

Transitioning to Microsoft Exchange Server 2007 on HP ProLiant servers and HP StorageWorks



Code notice	2
Executive summary.....	4
Introduction.....	4
Benefits of moving to Exchange Server 2007	4
Preparing for a transition to Microsoft Exchange Server 2007	5
Environment best practice verification	6
Transitioning to Microsoft Exchange Server 2007.....	6
Transitioning example	7
Legacy Exchange 2000 Server and Exchange Server 2003 configuration details	8
Microsoft Exchange 2000 Server	8
Microsoft Exchange Server 2003	9
Microsoft Exchange Server 2007 environment details	9
Microsoft Exchange Server 2007 installation.....	9
Preparing the Domain	9
Installing Exchange Server 2007.....	10
Moving the user Mailboxes	12
Move Mailbox Wizard	13
Exchange Management Shell Move Mailbox script.....	19
Exchange Server 2007 Move Mailbox performance	23
Decommissioning the legacy Exchange servers.....	24
Conclusion.....	24
Appendix A: Code sample PowerShell Move Mailbox script.....	25
Appendix B: Why 64-bit Microsoft Exchange Server?	29
For more information.....	30

Code notice

This document contains examples of code ("Code") that you may want to use in building or developing your own application. You may use this Code only as permitted herein. By your use of Code, you agree to the terms below. The Licensor, Hewlett-Packard Development Company, L.P. ("HPDC"), grants you the rights listed below.

You may use the Code either by referring to it when developing your own application, copying it in whole or in part into your application, or building upon a portion or all of it to create your own application based on it. While using the Code to build your own application, you may alter it, modify it, and create derivative works of the Code. You may also use the Code to test your application. You may distribute the Code provided that you comply with the conditions on distribution described below.

Conditions on Distribution

You may :

(1) reproduce and distribute an unlimited number of copies of the Code within your application, in source code form, internally within your organization, including subsidiaries and affiliates;

(2) you may reproduce and distribute an unlimited number of copies of the Code within your application, in source code form, externally provided that:

- (a) your application adds significant and primary functionality to the Code;
- (b) you distribute your application containing the Code under an End-User License Agreement, or in signed hard-copy form, with terms no less protective than those contained herein, but permitting your end users only internal distribution as described in item (1) above;
- (c) you do not use the HPDC name, trademarks, or logo; or the name, trademarks, or logo of Compaq Computer Corporation ("Compaq"); or the name, trademarks, or logo Hewlett-Packard Company ("HP") to market your application;
- (d) you include a valid copyright notice on your application; and
- (e) you agree to indemnify, hold harmless, and defend Compaq, HP, and HPDC from and against any claims or lawsuits, including attorneys' fees, that arise or result from the use or distribution of your application.

(3) reproduce and distribute an unlimited number of copies of the Code within your application, in binary form, internally within your organization, including subsidiaries and affiliates, or externally provided that:

- (a) you do not use the Compaq name, trademarks, or logo; or the name, trademarks or logo of Hewlett-Packard Company; or the name, trademarks, or logo of HPDC to market your application;
- (b) you include a valid copyright notice on your application; and
- (c) you agree to indemnify, hold harmless, and defend Compaq, HP, and HPDC from and against any claims or lawsuits, including attorneys' fees, that arise or result from the use or distribution of your application.

NO WARRANTY

THE CODE IN THIS DOCUMENT IS PROVIDED 'AS-IS', WITHOUT ANY EXPRESSED OR IMPLIED WARRANTY. IN NO EVENT WILL COMPAQ, HP, OR HPDC AND/OR THEIR SUBSIDIARIES OR AFFILIATES BE HELD LIABLE FOR ANY DAMAGES ARISING FROM THE USE OF THIS CODE. TO THE EXTENT PERMITTED BY LAW, COMPAQ AND HP AND HPDC, HEREBY DISCLAIM ALL WARRANTIES AND CONDITIONS WITH REGARD TO THE CODE, INCLUDING ALL IMPLIED WARRANTIES AND CONDITIONS OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE, AND NON-INFRINGEMENT. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE CODE IS WITH YOU. SHOULD THE CODE PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION. IN NO EVENT SHALL COMPAQ, HP, OR HPDC OR THEIR SUBSIDIARIES OR AFFILIATES BE LIABLE FOR ANY DIRECT, CONSEQUENTIAL, INCIDENTAL, SPECIAL, EXEMPLARY, OR PUNITIVE DAMAGES, OR DAMAGES FOR LOSS OF BUSINESS PROFITS, BUSINESS INTERRUPTION, OR LOSS OF BUSINESS INFORMATION, ARISING OUT OF THE USE OR INABILITY TO USE THE CODE (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE CODE TO OPERATE WITH ANY OTHER PROGRAMS), OR OTHER DAMAGES WHATSOEVER, EVEN IF COMPAQ, HP, AND/OR HPDC HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Reproduction, adaptation or translation without prior written permission is prohibited, except as allowed under the copyright laws.

Executive summary

This white paper provides an overview of the transition process associated with upgrading Microsoft® Exchange 2000 Server and Microsoft Exchange Server 2003 organizations to those using Microsoft Exchange Server 2007. Moving user mailboxes from Microsoft Exchange 2000 Server or Microsoft Exchange Server 2003 is a straight forward well-documented process. With proper planning and scheduling, the move-mailbox process can be accomplished with minimal disruption to the end-user community. Microsoft Exchange Server 2007 can co-exist with Microsoft Exchange 2000 Server servers and Microsoft Exchange Server 2003 servers allowing for interoperability between the Microsoft Exchange Server versions. This interoperability will allow for a phased deployment of Microsoft Exchange Server 2007 servers and moving of user mailboxes. User mailboxes can be moved using either the Exchange Management Console GUI or via a command line script using Windows® PowerShell. Both methods will be documented and described in this document. This document will also discuss many of the new features and benefits associated with Microsoft Exchange Server 2007.

Introduction

Microsoft defines migration as moving users from a foreign mail system such as Sendmail or Lotus Notes to Microsoft Exchange Server 2007. Transition is defined as moving users from Microsoft Exchange 2000 Server or Microsoft Exchange Server 2003 to Microsoft Exchange Server 2007. Exchange Server 2007 requires Windows Server 2003 x64 Edition, since Exchange 2000 Server or Exchange Server 2003 servers utilized the 32-bit version of Windows Server there is no in-place upgrade path to Exchange Server 2007.

Transitioning from a previous version of Exchange 2000 Server or Exchange Server 2003 requires the new Exchange Server 2007 servers be installed into the organization; user mailboxes can then be moved from the existing Exchange 2000 Server servers or Exchange Server 2003 servers to the new Exchange Server 2007 Mailbox servers.

Benefits of moving to Exchange Server 2007

Microsoft Exchange Server 2007 provides many new enhancements and features that can improve an organization's messaging environment. Enhancements in administration, performance, and availability can help companies reduce costs and leverage hardware more efficiently. Scalability and performance improvements associated with Exchange Server 2007 remove many of the limitations associated with previous versions of Microsoft Exchange servers, allowing optimal use of server CPU, memory, and disk subsystem resources. For more information on the benefits of Microsoft Exchange Server on Windows Server 2003 x64 Edition see *Appendix B: Why 64-bit Microsoft Exchange Server*. Exchange Server 2007 contains several availability options to help minimize server downtime; these options include Single Copy Cluster (SCC) common in Microsoft Exchange Server 2003 and new replication-based availability solutions called Local Cluster Replication and Continuous Cluster Copy. Improvements in message hygiene include spam filtering and anti-virus protection provided with the Microsoft Exchange Server 2007 Edge Transport server role. Several new technologies provide end users with ubiquitous access to e-mail. The technologies include:

- Integration between telephony systems, providing end users with unified voice and e-mail messaging via the Exchange Server 2007 Unified Messaging Server role.
- Outlook Anywhere provides Outlook client access via RPC over HTTP.
- Exchange Active Sync provides access to mobile devices.
- Outlook Web Access provides browser based access to user mailboxes.

Exchange Server 2007 provides corporations with a granular approach to message journaling that will help them comply with mandated message archiving and retention policies. The Hub Transport server role provides journaling options on a user level that will help organizations meet compliance requirements. Enhancements in administration are provided by Windows PowerShell and the Exchange Management Shell. Every administration task associated with Exchange Server 2007 can be accomplished through powerful and simple command scripts.

Preparing for a transition to Microsoft Exchange Server 2007

There are several preliminary requirements that must be met before installing the first Microsoft Exchange Server 2007 server into an existing Microsoft Exchange organization.

- Microsoft Exchange Server 2007 cannot co-exist with Microsoft Exchange Server 5.5 servers; the existing Microsoft Exchange Organization operating mode must be set to “Native Mode (no pre-Exchange 2000 servers)”.
- Microsoft Exchange 2000 Server servers must be running a minimum of Microsoft Exchange 2000 Server SP3
- Microsoft Exchange Server 2003 servers must be running a minimum of Microsoft Exchange Server 2003 SP2.

Active directory requirements that must be met prior to installing the first Microsoft Exchange Server 2007 server include the following:

- The Domain controller hosting the schema master role must be running Microsoft Windows Server 2003 Service Pack 1 (SP1).
- The Active Directory functional level must be set to Windows 2000 Server-native or higher.
- Active Directory sites that contain a Microsoft Exchange Server 2007 server must have an Active Directory global catalog server running Microsoft Windows Server 2003 Service Pack 1.

Administrators must determine if the Microsoft Exchange Server 2007 server environment will coexist with the existing Microsoft Exchange 2000 Server servers or Microsoft Exchange Server 2003 servers. Larger organizations that must perform the mailbox-move operation over an extended period of time; overnight or over a weekend for example; will be required to operate in a coexistence mode until all the users are transitioned to Microsoft Exchange Server 2007. Organizations may also be required to operate in coexistence mode due to legacy services that are no longer supported by Microsoft Exchange Server 2007.

The following is a list of services provide by previous versions of Microsoft Exchange that are not available in Exchange Server 2007.

