

Quest® Migration Manager 8.15

Tips and Tricks



© 2020 Quest Software Inc. ALL RIGHTS RESERVED.

This guide contains proprietary information protected by copyright. The software described in this guide is furnished under a software license or nondisclosure agreement. This software may be used or copied only in accordance with the terms of the applicable agreement. No part of this guide may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording for any purpose other than the purchaser's personal use without the written permission of Quest Software Inc.

The information in this document is provided in connection with Quest Software products. No license, express or implied, by estoppel or otherwise, to any intellectual property right is granted by this document or in connection with the sale of Quest Software products. EXCEPT AS SET FORTH IN THE TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT, QUEST SOFTWARE ASSUMES NO LIABILITY WHATSOEVER AND DISCLAIMS ANY EXPRESS, IMPLIED OR STATUTORY WARRANTY RELATING TO ITS PRODUCTS INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT SHALL QUEST SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, PUNITIVE, SPECIAL OR INCIDENTAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION OR LOSS OF INFORMATION) ARISING OUT OF THE USE OR INABILITY TO USE THIS DOCUMENT, EVEN IF QUEST SOFTWARE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Quest Software makes no representations or warranties with respect to the accuracy or completeness of the contents of this document and reserves the right to make changes to specifications and product descriptions at any time without notice. Quest Software does not make any commitment to update the information contained in this document.

If you have any questions regarding your potential use of this material, contact:

Quest Software Inc.

Attn: LEGAL Dept

4 Polaris Way

Aliso Viejo, CA 92656

Refer to our Web site (<https://www.quest.com>) for regional and international office information.

Patents


Quest Software is proud of our advanced technology. Patents and pending patents may apply to this product. For the most current information about applicable patents for this product, please visit our website at <https://www.quest.com/legal>.

Trademarks

Quest, the Quest logo, and Join the Innovation are trademarks and registered trademarks of Quest Software Inc. For a complete list of Quest marks, visit <https://www.quest.com/legal/trademark-information.aspx>. All other trademarks and registered trademarks are property of their respective owners.

Legend

 **CAUTION: A CAUTION icon indicates potential damage to hardware or loss of data if instructions are not followed.**

 **IMPORTANT, NOTE, TIP, MOBILE, or VIDEO:** An information icon indicates supporting information.

Migration Manager Tips and Tricks

Updated - February 2020

Version - 8.15

Contents

Introduction	6
Environment Assessment, Planning, and Testing	7
Environment Assessment and Planning	7
Trusts	8
SID Filtering	9
Network Topology Diagrams	9
Firewalls and Security	10
Exchange Topology Diagram	10
Active Directory Design	10
Active Directory Management (Delegation and Provisioning Model)	11
Exchange Design	11
Host Name Resolution	11
Group Policy Object Link Migration	12
Migration Cookbook	12
Mail Redirection between Source and Target Exchange Organizations	12
Outlook Offline Folders	13
Rolling Exchange Migration	13
Source Directory Preparation	13
Third-party Applications	14
Exchange Clusters	14
Cluster Server Migration	15
NetApp Filers and Other Storage Solutions	15
Backup Strategy	15
Service Attributes	16
Administrative Accounts	17
How Many Projects to Have	17
Delegating Migration Activities	17
Testing	18
Environment Preparation	19
Basic Migration Steps	20
Pilot Migration	20
Overview of Migration Scenarios	20
Active Directory Migration	21
Step 1. Establish Directory Synchronization (Optional)	24
Step 2. Migrate the Directory	25
Step 3. Process Distributed Resources	26
Step 4. Move End-user Workstations and Servers to the Target Domain	32

Step 5. Process BackOffice Servers	34
Step 6. Move BackOffice Servers to the Target Domain	35
Step 7. Stop Directory Synchronization (Optional)	35
Step 8. Disable Source Accounts	35
Step 9. Enable SID filtering	35
Step 10. Clean Up SIDHistory Attributes	36
Step 11. Clean Up Legacy Account Permissions	36
Step 12. Clean Up Service Attributes Used for Migration	36
Step 13. Decommission the Migrated Environments	36
Exchange Migration	37
Step 1. Establish Directory Synchronization_EX	39
Step 2. Start Public Folder and Calendar Synchronization	41
Step 3. Establish Free/Busy Synchronization (Optional)	41
Step 4. Synchronize Mailbox Data	42
Step 5. Switch to the New Exchange Mailboxes	42
Step 6. Update Outlook Profiles	42
Step 7. Change the Mail Exchanger or Alias Records	42
Step 8. Stop and Uninstall the Synchronization Agents	43
Step 9. Clean Up the Additional SMTP Addresses and Service Attributes Used for Migration ...	43
Step 10. Decommission the Migrated Environments	43
Active Directory and Exchange Migration	44
Parallel Active Directory and Exchange Migration	44
Exchange First, Then Active Directory Migration	48
Migration Steps	48
Rollback	48
Roll Back Active Directory Changes	49
Roll Back Changes Made to Exchange Organization	49
Considerations for Active Directory Migration and Resource Update	51
Using Import Lists to Modify Attributes or Merge and Rename Accounts	51
Linked Attributes and Group Migration	52
Subsequent Migrations	52
Skipping Attributes	53
Pre-installing the Resource Updating Agents	53
Migrating Services	53
Interoperability with Quest Collaboration Services	54
Considerations for Exchange Migration	55
Synchronization Agents	55
Directory Synchronization	57
Public Folder Synchronization	57
Configure Public Folder Jobs	57

- What Accounts to Use for Public Folder Synchronization 58
- How to Configure Public Folder Synchronization Collections 59
- Whether Hierarchy Reorganization is Required 59
- How Public Folder Deletions Are Handled 60
- When Public Folder Resynchronization is Required 61
- Mailbox Synchronization 61
- Mail Synchronization Agents and Jobs 61
- Using Collections to Map Exchange Migration Wizard 61
- Offline Folder (OST) Files and Remote Users Collections 63
- Restructuring Exchange Servers 65
- Calendar Synchronization 65
 - Calendar Synchronization Agents and Jobs 65
- One-Way vs. Two-Way Synchronization 66
 - Free/Busy Synchronization 66
- Mailbox Switch and Profile Update 67
 - Mailbox Switch 67
 - Client Profile Update 68
- Interoperability with Other Quest Products 68
- Preferred Settings for the Directory Synchronization Agent 69**
- Directory Synchronization Agent Placement 70**
- Indexing Service Attributes 71**
- Full Directory Resynchronization 72**
- Conclusion 73**
- Environment Preparation Checklist 74**
- Exchange Migration without Trusts 76**
- Active Directory Migration without Trusts 77**
- About us 79**
 - Technical support resources 79

Introduction

Target Audience

This document is intended for network administrators, consultants, analysts, Exchange architects, and any other IT professionals who are considering migrating Active Directory or Exchange using Quest Migration Manager.

Assumptions

All migration scenarios described in this document assume that:

- Trusts are established between each source and target domain involved in the Active Directory or Exchange data migration.
- SIDHistory is added to migrated accounts and used during the whole co-existence period to ensure that users will have the same access to resources when they start using their target accounts.

If you cannot establish trusts or use SIDHistory due to your corporate policy or other reasons contact Quest Professional Services for a custom migration scenario designed for your specific environment.

Environment Assessment, Planning, and Testing

Environment Assessment and Planning
Testing
Environment Preparation

Environment Assessment and Planning

The first step in migration is to assess your environment and design an appropriate migration plan. The following tools can help:

- Quest MessageStats—Reports on the current Exchange environment and assists with planning migration activities to the target Exchange organization. For more information, visit www.quest.com/messagestats/.
- Quest Reporter—Reports on the current Active Directory environment and assists with planning migration activities in the target Active Directory forest or domain. For more information, visit <https://quest.com/products/reporter/>.

