalID=1956&categoryID=47) which when plugged into PowerGUI will give you easy access to most of the scripts already mentioned in this article plus many more which you can see below.

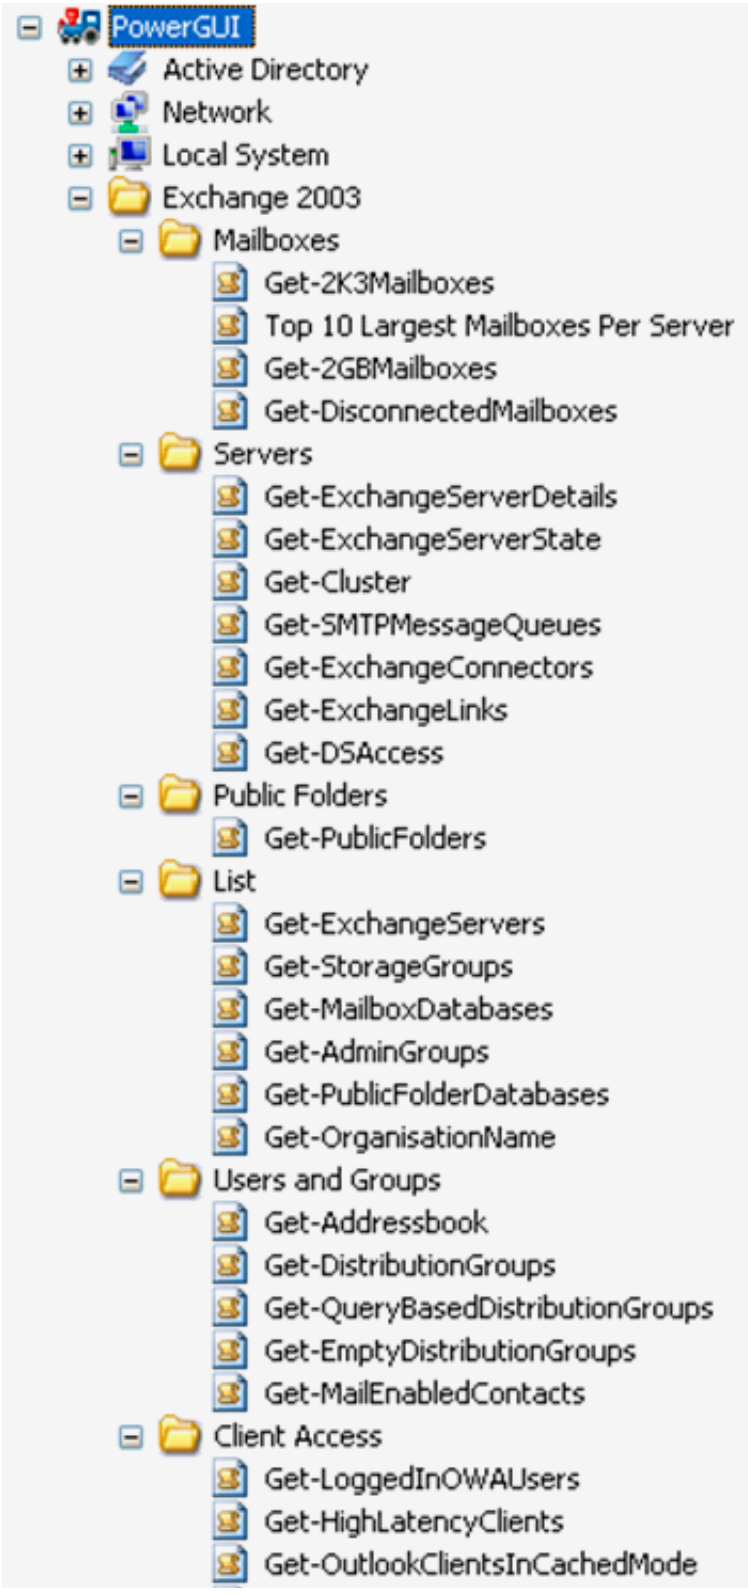

Simply click one of the script nodes and the script will go off and do its work in the background and return the results into the grid pane. For instance if you chose "Top 10 Largest Mailboxes Per Server" the exact same PowerShell script as previously mentioned will be used and you will see the results like below.