- Instant Messaging Service
- Lotus Notes Connector
- Key Management Service
- Microsoft Chat service
- Mobile Information server
- Instant Messaging Service
- Key Management Service
- Exchange Chat Service
- Exchange 2000 Conference Server

- Cc:Mail Connector
- Microsoft Mail Connector
- GroupWise connector
- X.400 connector

If any of the above services are required and are currently provided by Exchange 2000 Server or Exchange Server 2003 the organization must operate in a coexistence mode until the service can be retired.

Environment best practice verification

Use the Microsoft Exchange Best Practice Analyzer (EXBPA) to verify and resolve any issues that may arise before installing the first Exchange Server 2007 server. Below is an excerpt from the EXBPA report.

Software update available	Server: EXCHANGE2003
Server EXCHANGE2003 is running Exchange Server 2003 (original release). Service Pack 2 is available for download from http://www.microsoft.com/exchange .	
Tell me more about this issue and how to resolve it.	

This output informs us that that the server “EXCHANGE2003” does not currently have any Exchange service packs applied and, as discussed earlier, the Microsoft Exchange Server 2003 server must be at a minimum of Microsoft Exchange Server 2003 SP2.

Transitioning to Microsoft Exchange Server 2007

Microsoft Exchange Server 2007 server deployments are comprised of several different roles. The following server roles can exist in a Microsoft Exchange Server 2007 deployment:

- Client Access Server (CAS)
- Hub Transport (HUB)
- Edge Transport (Edge)
- Mailbox
- Unified Messaging (UM)

These roles can coexist on a single server or be installed onto dedicated servers. There are a couple of exceptions to this rule:

- Edge server role must be installed onto a dedicated server.
- If the mailbox server is installed into a cluster for Continuous Cluster Replication (CCR) or a Single Copy Cluster (SCC), the mailbox server role cannot be collocated with any other server roles.

Sizing and performance considerations will factor into the planning and deployment of Microsoft Exchange Server 2007 server roles.

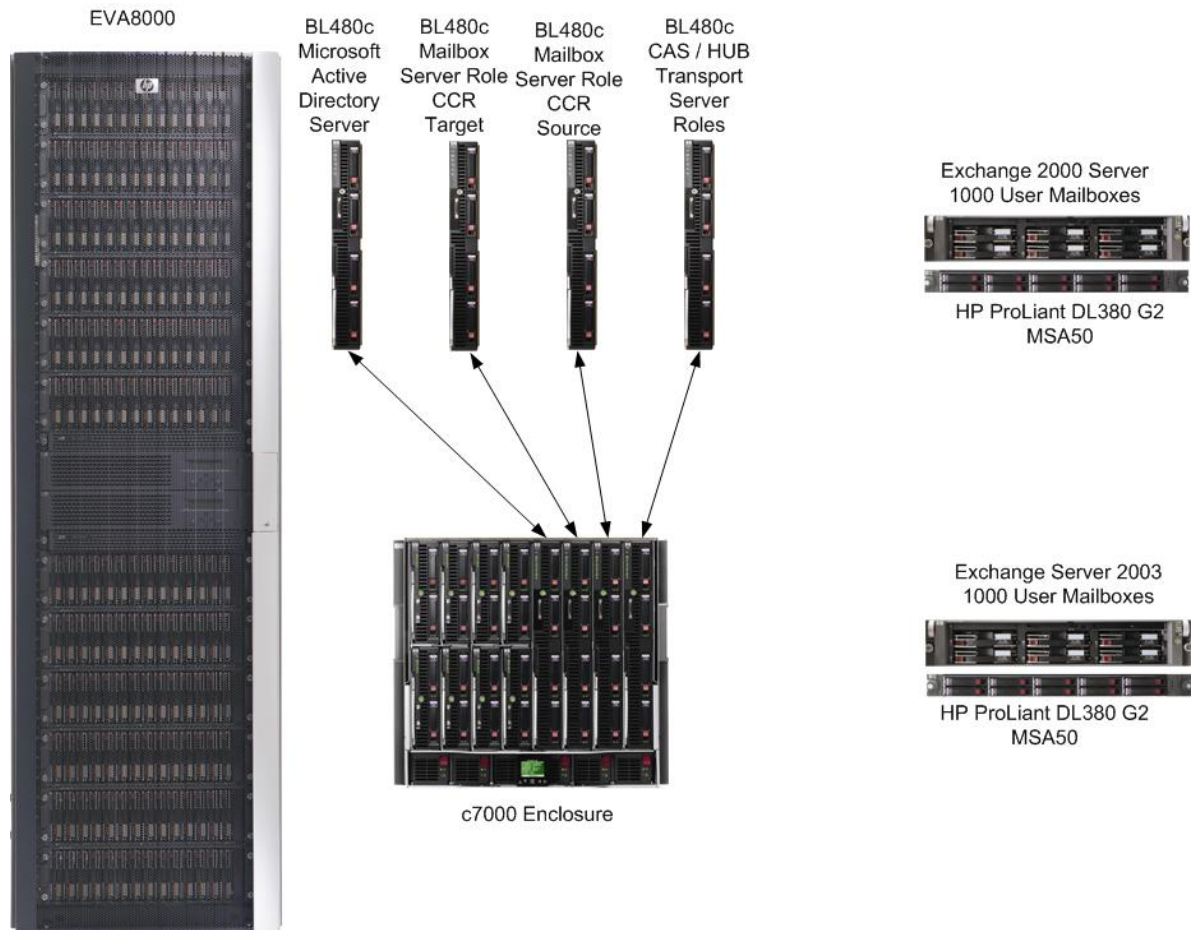
When installing Exchange Server 2007 server roles, the Client Access Server role is the first server role to be installed, followed by the Hub Transport server role, and finally the Mailbox server role. The Edge Transport and Unified Messaging server roles can be installed at any time after the CAS, HUB, and Mailbox server roles.

Transitioning example

The example described in this document will move 2000 Microsoft Exchange user mailboxes from Microsoft Exchange 2000 Server and Microsoft Exchange Server 2003. In the legacy Exchange organization there is a coexistence of two Microsoft Exchange servers: one Microsoft Exchange 2000 Server, and one Microsoft Exchange Server 2003 server. Each server currently hosts 1000 mailboxes with an average mailbox size of 100 megabytes per user. Both legacy Microsoft Exchange servers are currently deployed on rack mounted HP ProLiant DL380 G2 servers. The Exchange 2000 Server and Exchange Server 2003 databases are located on HP StorageWorks 50 Modular Smart Array (MSA50) Serial Attached SCSI (SAS) direct-attached storage enclosures. Each HP ProLiant DL380 G2 server is connected to one MSA50 disk enclosure.

The Microsoft Exchange Server 2007 environment is installed on the HP BladeSystem c-Class architecture as shown in Figure 1. The Exchange environment consists of four HP ProLiant BL480c servers. One ProLiant BL480c server hosts the Client Access Server (CAS) and Hub Transport server roles. The second ProLiant BL480c server will serve as the Exchange Server 2007 Mailbox server role for the 2000 mailboxes. The Exchange Server 2007 Mailbox server running on the BL480c server is the active server in a Microsoft Exchange Server 2007 Continuous Cluster Replication (CCR) cluster. The third ProLiant BL480c server is also configured as an Exchange Server 2007 Mailbox server and will be the passive node for the CCR cluster. The fourth ProLiant BL480c server is the Microsoft Active Directory global catalog server for the domain. An HP StorageWorks 8000 Enterprise Virtual Array (EVA8000) SAN will satisfy the storage requirements for the Exchange Server 2007 environment, including the CCR source and target database files.

Figure 1. Exchange Server 2007 sample deployment



Legacy Exchange 2000 Server and Exchange Server 2003 configuration details

In the legacy Microsoft Exchange environment, each user has a mailbox size of 100 megabytes and each server has one storage group and two databases. The users are equally divided between the servers, storage groups and databases.

Microsoft Exchange 2000 Server

SG1

- priv1.edb
 - Users = 500
 - MBX size = 100 megabytes / users
 - EDB size = 50 gigabytes
- priv2.edb
 - Users = 500
 - MBX size = 100 megabytes / users
 - EDB = 50 gigabytes

Microsoft Exchange Server 2003

SG1

- priv1.edb
 - Users = 500
 - MBX size = 100 megabytes / users
 - EDB = 50 gigabytes
- priv2.edb
 - Users = 500
 - MBX size = 100 megabytes / users
 - EDB = 50 gigabytes

Microsoft Exchange Server 2007 environment details

As part of the design for the Exchange Server 2007 environment, user mailbox quotas will be raised from 100 MB to 1GB per user. Because of this increase in mailbox size, the size of the databases will increase substantially. If the same storage group configuration were maintained, then each database would be 500GB and much too large to backup and restore within a reasonable time frame to meet many companies' service level agreements. In order to reduce the size of the databases, the Exchange Server 2007 design utilized 20 storage groups with one database per storage group (one database per storage group is a design requirement of CCR).

Each storage group will host a single 100 gigabyte information store database supporting 100 user mailboxes.

- SGX
 - PrivX.edb
 - Users = 100
 - MBX size = 1GB / user
 - EDB size = 100 GB
 - LUN size = 150 GB

The design could have used additional databases to further reduce the size of the databases; Exchange Server 2007 can support up to 50 storage groups. However, by using CCR, the databases are kept in sync on the passive node with log shipping. This means that the passive node can serve as the backup target. By using the passive node as the backup target, the potential overlap of the maintenance and backup windows that might occur on the active node is eliminated, which allows for larger backup windows and subsequently larger database sizes.

Microsoft Exchange Server 2007 installation

Preparing the Domain

The Microsoft Exchange Server 2007 setup.exe program provides switches that prepare Active Directory for the new Exchange Server 2007 installation. A partial list of the Setup switch options are displayed below:

- Setup /PrepareAD, /p – Prepares the Active Directory forest for the Exchange installation
- Setup /PrepareDomain, /pd – Prepares the local domain for the Exchange installation
- Setup /PrepareSchema, /ps – Prepares the Active Directory Schema for the Exchange installation

- Setup /PrepareLegacyExchangePermissions, /pl – Prepares the permissions in a legacy forest. Use this parameter only if you have servers that are running Exchange 2000 Server or Exchange Server 2003 in your organization.