The related topics detail our best practice recommendations for environment assessment and planning:

- [Trusts](#)
- [SID Filtering](#)
- [Network Topology Diagrams](#)
- [Firewalls and Security](#)
- [Exchange Topology Diagram](#)
- [Active Directory Design](#)
- [Active Directory Management \(Delegation and Provisioning Model\)](#)
- [Exchange Design](#)
- [Host Name Resolution](#)
- [Group Policy Object Link Migration](#)
- [Migration Cookbook](#)
- [Mail Redirection between Source and Target Exchange Organizations](#)
- [Outlook Offline Folders](#)
- [Rolling Exchange Migration](#)
- [Source Directory Preparation](#)
- [Third-party Applications](#)

- [Exchange Clusters](#)
- [Cluster Server Migration](#)
- [NetApp Filers and Other Storage Solutions](#)
- [Backup Strategy](#)
- [Testing](#)
- [Environment Preparation](#)
- [Service Attributes](#)
- [Administrative Accounts](#)
- [How Many Projects to Have](#)
- [Delegating Migration Activities](#)

Trusts

Why Use Trusts?

We recommend that you establish two-way external trusts between each source and target domain that will participate in migration.

If the forest functional level in both source and target forests is set to Windows 2003 or higher, you can establish forest trust between the forest root domains.

Trusts make it possible to resolve objects' security identifiers (SIDs), which in turn helps to distinguish objects and enable you to check whether everything is going right. Trusts also help provide co-existence of the source and target environments during the migration process, including uninterrupted access to the resources for both switched users and users not yet switched.

i | **NOTE:** Remember that external trusts between Active Directory domains that belong to different Active Directory forests are not transitive. You should establish trusts between each source Active Directory domain and target Active Directory domain individually.

If Trusts Are Not Established

When deciding whether to establish trusts, remember that if no trusts are established, the following restrictions apply:

- You will not be able to use a single administrative account for migration.
- You will have to switch users and resources at the same time. This means that when a user starts using its target account (normally, when the user's workstation is moved to the target domain), all resources must be updated, so that the target user has the same access to the resources as the corresponding source user.
- When working with remote Exchange servers, console establishes net use connections automatically; thus no trusts between the console machine where Migration Manager is installed and all Exchange servers where the synchronization agents are installed are needed. However if a net use connection between the console machine and remote Exchange server was already established using different account, you may need to manage this connection manually.
- The computer on which Migration Manager is installed must be a member of the domain in which Windows 2000 or Windows Server 2003-based Exchange cluster servers reside. If you have Windows 2000 or Windows Server 2003-based cluster servers in both the source and target domains, you need trusts to be established between the domains.

- If you migrate Exchange first and set the source user's account to be the Associated External Account for the corresponding Exchange mailbox, users will not be able to log on to the target mailboxes with their source accounts.
- Users will have to specify the target security account when they are switched to the target server. Because there are no trusts, their source accounts will not have permissions for the target mailboxes.

The migration scenarios described in the [Basic Migration Steps](#) topic assume that trusts are established between the source and the target domains.

SID Filtering

For information about SID filtering, refer to [Configuring SID Filtering](#).

Network Topology Diagrams

It is essential to have a solid understanding of your infrastructure and know your bandwidth at every location. Site topology diagrams with inter-site connections speeds, Active Directory design, OU hierarchy, domain controller placement, Flexible Single Master Operation (FSMO) role placement, and BackOffice server location, along with other infrastructure diagrams, help you see the whole picture and make good decisions.

Site topology diagrams and information about the number of users in each site help you decide, for example, whether to perform the migration centrally or to divide the migration into parts and allow some remote sites to be migrated locally (in other words, to delegate the migration activities).

Information about the locations of BackOffice servers can help you plan resource processing tasks and delegate these tasks to other administrators. Use the *Location of BackOffice Servers* report to find the computers that have the BackOffice service installed.

Site topology information also helps you decide which domain controllers to use for migration. You should choose the source and target domain controllers located in the same site as the Directory Synchronization Agent that will process migration and synchronization jobs between these domains to avoid traffic span across slow WAN links.

We recommend that the diagrams be created using visual tools (for example, Microsoft Visio) and printed out for convenience.

Also, if you want to migrate a large number of accounts (more than 500) to the target domain in one session, it is recommended that you select a target domain controller that owns the RID master role. Otherwise, the target domain controller may experience delays getting the next set of RIDs from the RID master when creating objects in Active Directory.

How RID Allocation Works

When a domain controller creates a security principal object such as a user or group, it attaches a unique Security ID (SID) to the object. This SID consists of a domain SID (the same for all SIDs created in a domain), and a relative ID (RID) that is unique for each security principal object created in a domain.

Each Windows domain controller in a domain is allocated a pool of RIDs that it is allowed to assign to the security principals it creates. When a domain controller's allocated RID pool falls below a threshold, that domain controller issues a request for additional RIDs to the domain's RID master. The domain RID master responds to the request by retrieving RIDs from the domain's unallocated RID pool and assigns them to the pool of the requesting domain controller. There is one RID master per domain in a forest.

Use Reporter's *FSMO Roles* report to help you determine FSMO placement.

Firewalls and Security

Since the Migration Manager agents are installed and updated from the console over RPC and the agents transfer data directly between source and target servers over RPC as well, RPC traffic must be allowed over the routers separating the subnets.

Make sure that the following ports are open on workstations, servers, routers, and firewalls: 135 and 137–139.

For more detailed information on what ports and protocols Microsoft operating systems and programs require for network connectivity, refer to Microsoft Knowledge Base article 832017, “Service overview and network port requirements for the Windows Server system,” at <http://support.microsoft.com/kb/832017>.

You can use the DCdiag and Netdiag utilities from Windows Support Tools to test network connectivity. To install Windows Support Tools, run **Setup.exe** from the **\SUPPORT\TOOLS** folder of Windows distributive CD. For more information about the utilities, refer to their online help and other documentation.

In Windows XP Service Pack 2, Microsoft introduced the Security Center, which includes a client-side firewall application. The firewall is turned on by default and configured to filter the packets sent to the ports 137–139, and 445. These ports are used by the File and Printer Sharing service, which must be installed and running on the computer to be updated.

To make sure Resource Updating Manager correctly updates computers running workstation Windows versions, add the File and Printer Sharing service to the firewall Exceptions list and unblock ports 137–139 and 445. Alternatively, consider deploying Resource Updating Manager agents using group policy or similar methods. For more information on resource processing requirements, refer to the [Step 3. Process Distributed Resources](#) topic.

When granting the required permissions to the administrative accounts in Active Directory, you should also make sure that permissions inherited from the parent are not blocked at any level in your Active Directory.

Exchange Topology Diagram

An Exchange topology diagram helps you plan directory, public folder, and mailbox synchronization jobs. This diagram should display Active Directory sites, Exchange administrative and routing groups, Exchange servers, bridgehead servers, and the number of public folders and mailboxes and their sizes on each Exchange server.

You can use this diagram when planning for directory, public folder, and mailbox synchronization to help you choose source and target Exchange servers that are in the same physical location, thus preventing large amounts of data from being transferred across slow WAN links.

We recommend that the diagrams be created using visual tools (for example, Microsoft Visio) and printed out for convenience.

Active Directory Design

It is recommended that your Active Directory design be certified by Microsoft or a certified partner with vast experience with Active Directory. Active Directory design (logical structure and site topology) must be completed prior to any migration.

Before the migration, you should check Active Directory replication to ensure that any changes to Active Directory (such as object creations and deletions) will be properly replicated between domain controllers and Global Catalog servers in different locations.

Active Directory Management (Delegation and Provisioning Model)

When planning a migration to Active Directory, it is recommended that your delegation model be engineered ahead of the project in order to identify all business rules that must be enforced. Quest Active Roles is the most comprehensive delegation and provisioning product available. The best practice is to install and configure Active Roles soon after implementing Active Directory and prior to migrating any accounts to ensure a pristine Active Directory from the start. For more details on Active Roles, see <https://quest.com/products/activeroles-server/>.

Exchange Design

It is recommended that the Exchange design be completed and certified by Microsoft or a certified partner with vast experience with Exchange.

We also recommend you fully complete the target Exchange deployment prior to starting the Exchange migration to avoid Exchange organization re-enumerations in Migration Manager when you make changes to the environment (such as create administrative groups or add new Exchange servers and information stores). Having all servers in place also lets you configure all synchronization jobs with the right target servers and avoid any mailbox moves in the future.