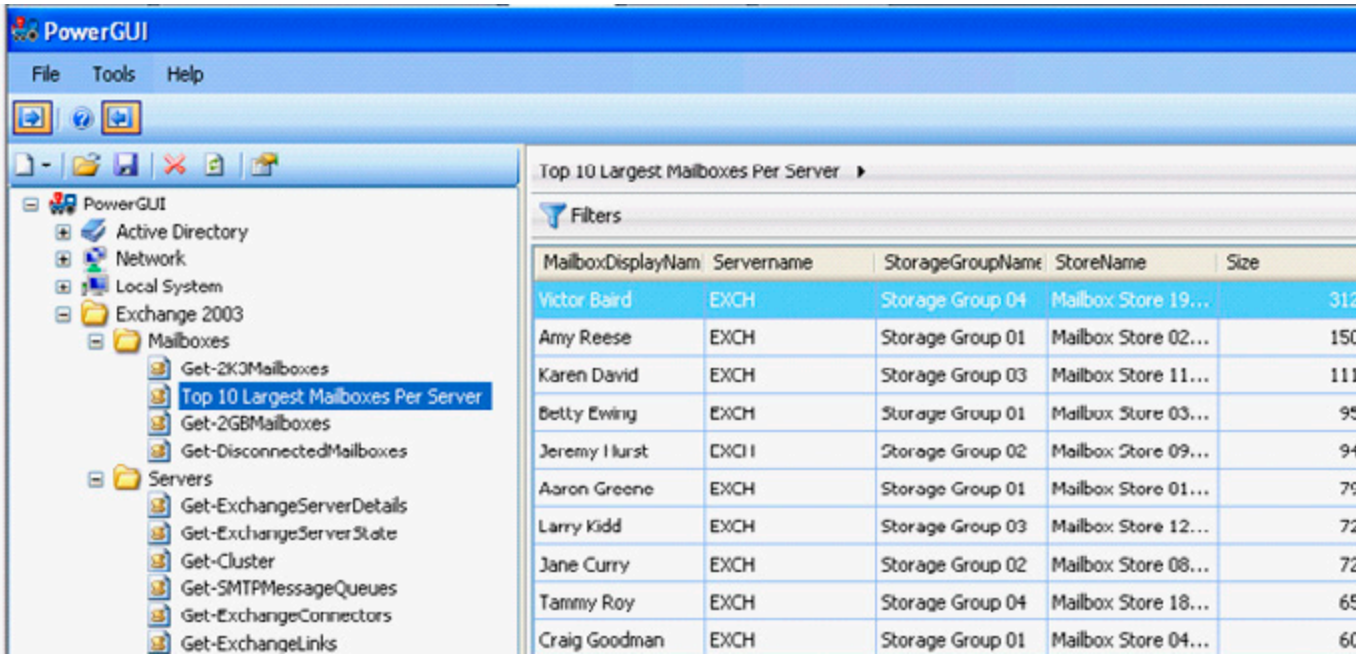

Included is a section for managing the user and groups part of Exchange management. These require the AD PowerShell cmdlets also provided free by Quest (http://www.quest.com/activeroles-server/arms.aspx). A good example for their use would be the node "Get-EmptyDistributionGroups'. It's pretty common for a distribution group to be set up for short-term use, say for a project and then at the end of the project people are removed from the group, but nobody ever thinks to tell the administrator that the group is no longer needed. Simply run this PowerShell script and it will give you a list of all distribution groups which are empty – it's even got a **Delete** button so that you can remove them from the same console.