Installing Exchange Server 2007

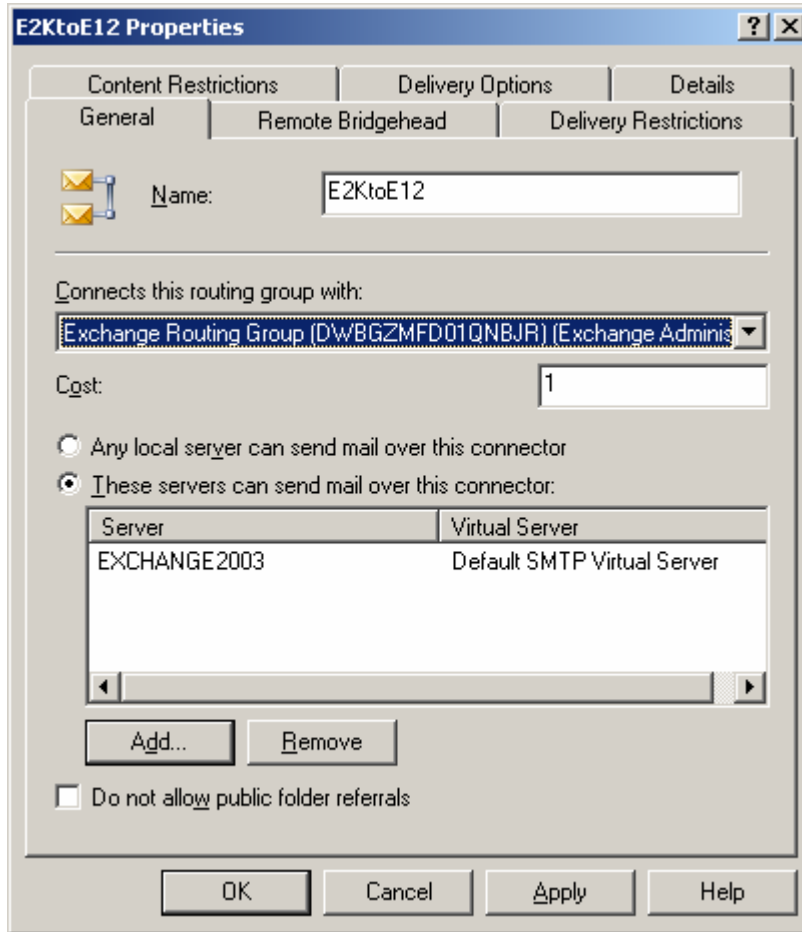
The first server installed was the CAS/HUB server. For this example, there are no Edge or Unified Messaging server roles. A majority node set cluster was then configured on the second and third servers that will host the mailbox server roles. Next, the Mailbox server role was installed as the destination for the move mailbox operation; this server will act as the source for the CCR cluster. Once the Mailbox server was installed on the second server, a CCR cluster was created and configured. The second Mailbox server will be the target server for the CCR cluster. The CCR cluster target Mailbox server will not be installed until after the move mailbox operation is completed.

Note: Documenting the installation and configuration of a CCR cluster is beyond the scope of this white paper. For more information on installation and configuration of CCR clusters for Exchange Server 2007 mailbox server, please refer to Microsoft's documentation.

In a pure Microsoft Exchange Server 2007 environment, there is no concept of either an administrative or routing group. However, once the first Exchange Server 2007 server is installed, a default Exchange Routing Group with a randomly generated ID suffix that contains the Microsoft Exchange Server 2007 Hub Transport servers is created to support mail routing between legacy Exchange 2000 Server and Exchange Server 2003 routing groups. This routing group is only visible through the Exchange Server 2003 Systems Manager interface. Unlike Exchange 2000 Server and Exchange Server 2003, which use routing groups to determine message routing paths, Exchange Server 2007 uses Microsoft Active Directory site topology to determine message routing paths.

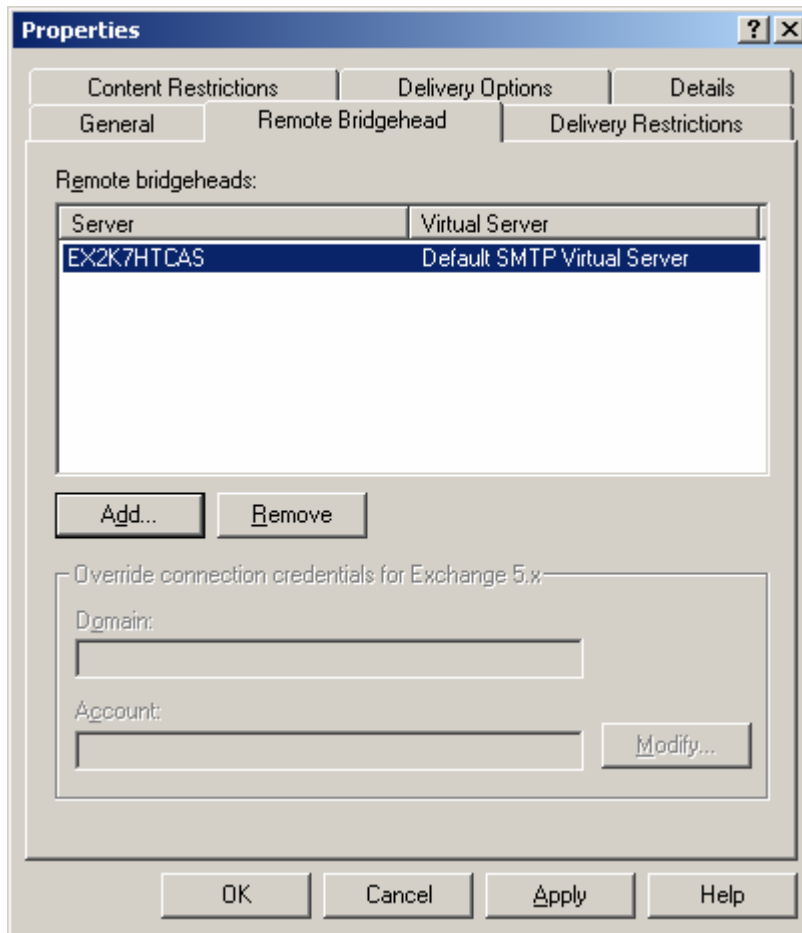
If you are transitioning a small environment and there is no need for coexistence, a routing group connector is not required to move the mailboxes from Exchange 2000 Server or Exchange Server 2003 to Exchange Server 2007. However, if Exchange Server 2007 will coexist with previous versions of Exchange 2000 Server or Exchange Server 2003, a routing group connector must be created between the existing legacy Exchange routing groups and the default routing group created for Exchange Server 2007 servers to enable mail routing. Figure 2 illustrates a routing group connector created with Microsoft Exchange System Manager on a Microsoft Exchange Server 2003 server between the legacy Exchange Server 2003 server routing group and the default Exchange Server 2007 routing group "Exchange Routing Group (DWBGZMFD01QNBJR)".

Figure 2. Routing Group Connector displays the default Exchange Server 2007 routing group



Notice in the Remote Bridgehead property page in Figure 3, the name of the server specified is EX2K7HTCAS; this server hosts the Hub Transport and Client Access server roles for the new Microsoft Exchange Server 2007 environment.

Figure 3. Remote Bridgehead property displays the Exchange Server 2007 Hub Transport server role




Moving the user Mailboxes

Once the Microsoft Exchange Server 2007 CAS, HUB, and Mailbox server roles have been installed, one can begin moving users from the legacy Microsoft Exchange Server environments to Microsoft Exchange Server 2007. The method used to move users for this intra-org transition is the move mailbox operation. The move mailbox operation can be executed using either the Exchange Management Console (EMC) or via the command line using the Exchange Management Shell (EMS), an extension of Windows PowerShell. This white paper will outline using both techniques to run the move mailbox operation.

Move Mailbox Wizard

In the example documented in this paper, the first 100 users, storage group 1, are moved using the Exchange Management Console and the remaining 1900 users using the EMS. In Figure 4, 100 recipient mailboxes have been selected under the Recipient Configuration section of the EMC.

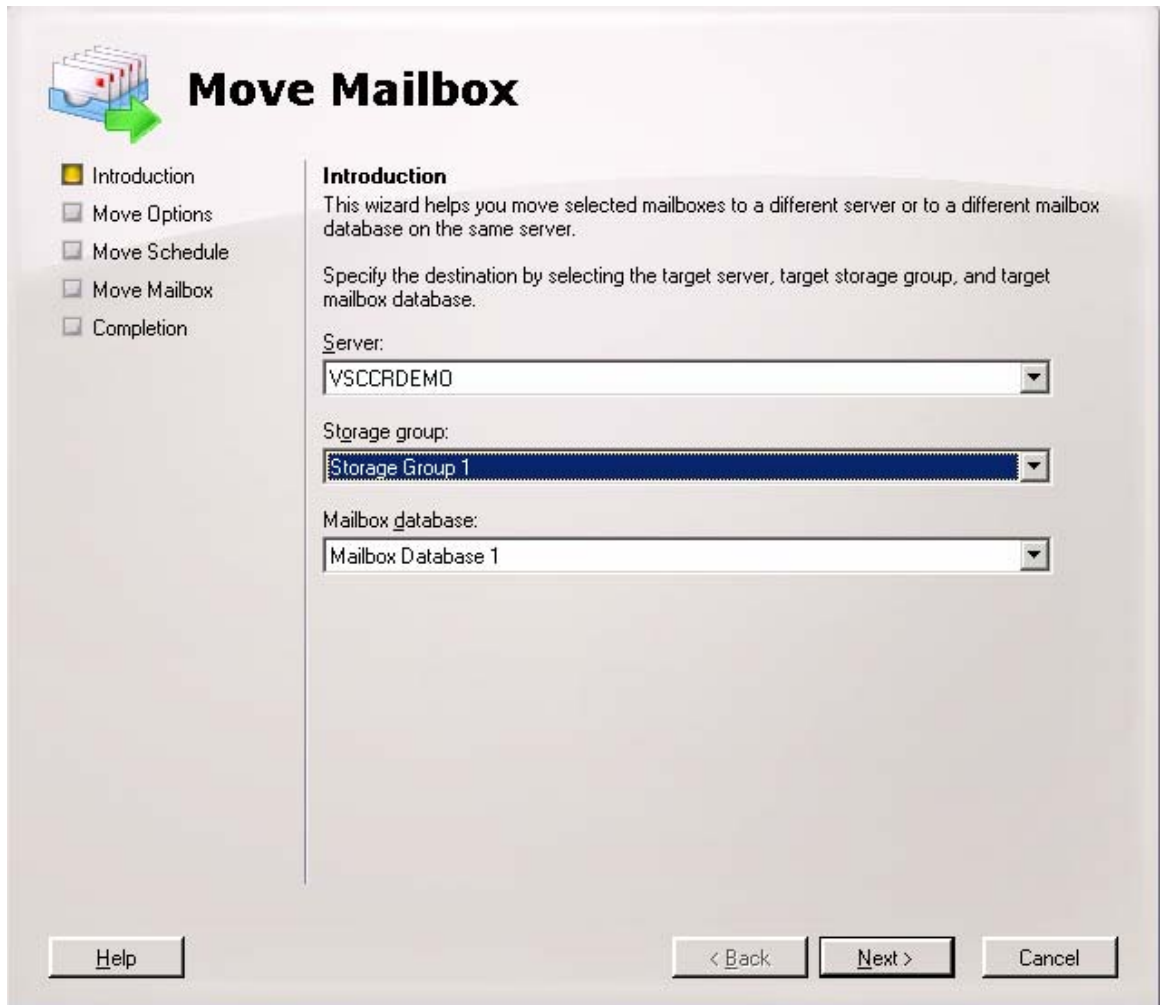
Figure 4. Microsoft Exchange Server 2007 EMC Recipient Configuration view