In an inter-org migration, running the target Exchange infrastructure in native mode from day one also allows you leverage Exchange features such as administrative groups and routing groups. Running Exchange in native mode from the very beginning allows you re-configure and optimize mail routing for your target organization according to Microsoft Exchange best practices before users start using the target Exchange servers.

Host Name Resolution

Before the migration, you should check that the Domain Naming System (DNS) and the Windows Internet Naming Service (WINS) are functioning properly so that machine names can be successfully resolved into IP addresses.

You can use the DCDiag and NetDiag utilities from Windows Support Tools to test network connectivity and identify the host name resolution problems. To install Windows Support Tools, run **Setup.exe** from the **\SUPPORTTOOLS** folder of Windows Server distribution CD. For more information about the utilities, refer to their online help and other documentation.

DNS Name Resolution

The DNS names of the source and target servers and of the servers on which ADAM or AD LDS project partition and SQL configuration database are located must be successfully resolved to IP addresses from the servers running the Directory Synchronization Agents as well as from the console.

Make sure that the DNS is configured and functioning properly in your environment.

NetBIOS Name Resolution

Since the agents installed on the source and target Exchange servers communicate with each other, Exchange source-target server pairs must be able to resolve each other's NetBIOS names. In other words, each server must be able to "see" the other servers by NetBIOS.

Windows Internet Naming Service (WINS) is usually used to resolve servers' NetBIOS names to IP addresses. If WINS is not configured in an environment, host files can be used instead.

You should check the host NetBIOS name resolution and make sure that the servers' NetBIOS names can be resolved from the console as well.

Group Policy Object Link Migration

Migration Manager does not migrate group policy objects (GPOs) from one Active Directory domain to another, regardless of whether the source and the target domains belong to the same or different forests.

However, you can use Microsoft Group Policy Management Console (GPMC) to migrate group policy objects from one domain to another in conjunction with the GPMCEExport utility from the Migration Manager for Active Directory Resource Kit.

The GPMCEExport utility allows you create a mapping file in the format required by GPMC. When you later migrate group policy objects from one Active Directory domain to another, you can use the created mapping file to let GPMC automatically translate all source security principles (specified in GPO security descriptors) into the corresponding target objects.

For more information on the GPMCEExport utility, refer to the *Quest Migration Manager for Active Directory Resource Kit User Guide*.

Quest Reporter provides a number of reports on group policy that can be useful for analyzing the current group policy design and planning for GPO migration. These reports are:

- **Group Policy Object Summary** – Displays links and other details for Group Policy Objects, but no settings or security.
- **Group Policy Objects** – Displays the details for selected Group Policy Objects, such as the name, security, links, settings, and other properties.
- **Linked Group Policy Objects** – Displays the details for Group Policy objects that are linked to objects in the directory.
- **Unlinked Group Policy Objects** – Displays the details for GPOs for which no links could be found. You may want to delete these unlinked GPOs if they are no longer required.

You can also find additional information on using GPMC for GPO at <http://technet.microsoft.com>. Try searching for “Migrating GPOs Across Domains”.

Migration Cookbook

We recommend that you create a migration cookbook: a detailed, step-by-step guide of the migration activities specific to your environment that outlines a repeatable process that can be tested and then mirrored throughout the environment.

This is particularly useful in the case of delegated migration, which is when a number of independent administrators located in different sites migrate the accounts they are responsible for. The cookbook can be given to all of the delegated administrators to standardize the migration tasks.

Mail Redirection between Source and Target Exchange Organizations

To ensure that users do not lose their mail during the migration period, and to make the migration from source Exchange organization to target smooth, Migration Manager establishes mail redirection between the source and the target Exchange servers.

Accordingly, Migration Manager requires the source and target Exchange organizations be connected using SMTP connectors. For step-by-step instructions for setting up, configuring, and testing the connectors, refer to the *Exchange Environment Preparation* documents for your particular environments.

Additional SMTP addresses are used for mail redirection. You should analyze your environment for SMTP namespaces and for redirection implement SMTP address templates that are not being used.

Outlook Offline Folders

One widely-used Microsoft Outlook feature is offline access to mailbox folders. The offline folders (OST) file is stored on a user's computer and keeps a local replica of the corresponding folders in the user's Exchange mailbox.

There may be a number of the remote users in your source environment who work with local offline folder (OST) files and only occasionally connect to their Exchange mailboxes. Because each OST file is associated with only one Exchange mailbox and cannot be used with any other mailbox, a user cannot continue using the same OST file with the new mailbox after migration.

Migration Manager for Exchange allows you keep the existing OST files without OST file resynchronization when a remote or laptop user's mailbox is switched and the Outlook profile is updated.

The mailboxes of these users should be grouped in *Remote Users Collections* and processed separately from other mailbox collections after the directory synchronization has been completed and before the mailbox synchronization is started.

The mailboxes of the Remote Users Collections are processed by the Mail Source Agent only. The agent creates target mailboxes corresponding to the source mailboxes in a Remote Users Collection. While a mailbox is being processed by the agent, it is unavailable to the user. Therefore, it is recommended to schedule process of Remote Users Collections for a time when the users normally do not use their mailboxes, such as 8 p.m. to 8 a.m. or during the weekend.

Rolling Exchange Migration

Exchange migration is a complex task that can take a long period of time. Because migration is performed in the production environment, it needs to be as transparent as possible to the end users.

A *rolling* Exchange migration strategy allows you to gradually collapse migrated Exchange servers and, at the same time, provide coexistence between Exchange organizations.

This strategy is useful if you need to reuse current production Exchange server hardware in the new Exchange environment. It is also useful if you must remove the administrative overhead and cost of supporting the source Exchange servers that no longer host active mailboxes by gradually decommissioning the source Exchange environment and its hardware.

Rolling Exchange migration requires more planning regarding decommissioning activities, the directory cleanup process, etc., as compared to the standard Exchange migration scenario outlined in the [Basic Migration Steps](#) topic.

If you need to decommission migrated production Exchange servers and re-use the freed-up hardware in the new Exchange environment, contact Quest Professional Services for a custom migration scenario designed for your specific environment.

Source Directory Preparation

Handling Disabled and Inactive Accounts

Migration Manager can automatically filter out disabled accounts during migration. We recommend you use this option and not bring disabled accounts to Active Directory.

You might also want to filter out user accounts that have not been logged on to for a long period of time and disable them prior to migration. To locate such accounts, use the following reports in Quest Reporter:

- **Accounts that have never logged on** – Displays user accounts that have never logged on
- **Inactive Accounts** – Shows all accounts that are disabled or deemed inactive during the time period you specify

Handling Accounts with Duplicate Names

There may be user accounts already created in the target Active Directory domain with the same names as the source accounts. If object matching by name is turned on for the domain pair (the rule is turned on by default), Migration Manager considers two objects to be duplicates if the source object's SAMAccountName is equal to the target's SAMAccountName.

Although Migration Manager reports on duplicate accounts during the migration session and allows you rename, merge, or skip them on the fly, the best practice is to handle duplicates before migration. You should find duplicate user and group names and determine which accounts belong to the same person and which do not. If two accounts, source and target, have the same name and belong to the same person, they should be merged during migration. If not, you should either rename one of them or skip them during migration.

To find duplicate user and group names, use the *Duplicate Users* and *Duplicate Groups* reports from Quest Reporter.

Third-party Applications

You should check your source environment for all business-critical third-party applications, such as ERP systems, fax software, and meta-directory services. These applications should be deployed in the test lab first and properly tested before you start migration in the production environment. Refer to the [Test the Third-party Applications](#) subsection.

Exchange Clusters

Migration Manager supports multi-node clusters running multiple Exchange Virtual Servers. Migration Manager detects such systems and configures agent services for automatic failover together with the Exchange services.

i **NOTE:** By default, the **QMMEXNode\$** and the **QMMEX\$<ExchangeServerName>\$** shares have the same path. Otherwise, to be able to install the agents on the active node using the Install Agents Wizard take the following steps:

1. Browse to the Exchange cluster server in the **Target Exchange Organizations** node of the Migration Manager Console.
2. Right-click the Exchange cluster server and select **Properties** from the shortcut menu.
3. Click **OK**.
4. Run Cluster Administrator.