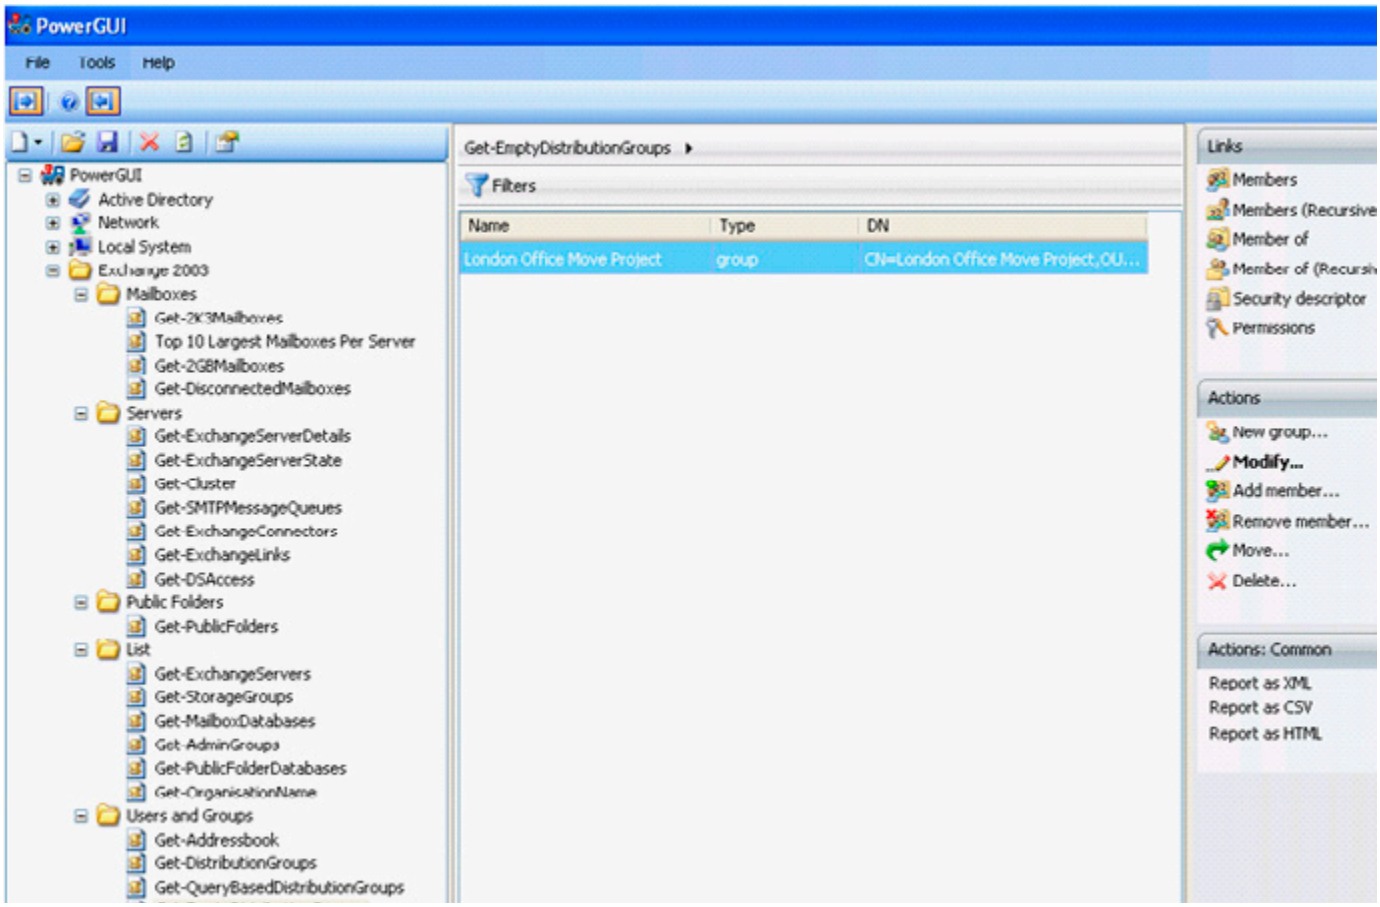

Conclusion

So you have seen that just because you may not be using Exchange 2007 you don't have to be a second-class citizen in terms of managing Exchange from scripts or the command line.

Using PowerShell and underlying technologies like WMI and Active Directory you can quickly and simply gather information from your Exchange environment where using the standard GUI tools for the same tasks requires far more manual effort. Make one further step by scheduling these scripts to run and you will soon be automating large sections of your admin tasks and be free to get on with those interesting projects you never have time to work on.

About Red Gate

You know those annoying jobs that spoil your day whenever they come up?

Writing out scripts to update your production database, or trawling through code to see why it's running so slow.

Red Gate makes tools to fix those problems for you. Many of our tools are now industry standards. In fact, at the last count, we had over 650,000 users.

But we try to go beyond that. We want to support you and the rest of the SQL Server and .NET communities in any way we can.

First, we publish a library of free books on .NET and SQL Server. You're reading one of them now. You can get dozens more from www.red-gate.com/books

Second, we commission and edit rigorously accurate articles from experts on the front line of application and database development. We publish them in our online journal Simple Talk, which is read by millions of technology professionals each year.

On SQL Server Central, we host the largest SQL Server community in the world. As well as lively forums, it puts out a daily dose of distilled SQL Server know-how through its newsletter, which now has nearly a million subscribers (and counting).

Third, we organize and sponsor events (about 50,000 of you came to them last year), including SQL in the City, a free event for SQL Server users in the US and Europe.

So, if you want more free books and articles, or to get sponsorship, or to try some tools that make your life easier, then head over to www.red-gate.com