In the Mailbox display pane you can see that Recipient Type Details indicates Legacy Mailbox for the selected users. In the Actions pane, select Move to launch the Move Mailbox wizard shown in Figure 5.

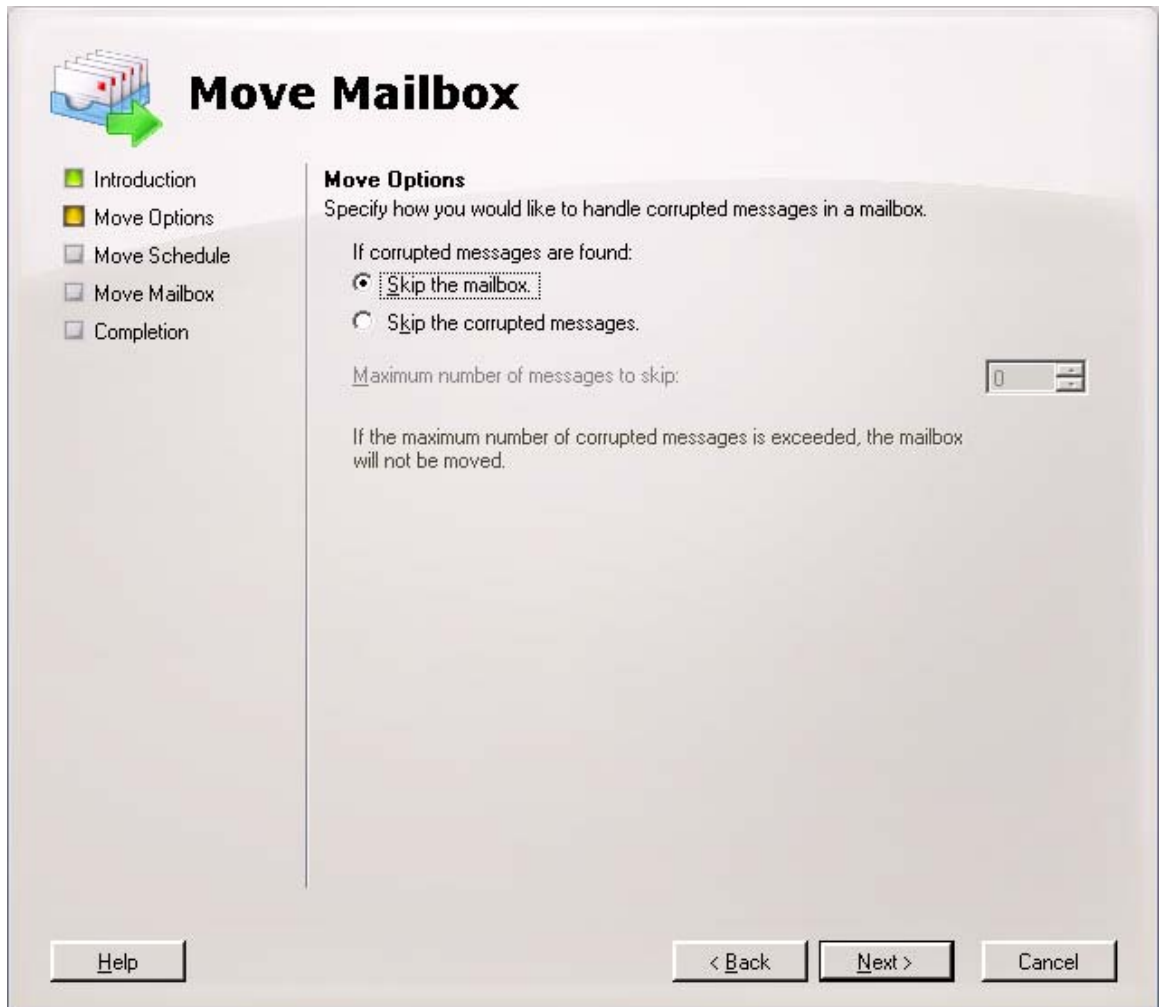
Selecting **Move** in the Actions pane launches the Move Mailbox Wizard. The Introduction screen shown in Figure 5 prompts the administrator to specify a target Mailbox Server, target Storage group, and a target Mailbox database as a destination for the selected users.

Figure 5. Move Mailbox Introduction Screen



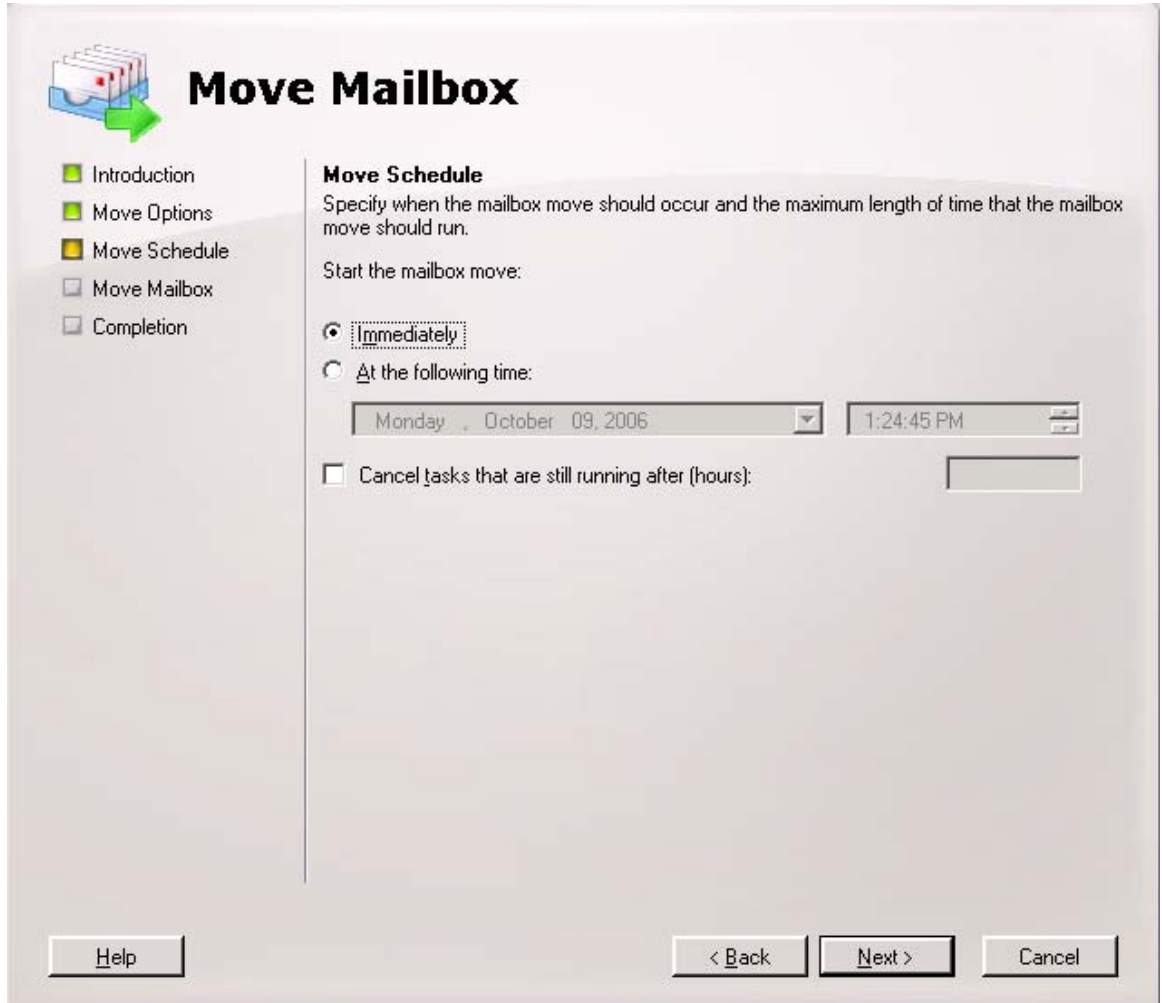
The second screen in the Move Mailbox wizard is the Move Options screen shown in Figure 6. This screen provides the administrator with several options for moving the mailbox. The Move Options page allows users to instruct the Move Mailbox wizard how to handle corrupt messages while moving legacy mailboxes to Exchange Server 2007. The application can either skip mailboxes with corrupt messages or skip corrupt messages until a user defined threshold titled “Maximum number of messages to skip” is exceeded.