Bring the **QMMEX\$<ExchangeServerName>\$** resource online.

However, if several Exchange virtual servers are running on a single cluster node, the agents can be installed and run only on one Exchange virtual server at a time. Thus, such Exchange servers can only participate in migration with Migration Manager consecutively, one by one.

There are also some Exchange virtual server limitations on clusters that have more than two nodes. Refer to Microsoft Knowledge Base article 329208, "XADM: Exchange Virtual Server Limitations on Exchange 2000 Clusters and Exchange 2003 Clusters That Have More than Two Nodes," at <http://support.microsoft.com/default.aspx?scid=kb;en-us;329208>.

Migration Manager fully supports Exchange 2007 Cluster Continuous Replication (CCR). Migration Manager detects such systems and configures agent services for automatic failover together with the Exchange services. The rest two types of Exchange 2007 Continuous Replication that are Standby Continuous Replication (SCR) and Local Continuous Replication (LCR) are also supported, but in case of failover a full resynchronization of all mailboxes, calendars and public folders involved in the migration must be performed.

Cluster Server Migration

Cluster server migration includes the following tasks:

- Move a cluster server to another domain.
- Process the cluster server's ACLs.

Cluster servers require special treatment when you move clusters to the target domain and update cluster resources (such as cluster shares, cluster database, and cluster printers). Follow the recommendations and guidelines in the [Active Directory Migration](#) topic to update and move clusters.

NetApp Filers and Other Storage Solutions

Since Quest Migration Manager version 8.4, Network Attached Storage (NAS) devices (such as NetApp Filer) and Storage Area Network (SAN) devices can be updated not only using the command-line updating tool (vmover.exe) but using the console (with new version of Resource Updating Manager).

Since these devices usually store a large amount of data and are actively used by a large number of users, you should carefully plan for their update timeframe and procedures.

We also recommend that all hardware updates (such as NAS and SAN) be carefully tested in a test lab first.

Backup Strategy

We recommend that you back up your source and target Exchange infrastructures before implementing Migration Manager in your production environment. We also recommend that source and target Active Directory data be backed up at least twice a day during migration.

Quest Recovery Manager for Active Directory is the most comprehensive tool that provides granular backup and restoration of Active Directory objects. For more information about Recovery Manager for Active Directory, see <https://quest.com/products/recovery-manager-for-active-directory/>.

Transaction Log File Cleanup

When Migration Manager for Exchange synchronizes mail and public folders, for every megabyte of data migrated from the source to the target, a transaction log file of equal size is generated on the target Exchange server. Exchange-aware backup applications purge the transaction logs after the backup completes. By the time the backup finishes, all logged transactions have already been applied to the store and backed up to tape, making log cleaning safe.

Run normal backup procedures to delete the transaction logs throughout the migration on both the source and the target Exchange servers.

If normal backup operations do not delete the transaction logs, then you should ensure that appropriate disk space is reserved for the expected transaction log growth.

Alternatively, you can enable circular logging during the migration to avoid transaction log growth. However, Microsoft recommends that circular logging be turned off on Exchange servers. If circular logging is turned on, bear in mind that large transaction logs can still be generated on the Exchange server, and watch closely that the logs are properly cleaned after backup. Keep in mind that if you use circular logging, only full backups can be performed.

For more information, refer to Microsoft Knowledge Base article 147524, "XDAM: How Circular Logging Affects the Use of Transaction Logs," at <http://support.microsoft.com/default.aspx?scid=kb;en-us;147524>.

Migration Manager for Exchange Agents and Backup Schedules

If any backup tools are installed on the servers where Migration Manager for Exchange agents are to run, the schedules for these tools and the agents should not overlap; that is, the backup utility and the agents should be scheduled to work during different hours.

In any case, use of any backup utility together with Migration Manager should be tested in the laboratory before use in the production environment.

ADAM/AD LDS Database and SQL Configuration Database Backup

All migration project configuration data are stored in the ADAM or AD LDS database and the SQL configuration database. It is very important that you back up these databases regularly, because losing them could be disastrous to your migration project.

For procedures for backing up an ADAM or AD LDS instance, refer to the ADAM or AD LDS online help.

You can also use Quest Recovery Manager for Active Directory to back up and restore ADAM or AD LDS databases. For more information about Recovery Manager for Active Directory, see <https://quest.com/products/recovery-manager-for-active-directory/>.

As an alternative to backing up the ADAM or AD LDS database, you can install a replica of the ADAM or AD LDS instance on another server in the network and simply connect to that instance if the first one is lost or corrupted.

To back up a SQL configuration database located on a Microsoft SQL Server or a server running Microsoft Desktop Engine (MSDE), use standard SQL backup procedures. Refer to the Microsoft article "Backing Up and Restoring Databases" at [http://msdn2.microsoft.com/en-us/library/aa196685\(SQL.80\).aspx](http://msdn2.microsoft.com/en-us/library/aa196685(SQL.80).aspx) and Microsoft SQL Server documentation for more details.

Service Attributes

The Directory Synchronization Agent in Migration Manager uses Active Directory attributes called service attributes to store necessary information. For each synchronized object, the Directory Synchronization Agent sets the *auxiliary* attribute and the *matching* attribute. Different attributes are used as auxiliary and *matching* attributes for different object classes in Active Directory.

For the attributes used as service attributes by default, refer to the *Quest Migration Manager for Active Directory—User Guide*.

CAUTION: Make sure that attributes to be used as service attributes during directory synchronization are not used in either the source or the target domain and do not contain any values. Otherwise, the Directory Synchronization Agent will not be able to match the objects that contain any data in service attributes. Thus, such objects cannot be migrated and synchronized.

Administrative Accounts

Migration Manager requires administrative access to the source and target domains and Exchange organizations involved in migration. Administrative access is also required to the console machine on which Migration Manager is installed.

When you install Migration Manager, create domain pairs, install synchronization agents, or create synchronization jobs, make sure you specify an administrative account that has the required privileges. The required permissions are listed in the *Migration Manager—System Requirements and Access Rights* document.

You can create and use different administrative accounts for different components; however, due to Migration Manager's distributed architecture, the best practice is to create a single administrative account, grant this account all necessary permissions, and use it for all migration activities. This greatly simplifies project management and minimizes the number of issues related to the lack of permissions.

! CAUTION: This powerful account must be maintained closely and should be deleted after the project is complete. It is recommended that this account be owned by one individual and one backup individual (or as few individuals as possible.)

Refer to the *Appendices* section of the *Migration Manager—System Requirements and Access Rights* document for the guidelines on creating a single administrative account and granting required permissions.

How Many Projects to Have

Migration Manager allows you to configure a separate directory synchronization job for each pair of source and target domains. With Migration Manager you can also split your source forest or Exchange organization or merge a number of forests or Exchange organizations into one target.

The best practice is to use a single migration project for all your migration activities, for the following reasons:

- Because you can work with only one project at a time from the console, configuring all the domain pairs, synchronization jobs, synchronization agents, and so on is easier if you use a single project.
- The project database is the central storage for the object mapping information that is used during resource processing to translate rights and permissions. Using multiple projects requires you to update resources separately for each project for which you have to update rights, permissions, local group membership, etc.
- You can delegate rights to perform migration tasks (such as configuring domain pairs, migrating accounts, and updating resources) to different persons and migration teams. A single project helps you track who has what rights (roles) in the project more easily. For more information about delegation, refer to the [Delegating Migration Activities](#) topic.

Delegating Migration Activities

When you install Migration Manager and create a migration project, the account used to create the project is granted the Full Admin right over the project. This account is then able to delegate certain rights within the project to the other administrators to perform their specific migration tasks.

Using Migration Manager, you can delegate the following migration tasks to the other administrators:

- Account migration
- Resource update

Delegating account migration is useful when you want the actual account migration to be done by the other persons but you do not want these persons to know the administrative credentials to access the source and target domains. You can also limit the scope of the delegated migration to certain OUs in the source and target domains. For more information on delegated migration, refer to the *Quest Migration Manager for Active Directory—User Guide*.