Figure 6. Move Mailbox Wizard Move Options



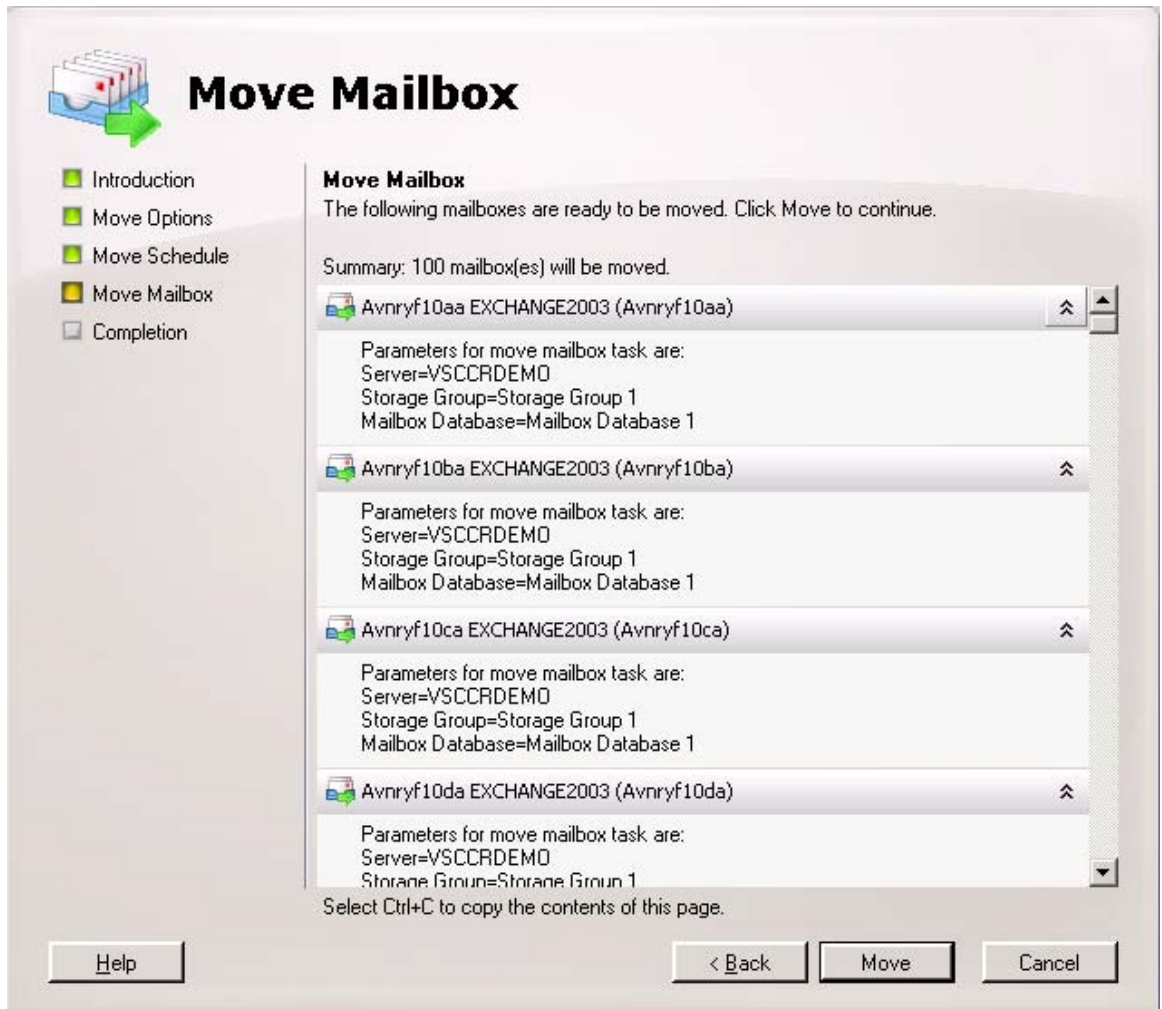
The third screen in the Move Mailbox Wizard, illustrated in Figure 7, is the Move Schedule screen; this option allows the end user to schedule the move and define how long the move process should run. An administrator can choose to start the move immediately or schedule a Move Mailbox operation for a certain time. For example an administrator could schedule a move for Friday at 5:00 PM. The Move Schedule also allows the administrator to determine how long the move mailbox operation should run. In the previous example, if the administrator scheduled the run to begin at 5:00 PM on Friday the move mailbox process could be configured to cancel all tasks after 48 hours, ending the move mailbox process on Sunday at 5:00 PM.

Figure 7. Move Mailbox Wizard Move Schedule



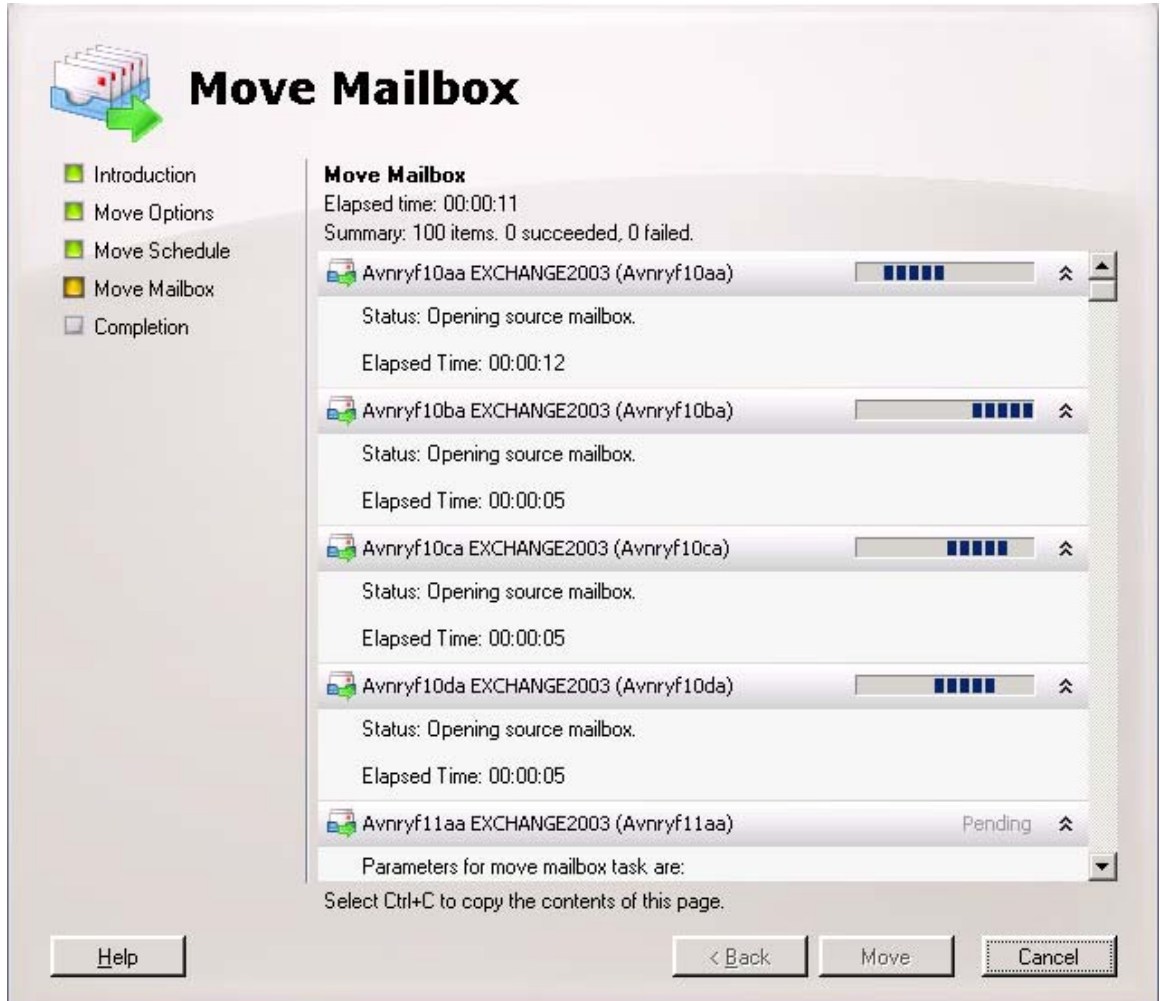
The fourth screen is the Move Mailbox screen shown in Figure 8; this screen provides a summary of the user mailboxes selected to be moved. Once the “**Move**” button is selected, the mailboxes will be moved from the source server to the target server.