Delegating resource updating tasks to other administrators is useful when you cannot get administrative access to computers located in remote sites or resource domains. Running resource update locally is also useful if computers to be updated are located across slow WAN connections. In such cases, you can delegate the resource updating tasks to the remote site or to other domain administrators who have the required level of access and are located within an area of good connectivity to the computers to be updated.

When you delegate resource processing, we recommend you use resource processing task setups rather than INI files. Refer to the [Step 3. Process Distributed Resources](#) topic for more details.

Testing

It is recommended that you test the migration strategies using an environment that mirrors production as closely as possible. This is critical for testing scalability. After thorough testing, begin the pilot phase and then full deployment.

Creating a Test Environment

Migration strategies and activities should be tested in a lab that closely resembles your current production environment, including resources, servers, applications, legacy design, and bandwidth. So that you can completely run through your migration scenarios, we recommend you create a copy of your live Active Directory and Exchange in the test lab. This can be done by doing any of the following:

- Move one of your live domain controllers into the isolated network that will become your test lab and assign it the PDC-emulator role. You may want to create an additional domain controller in production and then use it for testing purposes in the lab. Note that you will also have to seize all FSMO roles to this domain controller and make it a Global Catalog as well after you moved it into the isolated network.
- Create an image of your live domain controllers and Exchange servers using third-party tools and restore the image to another box with a similar hardware configuration in the lab.
- Export the source Active Directory into a CSV or LDIF file using third-party tools (such as CSVDE, LDIFDE, or LDAPBrowser) and import it into the test environment.
- Restore the directory data from backup into the test environment.

Test the Third-party Applications

All third-party applications running in the source environment (such as ERP systems, fax software, and meta-directory services) must be tested to prove that:

- The applications will function correctly in the target environment if you move the server they are running on to the target domain and re-assign permissions for target user accounts.
- Migrated users will still be able to work with the applications after permissions have been processed for target user accounts.

Depending on the results, you will need to decide whether to move the server running the application to target or to re-deploy the application there. In either case, additional planning is needed.

We recommend you contact each third-party software vendor to verify that they support the migration process or ask them to provide a migration procedure.

Test the Specific Hardware Update Procedures

It is recommended that you test all update procedures for your specific hardware devices (such as Network Attached Storage and Storage Area Network devices) in a test lab prior to starting resource update in the production environment.

Environment Preparation

The source and target environments must be properly prepared before starting any migration activities. [Environment Preparation Checklist](#) of this document contains a checklist of the tasks to prepare the environments.

Basic Migration Steps

[Pilot Migration](#)
[Overview of Migration Scenarios](#)
[Active Directory Migration](#)
[Exchange Migration](#)
[Active Directory and Exchange Migration](#)
[Rollback](#)

Pilot Migration

Since most lab environments are not true mirror images of production, there may be some differences in behavior when you start migration in production. Therefore, the best practice is to begin your migration with a *pilot migration*: select a pilot group of accounts that includes a wide spectrum of groups, users, and service accounts, and migrate that group first. This limits the user impact if unexpected issues occur. Do not limit the pilot to just IT or the migration project team members because they generally have modified configurations and won't represent a good cross-section of your users.

During the pilot migration, follow the procedures in the cookbook you prepared during the planning stage. If no problems occur, you can start migrating the rest of the environment. However, if a problem should arise, you should resolve it first, perform additional testing, update the cookbook to include the correct procedures, and then repeat the pilot.

Overview of Migration Scenarios

Migration Manager allows you to migrate multiple Active Directory domains and Exchange organizations at a time. Three basic migration scenarios are supported by Migration Manager:

- **Active Directory Migration:** You can use Migration Manager for Active Directory to migrate users and groups to another domain in the same or a different Active Directory forest. Depending on where the target domain resides, the migration scenario can be either *inter-forest* or *intra-forest*.
 - a. The *inter-forest* migration scenario assumes you need to migrate users and groups to another domain in a different Active Directory forest. Either Exchange Server is not installed in the source forest or for some reason you do not want to migrate Exchange data to the target.
 - b. The *intra-forest* migration scenario is used to migrate users and groups to another domain in the same Active Directory forest. If the source users already have Exchange mailboxes, you can simply reconnect these mailboxes to the new user accounts during migration.

- **Exchange Migration:** You can use Migration Manager for Exchange to migrate messaging and public folder data only while leaving the security accounts in the original Active Directory forests. In this case, users from one or several forests will have mailboxes in a separate Exchange organization. This deployment type is sometimes referred to as Exchange Resource Forest. This scenario is also applicable if Active Directory objects (such as user accounts, groups, and contacts) have already been migrated to the target domain by means of other migration tools, such as Active Directory Migration Tool (ADMT). In this case users log on to the target domain and you only need to migrate their Exchange data (such as mailboxes and public folders) to the target Exchange organization and decommission the source Exchange environment.
- **Active Directory and Exchange Migration:** In this scenario, you migrate both user accounts and mail data to the new Active Directory forest and Exchange organization. You need both Migration Manager for Active Directory and Migration Manager for Exchange for this scenario.

Active Directory Migration

The Active Directory migration scenario assumes you need to migrate only user accounts and resources (such as end-user workstations, file and print servers, and BackOffice servers) to another domain in the same or a different Active Directory forest. Either Exchange Server is not installed in the source forest or for some reason you do not want to migrate Exchange data to the target.

The Active Directory Migration scenario is illustrated in the figure below:

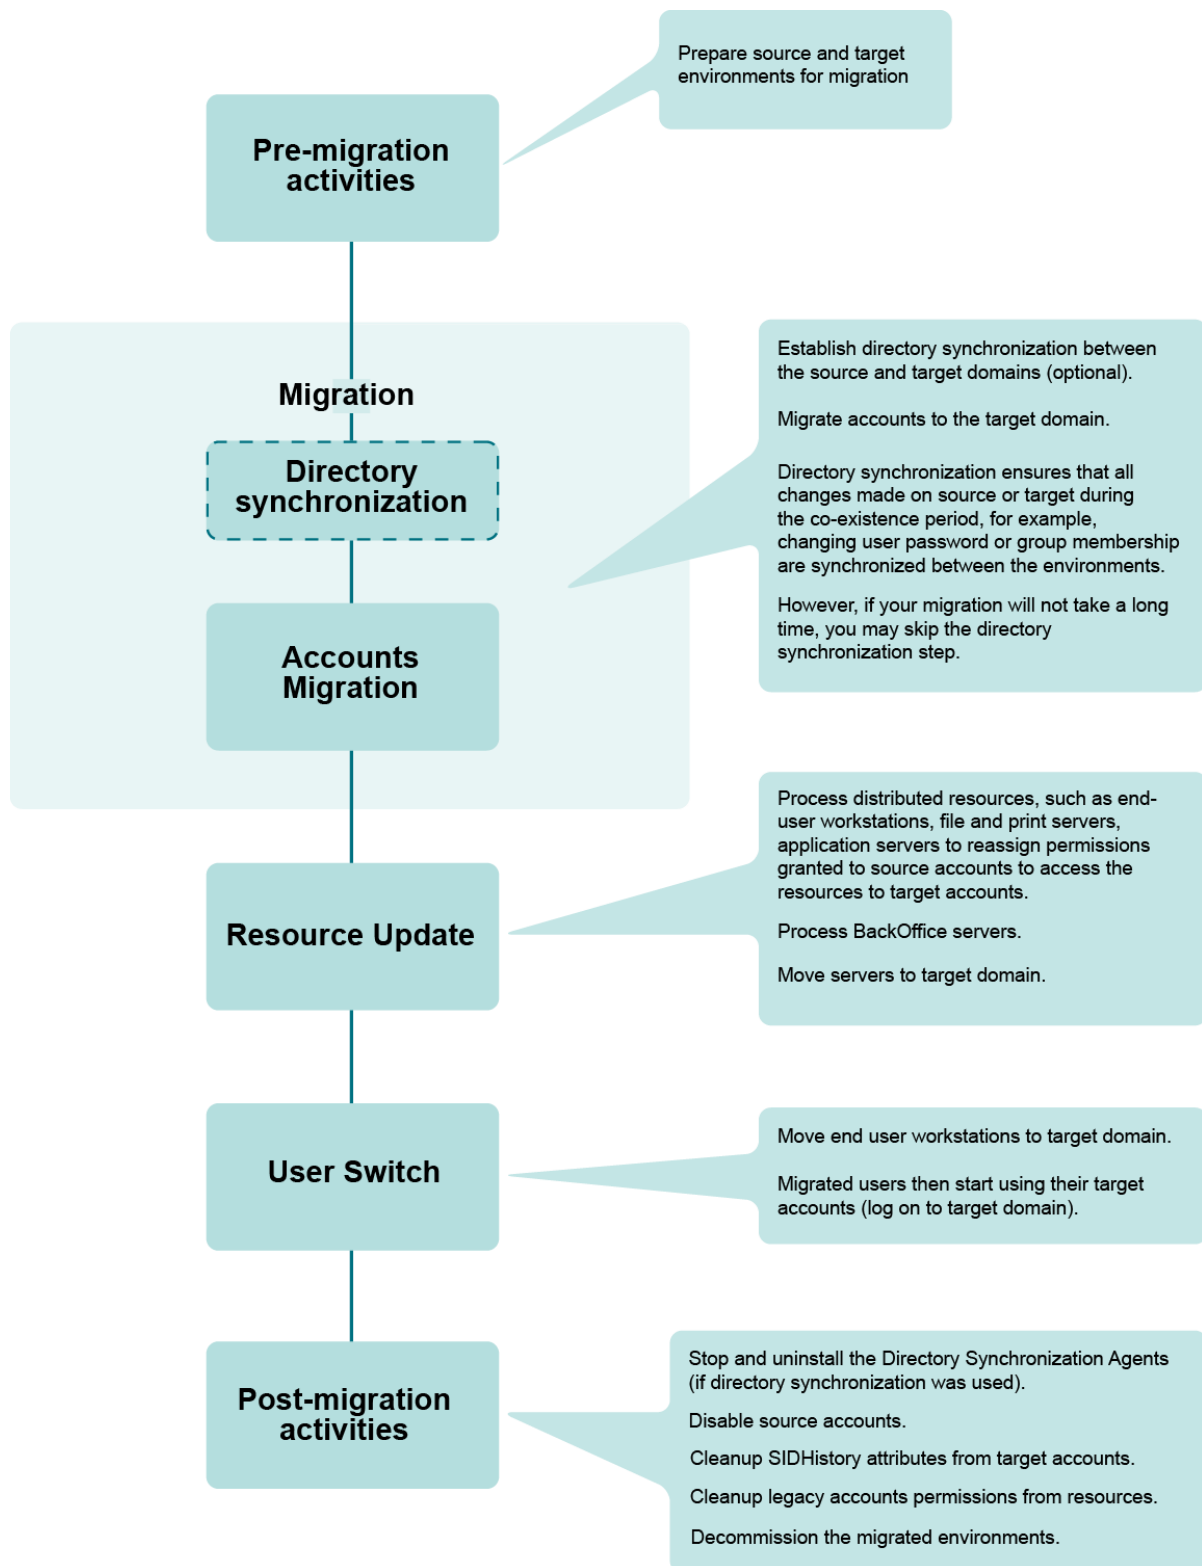


Figure 1: Overview of Active Directory migration process.

To migrate Active Directory, complete the following steps.

1. **Establish directory synchronization between the source and target domains (optional).** Configure the Directory Synchronization Agent to synchronize only object properties (including passwords and group membership) for all accounts within the specified synchronization scope. This ensures that all properties of source and target users are kept in sync during the co-existence period. If the coexistence period is short, you can skip this step.
2. **Migrate the directory.** Migrate accounts from the source to the target domain in migration sessions. You can delegate rights to perform the account migration to the other administrators in your environment.
3. **Process distributed resources.** Update distributed resources, such as end-user workstations, file and print servers, and application servers using Resource Updating Manager, adding permissions to the resources for the target users. When user workstations are updated, user profiles should also be updated, so the migrated users will get the same profile as the corresponding source users when they log on to the target domain for the first time. You can delegate rights to perform resource update to other administrators.
4. **Move end-user workstations and servers to the target domain.** This step is actually the user switch, because when you move a workstation to a target domain using Resource Updating Manager, the last logged-in domain on users' workstations is changed from the source to the target domain and thus users start to log in to the target domain. Move file and print servers and other application servers that have been processed by Resource Updating Manager to the target domain.

i | **NOTE:** Migration of the users can be performed in batches. If you choose this approach, repeat steps 2 through 4 for each group of users you migrate.

5. **Process BackOffice servers.** Update Microsoft BackOffice servers, such as Exchange, SQL, and SMS Server, using the corresponding processing wizards. You can delegate rights to perform BackOffice server update to other administrators.
6. **Move BackOffice servers to the target domain.**
7. **Stop directory synchronization (optional).** If you established directory synchronization between the source and the target domains in step 1, stop and uninstall the Directory Synchronization Agents.
8. **Disable the source accounts.** We recommend that you wait some time after disabling the source accounts to make sure that all users are using their target accounts before you proceed to step 9.
9. **Enable SID filtering.** After you enable SID filtering, wait some time to ensure that all target users can access the resources they used before the migration.
10. **Clean up SIDHistory from the target accounts.** Once all users are fine and have access to resources they had before, proceed to clean up SIDHistory.
11. **Clean up legacy account permissions from resources.** Note that cleanup is hard to undo. It is recommended that you clean up permissions only when you are sure that all users are using their target accounts for all applications and have no problems accessing resources.
12. **Clean up service attributes used for migration.**
13. **Decommission the migrated environments.**

Step 1. Establish Directory Synchronization (Optional)

The following are the important issues to consider when synchronizing directories:

- Whether you need to establish directory synchronization
- Whether you need one-way or two-way directory synchronization
- Whether you should allow the Directory Synchronization Agent to create the objects

Do You Need to Synchronize the Directories?

Depending on your project requirements, you may or may not use the ongoing directory synchronization capabilities. You can skip this step if you are going to migrate a small number of objects from one Active Directory forest to another and the coexistence period (when there are active user accounts in both the source and target environments) is expected to be short (for example, a weekend).

However, directory synchronization is required if any of the following applies:

- You need to migrate a large number of objects from one Active Directory forest to another.
- The objects and resources are migrated in groups (for example, by departments).
- The source and target environments must coexist during a period of time (that is when administration happens in both environments).
- You need to migrate user mailboxes from one Exchange organization to another. Refer to the **Exchange Migration** and **Active Directory and Exchange Migration** scenarios for more details on directory synchronization requirements when Exchange migration is involved.

We recommend that you establish directory synchronization between source and target domains to ensure that objects' properties are kept in sync and changes made to the objects in one directory are replicated to another directory.

Should You Choose One-Way vs. Two-Way Directory Synchronization?

In most cases, you only need one-way directory synchronization to be established between the source and the target domain and directed from source to target. This is the case when you migrate objects in groups to the target domain while continuing administration only on the source (make changes only to source objects) during the whole coexistence period. One-way directory synchronization ensures that all changes made to the source objects (such as passwords and group membership) that have been migrated to the target are replicated to the target directory.

However, sometimes two-way directory synchronization is needed. This is the case when two production directories are merged (for example, when two different companies merge), administration is performed on both sides simultaneously, and you want changes made to the objects on either side to be replicated to the other side.

Should You Have the Directory Synchronization Agent Create the Objects?

The Directory Synchronization Agent matches source and target objects by the criteria you specify. However, if the agent cannot find a matching object in the directory (no matching criterion is met), it can create one and then synchronize object properties.

The best practice for the Active Directory Migration scenario is to establish ongoing directory synchronization between the domains without object creation capability and to migrate the objects from the source to the target domain using migration sessions.