Figure 8. Move Mailbox Wizard Move Mailbox Confirmation



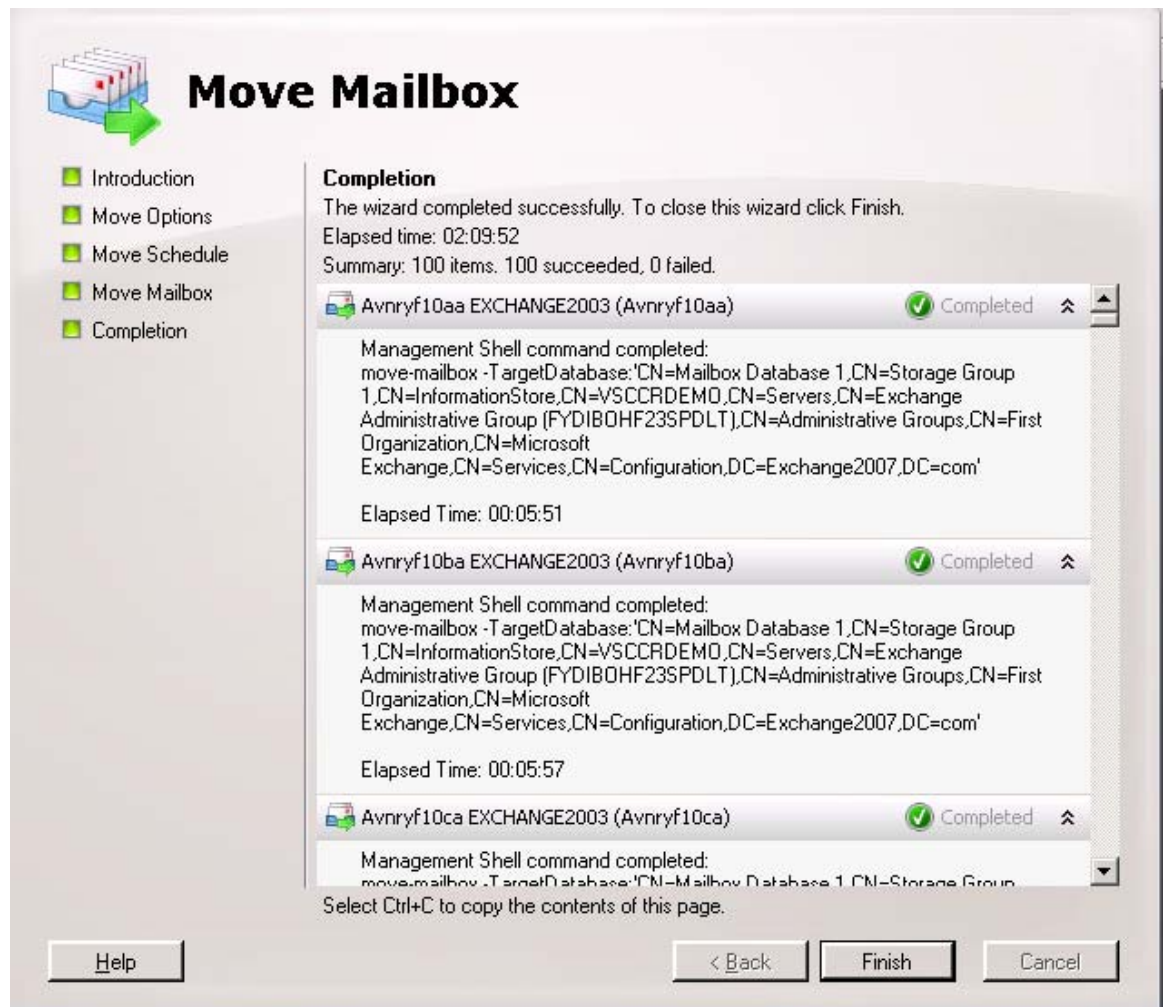
During the move mailbox process, the Move Mailbox screen displays the current status of the move operation, illustrating the number of items processed and the elapsed time for each mailbox. The move mailbox application is capable of processing four mailboxes concurrently. In Figure 9, you can see there are four mailboxes in the process of being moved and the fifth is in a pending state.

Figure 9. Move Mailbox Wizard Move Mailbox Status



Once the move mailbox process is complete, the Completion screen is displayed as shown in Figure 10. The Completion screen displays the elapsed time for the total operation, elapsed time for each mailbox, and a summary of the number of items successfully moved and the number of failed items.

Figure 10. Move Mailbox Wizard Completion



During the move mailbox operation, 100 mailboxes were moved in 2 hours 9 minutes and 52 seconds. Each mailbox was approximately 100 megabytes totaling 10 gigabytes of data. This results in a rate of 1.3 minutes per mailbox or 4.6 gigabytes per hour. The rate at which mailboxes are moved may vary widely dependant upon the network topology, system configuration, and I/O subsystem of the source and destination servers.

Exchange Management Shell Move Mailbox script

With Microsoft Exchange Server 2007, all management tasks can be performed via the Exchange Management Shell (EMS), which is an extension of Windows PowerShell. The Exchange Management Console (EMC) GUI is actually built on top of the EMS so that any command executed inside the GUI can also be scripted on the command line. In addition, to help lower the learning curve for administrators, any tasks performed using the EMC will provide a reference to the applicable management shell command completed. This output can be used as a basis for developing a PowerShell script that can be used to move users.

In Figure 10, the Move Mailbox Completion screen shows the Management Shell command completed for each mailbox moved during the move mailbox operation for the 100 mailboxes. Based on this output, a script was developed to move the remaining 1900 mailboxes from the Exchange 2000 Server and Exchange Server 2003 servers to the Exchange Server 2007 server.

The script was developed to provide two modes of operation, moving users to a single storage group or moving users to multiple storage groups. For the multi-storage group mode, the script assumes that only one database exists per storage group. For this example, the 1900 users will be migrated to 19 storage groups and the second mode of operation was executed.

When choosing the second mode, the script has several required parameters. There are three text file inputs to the script. The first file contains a list of other txt files.

For example,

C:\test\sg1names.txt

C:\test\sg2names.txt

.
. .
.

C:\test\sgNnames.txt

Each of the files listed in the first text file contains the recipients' e-mail addressees for a single storage group. So "C:\test\sg1names.txt" would contain a list of 100 e-mail addresses referencing the 100 users that will be migrated to storage group 1. The second and third file inputs contain the list of storage groups and databases respectively that will host the mailboxes on the Exchange Server 2007 server.

To generate the recipient email address text files, the CSVDE utility was used to export Active Directory user information. The CSV output was then imported into Microsoft Excel for manipulation and sorting. This allows an administrator performing the move operation to define how and where the user mailboxes will be moved. Using CSVDE, an administrator can export and sort user data based upon user attributes stored in Active Directory. For example, users can be sorted and grouped by organizational unit, current Microsoft Exchange server, mail stores, or any other Active Directory user attribute.

The script then reads in each of the files and generates the applicable command to execute within EMS. The main function call in the script is the move-mailbox operation. The script excerpt below shows a piece of code used to generate the move-mailbox operation. The full EMS script used to move user mailboxes for this example is provided in Appendix A.

```
$targetdatabase = get-mailboxdatabase -identity "$mbxName\$SGName\$DBName"
foreach($User in $FileInput)
{
    get-mailbox -identity $user | move-mailbox -TargetDatabase:$targetdatabase -
confirm:$False -maxthreads:4
}
```

The \$ refers to variables used in the script. The \$targetdatabase uses the get-mailboxdatabase function and the input files to get the database object for the current batch of users. Then for each user that will be moved to the current target database, the mailbox object is retrieved using the get-mailbox command. The user mailbox object is then piped as input to the move-mailbox command that

takes the target database object as a parameter. The `--maxthreads:4` parameter was also used to speed up the move mailbox process. This process is nested within several iterations that will traverse through all of the users, storage groups, and databases specified in the input files.

The output of the script is shown in Figure 11.

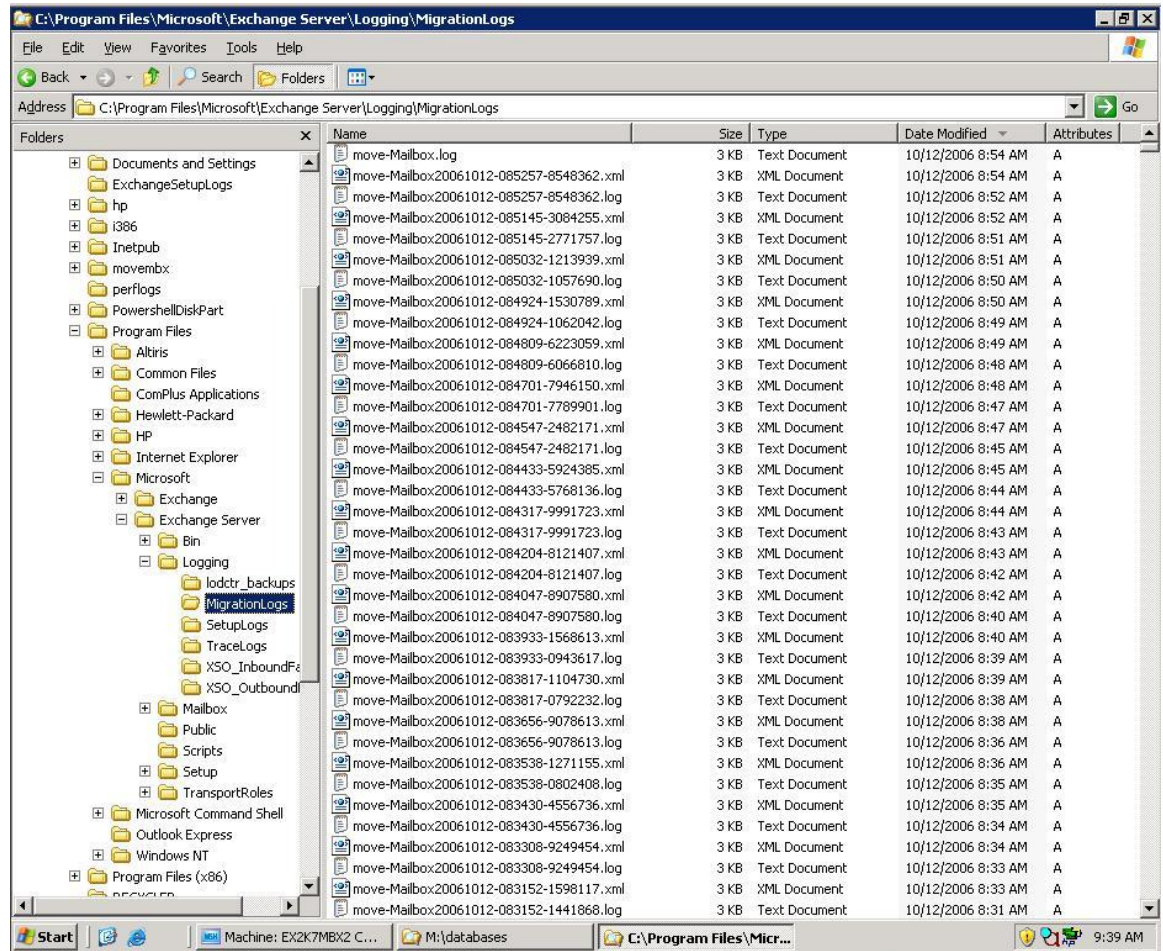
Figure 11. Power Shell Move Mailbox script output

```
Machine: EX2K7MBX2 CWD: C:\movembx
DistinguishedName           : cra3p2aa EXCHANGE2K
                             : CN=Eocra3p2aa EXCHANGE2K,OU=EXCHANGE2K,OU=Lo
Eocra3p2ba
Moving messages. Inbox (68/250)
oooooooooooooooooooooooooooo 1
PrimarySmtpAddress          : Eocra3p2aa@Exchange2007.com
SourceServer                 : exchange2k.Exchange2007.com
SourceDatabase               : priv2
SourceGlobalCatalog         : ad.Exchange2007.com
TargetGlobalCatalog         : ad.Exchange2007.com
TargetDomainController      : ad.Exchange2007.com
TargetMailbox                :
TargetServer                 : USCCRDEMO.Exchange2007.com
TargetDatabase               : USCCRDEMO\Storage Group 8\Mailbox Database8
MailboxSize                  : 102142KB
IsResourceMailbox           : False
SIDUsedInMatch              :
SMTPProxies                  :
SourceManager                :
SourceDirectReports          :
SourcePublicDelegates        :
SourcePublicDelegatesBL     :
MatchedTargetNTAccountDN    :
IsMatchedNTAccountMailboxEnabled :
MatchedContactsDNList      :
TargetNTAccountDNToCreate   :
TargetManager                :
TargetDirectReports          :
TargetPublicDelegates        :
TargetPublicDelegatesBL     :
Options                      : Default
SourceForestCredential       :
TargetForestCredential       :
TargetFolder                 :
RsgMailboxGuid               :
RsgMailboxLegacyExchangeDN  :
RsgMailboxDisplayName        :
RsgDatabaseGuid              :
MoveType                     : IntraOrg
MoveStage                    : Completed
StartTime                    : 10/10/2006 9:17:35 AM
EndTime                      : 10/10/2006 9:19:00 AM
StatusCode                   : 0
StatusMessage                 : This mailbox has been moved to the target da
                             : tabase.
```