The other benefits of creating objects in the target domain by migrating them in migration sessions are:

- You can migrate the OU hierarchy in the migration session.
- You can specify the target OU in which the accounts will be created.
- You will be able to undo the migration session, if needed, by rolling back all changes made to the target directory.
- You can use import lists to rename and merge accounts on the fly.
- You can modify the accounts' attribute values to be applied in the target domain.
- You can test your migration session using test mode and then use this session as a template for the live migration sessions.
- You can granularly select objects by their group membership.

The properties of the migrated objects will then be kept in sync by the Directory Synchronization Agent.

Keep in mind that if you establish directory synchronization and do not allow the Directory Synchronization Agent to create objects in the target directory, this does not prevent the Directory Synchronization Agent from searching for target objects that match the source objects by the matching criteria specified. If the agent finds a matching object for the source object in the target domain, it will merge and synchronize these objects regardless of whether the object was migrated in a migration session.

However, if you are going to implement the **Exchange migration** scenario, you may want to allow the Directory Synchronization Agent to create the disabled mailbox-enabled objects in the target domain. For more information about the Exchange migration scenario, refer to the [Exchange Migration](#) topic.

Step 2. Migrate the Directory

Using Migration Sessions

Migrate the accounts from one source domain to another using migration sessions.

Depending on your project requirements, you can migrate accounts in one or more migration sessions. You can also delegate account migration to other administrators by creating and pre-configuring the delegated migration sessions. For more information about delegated migration, refer to the [Delegating Migration Activities](#) topic earlier in this document.

We recommend migrating user accounts and groups together in one session in order to transfer group membership to the target accounts during a session. However, if you plan to migrate user accounts and groups in portions using multiple sessions, we recommend you use the following order to successfully transfer group membership information:

1. Migrate user accounts in multiple sessions.
2. Migrate universal, local and global groups.

If you migrate groups before users, group membership may not be resolved if a user account that is a member of the source group is not yet migrated to the target domain. Additional steps will be required to update group membership on the target after the user account is migrated, such as re-migrating and merging groups or running full directory resynchronization (if directory synchronization is established for that pair of domains).

Adding SIDHistory

We also recommend that you always add the SIDHistory attribute to the target accounts. This greatly helps to ensure that there is no impact to end users during the co-existence period of two environments, making your migration smoother.

Migrating Other Active Directory Objects and Preserving OU Hierarchy

Migration Manager allows you to migrate not only user and group accounts, but also other Active Directory objects, including organizational units (OUs), contacts, print queues, shares, and computers. If you want to

preserve the existing OU hierarchy, select the OUs that you want to create in the target domain as well as objects to be migrated within these OUs.

Migrating Computer Accounts

We also recommend you select and migrate computer accounts in migration sessions. This allows you to specify the target OU where computer accounts will reside. When you later move computers from the source to the target domain, the existing (migrated) computer accounts will be used and no new computer accounts will be created in the default Computers OU in the target domain. You can also use these migration sessions to populate the computer list in Resource Updating Manager later with the computers to be processed. This eliminates the need to select individual computers that must be updated from all computers in the source domain.

Step 3. Process Distributed Resources

Distributed resources are end-user workstations, file and print servers, servers running IIS, scheduled tasks, and other services and applications. To ensure that resources will still be available to users when they start using their target accounts and when you have cleaned up SIDHistory, permissions granted to source accounts to access the resources must be re-assigned to the target accounts.

Service accounts and accounts used to run scheduled tasks must also be changed to the corresponding target ones to ensure that services and scheduled tasks will run correctly after the source accounts are disabled.

i **NOTE:** Service accounts should be handled separately since the server must be restarted to use the new account. You must coordinate this activity before disabling the source account or the service may stop running. You can use the **Services not Running as System Account** report to determine which accounts are used to launch services other than local system on the selected computers.

Update Methods

In order to re-assign user rights and permissions set to the objects, update local group membership, and change service and scheduled tasks' accounts, distributed resources must be updated. You can do this in any of the following ways:

- Interactively, by using Resource Updating Manager (recommended). In this case the updating agent is installed on each computer selected for processing and controlled from Resource Updating Manager remotely. This also allows you track processing status and results.
- Locally, by using the command-line updating tool Vmover.exe and INI files. You can run the update on a computer either directly, from a command line, or via a user logon script.

Items that Can be Updated

The following items can be updated:

- Local group membership
- User rights
- Services
- Scheduled tasks
- Registry
- Local profiles
- Roaming profiles
- Shares

- Printers
- File system
- IIS
- DCOM
- COM+

"Leave Source Accounts' Permissions" Option

We strongly recommend selecting the **Leave source accounts' permissions** option when updating the resources. This ensures that source users will still be able to access the resources after they have been processed. Then if you need to roll back and use your source accounts again, rollback will be immediate and therefore have no impact on source users. The permissions for these legacy accounts can be cleaned up after the migration is over (see step 10 below).

Factors to Consider

Here are the factors to keep in mind when doing the resource updating:

- The more objects the computer has (in most cases, this means the more files and folders), the longer it will take to process. Thus, processing a file server takes much longer than processing a workstation.
- Updating file system permissions requires a lot of disk access (I/O) operations and may slow the server for a period of time.
- Each computer in a set is processed by its own agent. Thus, all the computers are processed in parallel and it takes approximately the same time to process a dozen of workstations as a thousand.
- Expect about 10% of your workstations to require troubleshooting because they are offline, the Server service is not running, the domain administrators were removed from the local *Administrators* group, etc.

Accordingly, it is recommended that you create separate groups for end-user workstations and servers and that you process workstations first and then servers. See the Resource Updating Manager online help for details. You may also want to perform server processing during non-business hours to ensure that no users are affected by a possible server slowdown.

Creating a Computer List

In Resource Updating Manager you can create the computer list in either of the following ways:

- Using information of the Computer Browser service (running on the console where Resource Updating Manager is started).
- Using the domain controller information. This allows for creating the list of computers from the domain controller, and therefore it will not be affected if there are problems with the Computer Browser service or WINS.

Administrative Rights

Administrative rights over the resources are required for successful resource updating (see the *Quest Migration Manager - System Requirements and Access Rights* document). To obtain administrative rights you can do any of the following:

- Add the account you are currently logged on to the console with to the local Administrators group on each computer to be updated.

i **NOTE:** By default, if the computer has been joined to the domain, the Domain Admins group of the domain is a member of a computer's local Administrators group. You will get administrative access to the computer if the account you are using is a member of the source Domain Admins group.

- Use another account that is a member of a computer's local Administrators group to connect to the computer. In Resource Updating Manager you can specify the credentials for each domain.

Server Service

Server service is automatically installed when you install the File and Printer Sharing service on the computer. Server service must be running on a computer in order for it to be updated from either Resource Updating Manager or the command-line updating tool VMover.exe.

If for some reason (for example, due to your corporate policy) Server service is not allowed to run on the computers, you need to install or enable Server service temporarily, update the computers, and then disable or uninstall the service.

Client-side Firewalls

In Windows XP Service Pack 2, Microsoft introduced the Security Center, which includes a client-side firewall application. The firewall is turned on by default and configured to filter the packets sent to ports 137–139 and 445. These ports are used by the File and Printer Sharing service that must be installed and running on the computer to be updated.

In order to successfully update Windows XP Service Pack 2 computers from Resource Updating Manager, the File and Printer Sharing service must be added to the firewall Exceptions list and ports 137–139 and 445 must be unblocked.

If the File and Printer Sharing service was installed and running on a computer and a shared folder or printer were then created on that computer, these ports will remain open after you apply Service Pack 2 on Windows XP computer. However, if there were no shared folders or printers created prior to Service Pack 2 installation, these ports will be blocked. You need to make sure that ports 137–139 and 445 are not blocked by the firewall.

To install the File and Printer Sharing service:

1. Open the **Local Area Connection** settings as follows: right-click **My Network Places** icon on the desktop and select **Properties**. In the dialog that appears, right-click the **Local Area Connection** and select **Properties**.
2. In the **Local Area Connection Properties** dialog, click **Install**.
3. In the **Select Network Connection Type** dialog, select **Service** and click **Add**.
4. Select the **File and Printer Sharing for Microsoft Networks** in the **Select Network Service** dialog and click **OK**.