Setting the thread value to 4 in the Exchange Management Shell script resulted in the mailbox move operation completing a move in 1.4 minutes per mailbox. This performance is similar to the Move mailbox operation in the Exchange Management Console which operates on four mailboxes concurrently during the move operation.

Upon completion of the move process by either the Move Mailbox wizard or the Windows PowerShell script, a log file is generated. The Move Mailbox wizard creates a single log file for all the selected users moved, while the script creates a log file for each user. By default the logs are stored in C:\Program Files\Microsoft\Exchange Server\Logging\Migration Logs, as shown in Figure 12.

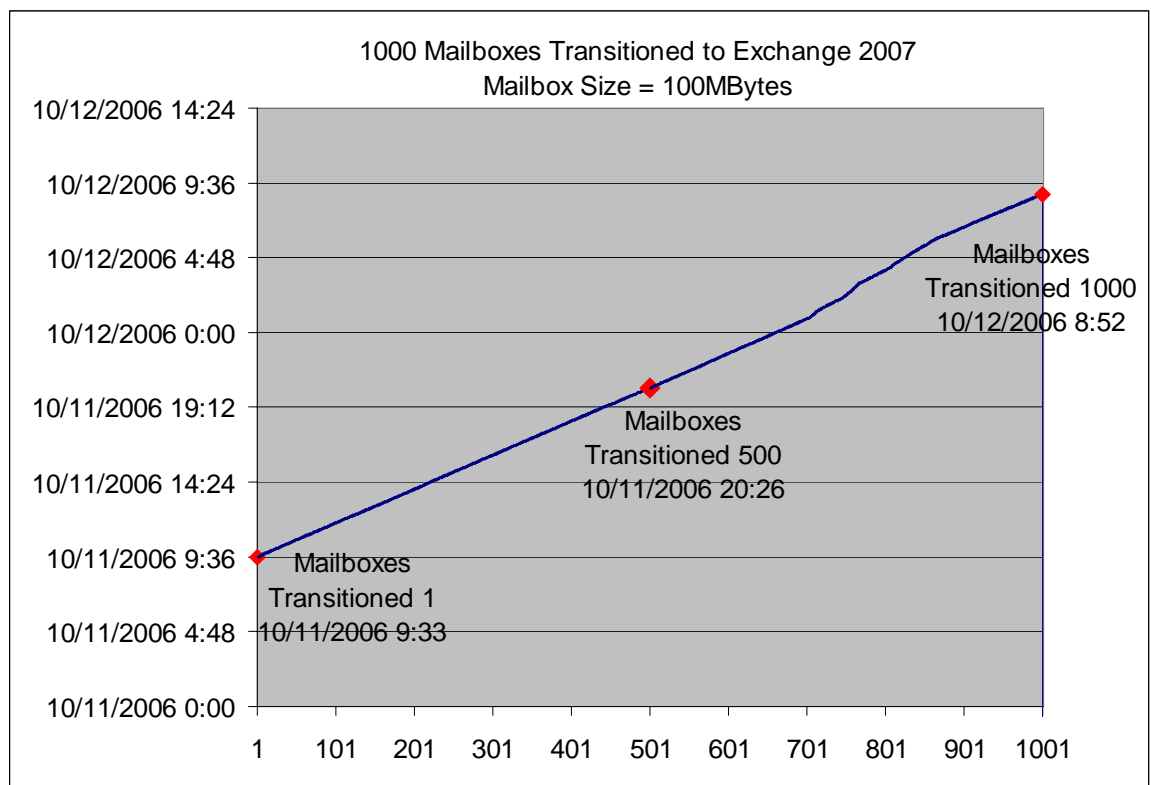
Figure 12. Move Mailbox logs



Exchange Server 2007 Move Mailbox performance

The chart in Figure 13 illustrates the time to move 1000 mailboxes, half of the mailboxes in the organization, 500 from Exchange 2000 Server and 500 from Exchange Server 2003 to a single Exchange Server 2007 Mailbox server. Beginning at 9:33 AM on 10/11/2006 until 20:26 on 10/11/2006, five hundred mailboxes were moved from an Exchange 2000 Server to the Exchange Server 2007 server at a rate of 0.77 mailboxes per minute. From 20:26 on 10/11/2006 until 8:52 AM on 10/12/2006 five hundred mailboxes were moved from an Exchange Server 2003 server to the Exchange Server 2007 server at a rate of 0.67 mailboxes per minute. The total transfer time to move 1000 mailboxes with an average size of 100 megabytes resulted in 23 hours and 19 minutes at a rate of 1.4 minutes per mailbox or 0.72 mailboxes per minute. Each mailbox was approximately 100 megabytes requiring 100 gigabytes of data to be transferred at a rate of 4.29 gigabytes per hour.

Figure 13. Move Mailbox Performance 1000 users



While the focus of this paper is not on the performance of moving user's mailboxes, it is important to understand the impact on the existing Microsoft Exchange servers and the new Microsoft Exchange Server 2007 servers while mailboxes are being moved to the new environment. This data can prove valuable when planning a transition to ensure the move operation doesn't impact performance during normal working hours.

The new Microsoft Exchange Server 2007 mailbox server, hosted on a ProLiant BL480c server blade configured with 3.2 GHz Dual-processor Dual-Core Intel® Xeon® with Hyper-Threading enabled, averaged 21.6% CPU utilization during the move mailbox process. CPU utilization was also captured on the Exchange Server 2003 server: a single-processor 1.4 GHz Intel Pentium® III. This server averaged 17% CPU utilization while mailboxes were moved from Exchange Server 2003 to the

Exchange Server 2007 server. The CPU utilization on the destination Microsoft Exchange Server 2007 mailbox server can be attributed to the content indexing process. The content indexing process requires extensive CPU resources and can impact user performance. The content indexing process can be disabled during the move mailbox process and re-enabled during off peak hours.

Decommissioning the legacy Exchange servers

In the example used for this paper, there was a Microsoft Exchange 2000 Server and a Microsoft Exchange Server 2003 server. These servers can be decommissioned once all the users have been successfully moved to the Microsoft Exchange Server 2007 server. If operating in coexistence mode and all the legacy Microsoft Exchange Server resources have been transitioned to Microsoft Exchange Server 2007, any routing group connectors can be deleted. To decommission the legacy Exchange servers, use the associated Microsoft Exchange Server setup program to uninstall the Microsoft Exchange Server.

In the example documented in this white paper, the Exchange 2000 Server was assigned the role of Routing Group Master. This role had to be assigned to the Exchange Server 2003 server before the Exchange 2000 Server could be uninstalled. The last server that remained was the Exchange Server 2003 server. This server hosts the Recipient Update Service (RUS). The remove server option in the Microsoft Exchange Server 2003 setup utility could not be assigned to this server until the RUS service was assigned to a different server in the domain. Using the Microsoft Exchange Server 2003 Exchange System Manager application, the RUS service was assigned to the Exchange Server 2007 mailbox server; this allowed for the removal of the last legacy Exchange Server using the Exchange Server 2003 server setup utility. At the time of this writing, a procedure for assigning the RUS to another server was not documented. Microsoft Exchange Server 2007 does not use RUS for provisioning recipients and this role is not required in a pure Exchange Server 2007 environment.

Conclusion

Microsoft Exchange Server 2007 provides a feature-rich messaging environment designed to simplify administration, increase performance and scalability, optimize the use of server resources, and provide end users with ubiquitous access to their mail and messaging services.

Microsoft Exchange Server 2007 can co-exist with Microsoft Exchange 2000 Server and Microsoft Exchange Server 2003 server during the transition process. The same tools and procedures are used to move either Exchange 2000 Server or Exchange Server 2003 mailboxes. Mailboxes can be moved using either the Microsoft Exchange Server 2007 Exchange Management Console or via a scripted process using the Microsoft Exchange Management Shell. The speed at which mailboxes are moved from Exchange 2000 Server and Exchange Server 2003 to Exchange Server 2007 is comparable when using either the Move Mailbox wizard in the EMC or using the EMS with the thread setting of 4. The move mailbox process for 1000 mailboxes and 100 gigabytes of data was accomplished in approximately 23 hours and 19 minutes at a rate of 0.72 mailboxes per minute and 4.29 gigabytes per hour. During the move mailbox process, the Microsoft Exchange Server 2007 Mailbox server experienced a moderate and steady CPU load of 21.6% on a dual-processor, dual-core 3.2 GHz ProLiant BL480c server blade.

Once all legacy Microsoft Exchange 2000 Server and Microsoft Exchange Server 2003 services have been decommissioned, the legacy Microsoft Exchange Servers can be retired using the associated setup and remove utility for each version.

The transition from Microsoft Exchange 2000 Server or Microsoft Exchange Server 2003 to Microsoft Exchange Server 2007 is a straightforward process that can be accomplished using the tools and documentation provided with Microsoft Exchange Server 2007.

Appendix A: Code sample PowerShell Move Mailbox script

```
#####  
##  
## This script is designed to perform the move mailbox operation  
## per storage group.  
##  
#####  
  
##Read initial parameter to determine whether this is a single storage group or multiple storage  
execution  
param (  
    $inputType = (read-host "Move Mailbox automation script.`n`nPress 1 to enter in the  
parameters to move mailboxes for a single database on the Exchange Server 2007 mailbox  
server.`nPress 2 to enter in the parameters to move mailboxes for multiple storage groups on the same  
Exchange Server 2007 mailbox server.`n[1] or [2]" )  
)  
  
##Function to gather command line variables  
function commandline()  
{  
  
    $Script:TxtFile = read-host ("Please enter the location of the txt file to import with the email  
addresses of the users whose mailboxes will be moved.")  
    $Script:mbxName = read-host ("Please enter the name of the Exchange Server 2007  
mailbox server where the users mailboxes will be moved.")  
    $Script:SGName = read-host ("Please enter the name of the storage group on the Exchange  
Server 2007 server.")  
    $Script:DBName = read-host ("Please enter the name of the mailbox database on the  
Exchange Server 2007 server.")  
    # $Script:AdminGroup = read-host ("Please enter the name of the Administrative Group for  
the Exchange Server 2007 server.")  
    # $Script:Domain = read-host ("Please enter the Exchange Server 2007 domain name. (Leave  
out the .com, .net, etc suffix)")  
}  
  
##function to gather command line variables for multiple storage group  
function commandlineMultiple()  
{  
  
    write-host ("The following parameters will remain consistent for each storage group move  
mailbox operation.`n")  
    $Script:mbxName = read-host ("Please enter the name of the Exchange Server 2007  
mailbox server where the users mailboxes will be moved.")  
    # $Script:AdminGroup = read-host ("Please enter the name of the Administrative Group for  
the Exchange Server 2007 server.")  
    # $Script:Domain = read-host ("Please enter the Exchange Server 2007 domain name. (Leave  
out the .com, .net, etc suffix)")  
  
    write-host ("The following parameters refer to files created by the administrator that indicate  
the specific variables for each storage group move mailbox operation. In order to script multiple  
storage group operations, three text files are required.`n`n1) The first text file contains a list of text
```

files that correspond to the users' email addresses and mailboxes for each storage group. 2) The second text file contains the list of Exchange Server 2007 storage groups where users' mailboxes will be moved to. 3) The third text file contains the list of databases corresponding to each storage group. This script currently only supports a 1:1 storage group to mailbox database ratio. If more than 1 database exists per storage group, then the first option in this automation script can be run to execute the move mailbox operation on a per database level. One important thing to note is that each line in the text files must correspond to the same line in the other text files.

```
$Script:TxtFile = read-host ("Please enter the location of the txt file that contains the list of txt files with the email addresses of the users whose mailboxes will be moved.")
$Script:SGFile = read-host ("Please enter the location of the txt file that contains the names of the storage groups on the Exchange Server 2007 server.")
$Script:DBFile = read-host ("Please enter the location of the txt file that contains the names of the mailbox database on the Exchange Server 2007 server.")
}
```

```
##Function to call the move-mailbox cmdlet for the multiple storage group operation
function movembxMultiple()
{
```

```
    ##Remove quotes around file name
    if($TxtFile[0] -eq "")
    {
        [int]$count = $TxtFile.get_length()
        $TxtFile = $TxtFile.remove($count-1)
        $TxtFile = $TxtFile.substring(1)
    }
```

```
    ##Remove quotes around file name
    if($SGFile[0] -eq "")
    {
        [int]$count = $SGFile.get_length()
        $SGFile = $SGFile.remove($count-1)
        $SGFile = $SGFile.substring(1)
    }
```

```
    ##Remove quotes around file name
    if($DBFile[0] -eq "")
    {
        [int]$count = $DBFile.get_length()
        $DBFile = $DBFile.remove($count-1)
        $DBFile = $DBFile.substring(1)
    }
```

```
$FileInputText = get-content -Path $Txtfile
$FileInputSG = get-content -Path $SGfile
$FileInputDB = get-content -Path $DBfile
$i=0
```

```
foreach($file in $FileInputText)
{
```

```

        if($File[0] -eq "")
        {
            [int]$count = $File.get_length()
            $File = $File.remove($count-1)
            $File = $File.substring(1)
        }

        $FileInput = get-content -Path $file
        $SGName = $FileInputSG[$i]
        $DBName = $FileInputDB[$i]

        $targetdatabase = get-mailboxdatabase -identity
"$mbxName\$SGName\$DBName"

        foreach($User in $FileInput)
        {
            get-mailbox -identity $user | move-mailbox -TargetDatabase:$targetdatabase
-confirm:$False -maxthreads:4
        }

        $i++
    }
}

##Function to call the move-mailbox cmdlet
function movembx()
{

    ##Remove quotes around file name

    if($TxtFile[0] -eq "")
    {
        [int]$count = $TxtFile.get_length()
        $TxtFile = $TxtFile.remove($count-1)
        $TxtFile = $TxtFile.substring(1)
    }

    $FileInput = get-content -Path $Txtfile
    $targetdatabase = get-mailboxdatabase -identity "$mbxName\$SGName\$DBName"

    foreach($User in $FileInput)
    {
        get-mailbox -identity $user | move-mailbox -TargetDatabase:$targetdatabase -
confirm:$False -maxthreads:4
    }
}

if ($inputType -eq "1")
{
    ##Moving mailboxes for a single storage group

```

```
##Call command line function to gather user input

commandline

##Call movembx function to move the mailboxes for the users gathered from the
commandline
    movembx
}
elseif ($inputType -eq "2")
{

    ##Moving mailboxes for multiple storage groups on the same Exchange server

    ##Call command line function to gather user input

    commandlineMultiple

    ##Call movembx function to move the mailboxes for the users gathered from the
    commandline

    movembxMultiple

}
else
{
    write-host "Please enter a valid option."
}
}
```

Appendix B: Why 64-bit Microsoft Exchange Server?

Microsoft Exchange Server 2003 has reached the limits of a 32-bit operating system. Performance testing indicates that, even with increasing processor frequency and efficiency, the system bottleneck for an Exchange server is the memory subsystem. An historical look at Microsoft Exchange Server benchmarking results shows that the highest performing systems have been trending down the scale from 8 processors common in Exchange 5.5, to quad processors in Microsoft Exchange 2000 Server and finally to dual-core two-processor system for Microsoft Exchange Server 2003. Currently, quad processor systems are only scaling to approximately 50% of the total CPU resources at the maximum MMB3 benchmark results. The scaling and performance benefits between dual-processor and quad-processor systems are minimal.

Microsoft Exchange Server 2007 is a 64-bit application that requires a 64-bit version of Windows Server 2003 and 64-bit server hardware. The hardware compatibility options are servers configured with Intel Xeon EMT64 processors or AMD Opteron™ AMD64 processors. The Windows Server 2003 x64 Editions can address 16 terabytes of virtual address space, 8 terabytes for kernel address space and 8 terabytes for user address space. The increased user mode address space allows for significant growth of the ESE buffer cache size. Increased ESE buffer cache size and larger 8KB pages associated with Microsoft Exchange Server 2007 are expected to significantly decrease the I/O requirements of a Microsoft Exchange Server 2007 server.

Note: The /3GB switch is not supported in the Windows Server 2003 x64 Editions.

In a 32-bit Windows Server 2003 environment operating with 4 gigabytes RAM, the operating system reserves 2 gigabytes for kernel memory and 2 gigabytes for user space. The store.exe process in Microsoft Exchange Server 2003 is the highest consumer of user space memory. On a heavily loaded Microsoft Exchange Server 2003 server, the store.exe process can easily consume more than 2 gigabytes of memory. This leads to virtual memory fragmentation and requires a restart of the Microsoft Exchange Server services on the system to reclaim the fragmented memory. The resolution for this problem is to implement the /3GB switch in the boot.ini file as described in Microsoft's KB article 328882 "[Exchange memory use and the /3GB switch](#)." This operating system switch allocates 1 gigabyte of memory to the kernel memory space and 3 gigabytes of memory to the user space, providing an additional 1 gigabyte of address space to the user memory space. The additional user memory space allows the store.exe process to leverage more than 2 gigabytes of virtual memory. However, reducing the kernel memory space can lead to additional memory management challenges associated with the availability of system page table entries. The recommendation is to use the /USERVA=3030 boot.ini switch in conjunction with the /3GB switch. The /USERVA=3030 boot.ini switch provides an additional 42 megabytes of memory to the kernel memory space. This is documented in Microsoft KB article 810371 "[Using the /Userva switch on Windows Server 2003-based computers that are running Exchange Server](#)".

Microsoft Exchange Server 2003 is an I/O intensive application which requires extremely low read and write I/O latency for optimal performance. Maintaining a low I/O latency is typically achieved by increasing the number of disk drives in an I/O subsystem. In recent years, disk drives have significantly increased in capacity while the performance of these drives has remained relatively constant. This leads to an imbalance where the number of spindles required for performance is much larger than the number of spindles required for capacity. A properly sized Microsoft Exchange Server 2003 server will typically result in excess storage capacity, a scenario that is usually frowned upon by most organizations.

For more information

To learn more about award winning industry hardware, please visit the HP website at www.hp.com

For more information about HP ProLiant servers, visit <http://www.hp.com/go/proliant>

For more information about HP BladeSystem server blades and architecture, visit: <http://www.hp.com/go/bladeSystem>

HP ActiveAnswers website: <http://www.hp.com/solutions/activeanswers/>

Microsoft KB article comparing 32-bit and 64-bit Windows editions <http://support.microsoft.com/kb/294418/en-us>

To help us improve our documents, please provide feedback at www.hp.com/solutions/feedback

© 2006 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Intel, Xeon, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. Microsoft and Windows are U.S. registered trademarks of Microsoft Corporation. AMD Opteron is a trademark of Advanced Micro Devices, Inc.

4AA0-9163ENW, Rev. 1, November 2006