To unblock the ports used by the File and Printer Sharing service:

1. Start the **Windows Firewall** snap-in from **Control Panel**.
2. Check **File and Printer Sharing** in the **Exceptions** tab and click **OK**.

Command-line Updating

If for some reason you cannot use the agents to update computers from Resource Updating Manager, you can update the computers locally, using the command-line updating tool VMover.exe.

VMover.exe is located in the **Aelita Shared\Migration Tools\Resource Updating\Agent** subfolder of the Migration Manager installation folder (64-bit VMover.exe is located in the **x64** subfolder). The update can be performed either directly from the command line or via a logon script.

To perform the updates, VMover retrieves the source-target account pairs from the INI file. This file can be created in Resource Updating Manager.

When using the INI file, VMover will perform the updates for all accounts migrated by that time the INI file was created.

Refer to the *Quest Migration Manager for Active Directory Resource Processing Guide* for more details.

Centralized Resource Update

Use Resource Updating Manager for centralized resource update. Determine the processing scope, i.e. select the computers on which you want to update resources. For that, you should create collections and include computers you want to process (for example, all workstations or all servers). Each collection will contain the computers to be processed, and the set of processing settings. You can create as many collections as needed. Remember to obtain administrative rights over the resources. For details, refer to Resource Updating Manager help.

Delegated Resource Update

We recommend you use delegated resource updating in sites or resource domains where you cannot get administrative access to computers.

Delegated and decentralized resource updating is useful if computers to be updated are located across a slow WAN connection and therefore sending multiple agents, no matter how small, would consume too much of the available bandwidth.

You can delegate the resource updating tasks to the remote site or to other domain administrators who have the required level of access and are located within an area of good connectivity to the computers to be updated.

Resource updating can be delegated by exporting the INI file and performing resource updating using this file by running the updating tools in stand-alone mode or running resource updating from the command line

Delegating Resource Processing Tasks

To delegate a task to other administrators

1. In Migration Manager, right-click **Tasks** under the **Resource Processing** node, select **New Task**, and select the task type.

i **NOTE:** You can create the resource processing tasks for:

- Active Directory Processing
- Exchange Processing
- SQL Processing
- SMS Processing

If you select **Resource Updating Manager** to perform distributed resource updating, the Resource Updating Manager will start. However, the task of that type is not created in the **Tasks** container.

2. Depending on the task you selected, the appropriate task configuration wizard will start. Select the **Delegate the resource processing task** option in the **Configuration Mode** step.
3. Specify the desired re-permissioning options for the task in the **Re-Permissioning Options** step.
4. In the **Delegate Task** step, select the accounts to whom you want to delegate the task and specify the **Resource Admin** role for these accounts.
5. Save the task configuration. The task will then appear in the **Tasks** container.

After the task is delegated, the delegated administrator can start Migration Manager from his or her location, connect to the migration project, select the task in the tree he or she was delegated rights to, configure the task (for example, specify the resources to be processed by the task), and run the task.

Creating Task Setup Packages

The setup package is a standard MSI installation file. There are two types of the setup packages:

- Packages containing the task configuration only, which are for administrators with Migration Manager installed on their workstations
- Packages containing the task configuration and executables needed to run this task, which are for administrators without Migration Manager installed

To create an MSI installation file for the task

1. In Migration Manager, select the task under the **Resource Processing | Tasks** node and right-click the task.
2. Click **Create Setup** on the shortcut menu.
3. Specify the output folder for the package.
4. Select the type of the MSI installation package.
5. **Click Start.**

Delegating Resource Update Using INI Files

Delegated resource updating can be performed by delegated administrators using the INI files received from the migration administrator.

Updating User Profiles

A user profile consists of two parts: the key in the system registry and the folder on a hard disk that contains user-specific data, desktop settings, and a registry hive stored in NTUSER.DAT file. Depending on where user data is stored (on a local hard disk or server), user profiles can either be local or roaming.

When a user configured with a roaming profile logs on to a workstation for the first time, a local copy of the roaming profile is created on the workstation. NTUSER.DAT is copied from the roaming profile folder to the corresponding local one as well.

User profiles (profile folders and registry keys for both local and roaming profiles) get updated when you run distributed resource updating from Resource Manager and select both the **Local profiles** and **Roaming Profiles** options.

After processing, the profiles are shared between the source and target user accounts. As a result, target users will use the same profiles as the corresponding source users do.

i **NOTE:** You can use the **ExportProfile** utility from the Migration Manager Resource Kit to associate the current source user's profile with his or her future (migrated) account. However, this does not grant the new account the permissions needed to use the profile. This has to be done by running distributed resource updating. Refer to the *Quest Migration Manager for Active Directory Resource Kit - User Guide* for more details on the **ExportProfile** utility.

Enabling the “Cross-Forest User Policy and Roaming User Profiles” Policy

If the end-user workstations are running Windows 2000 SP4 or higher, you should enable the **Allow Cross-Forest User Policy and Roaming User Profiles** policy before you move the workstations to the target domain. This is to allow users to use roaming profiles stored on a server that is still a member of the source domain.

You can configure this policy either locally on the client workstation or by using a domain or organizational unit-based Group Policy object (GPO). To do this locally on a workstation, complete the following steps:

1. Log on to the computer as a user with administrator rights.
2. Click **Start**, click **Run**, type **gpedit.msc**, and then click **OK**.
3. Double-click Computer Configuration, double-click Administrative Templates, double-click System, and then click **Group Policy**.
4. In the right pane, double-click **Allow Cross-Forest User Policy and Roaming User Profiles**.
5. Click **Enabled**, click **Apply**, and then click **OK**.
6. Quit the Group Policy tool.
7. Allow sufficient time for the computer policy to be automatically updated, or update it yourself in Windows 2000 by running the following command in the command line:
 - a. **secedit /refreshpolicy machine_policy**
 - b. In Windows 2003 and Windows XP, use the **gpupdate** command.

For more details on user policies, refer to *Microsoft Knowledge Base article 823862* at <http://support.microsoft.com/default.aspx?scid=kb;en-us;823862>.

Troubleshooting User Profile Update

Many system and service processes work on the workstation on behalf of users (for example, printer drivers and virus scanner services). If such a process does not properly close handles to the user profile hive it used, then the profile cannot be correctly unloaded when the user logs off. In addition, failure to close handles causes a problem with updating of the user profile: if the profile is locked by any service during processing, then the user may get a brand new profile when he or she logs on to the system with the target account for the first time after the workstation has been processed.

The **User Profile Hive Cleanup Service (UPHClean)** by Microsoft is intended to help troubleshoot such issues. For more information about UPHClean, refer to the following documents:

- Microsoft Knowledge Base article 837115, "Troubleshooting profile unload issues," at <http://support.microsoft.com/default.aspx?scid=kb;en-us;837115>
- The UPHClean readme file, available at <http://download.microsoft.com/download/a/8/7/a87b3d05-cd04-4743-a23b-b16645e075ac/readme.txt>

To download UPHClean, use the following link:

<http://www.microsoft.com/downloads/details.aspx?FamilyId=1B286E6D-8912-4E18-B570-42470E2F3582>

Update Cluster Servers

Migration Manager is capable of re-permissioning a Microsoft cluster. However, it requires a more involved procedure than what is required by non-clustered servers. Refer to the *Cluster Server Migration* section of the *Quest Migration Manager for Active Directory—Resource Processing Guide* for a detailed procedure on how to migrate Microsoft cluster servers with Migration Manager.

Resource Updating Troubleshooting

Workstations

After you run the updating session, some computers you specified may not be successfully updated. The computer may be turned off or inaccessible by the network from your location (network connectivity problems),

or you may not have enough rights to connect to it, and so on. Resource Updating Manager allows you to sort computer objects by last error status. If a computer was processed successfully, the last error field is empty and the computer can be removed from the list. Otherwise, you will see error descriptions. The typical errors are:

- **Access is denied**—You get this error when you do not have administrative rights over the server or workstation. Refer to the **Administrative Rights** section in the [Step 3. Process Distributed Resources](#) topic.
- **The network path was not found**—The cause is probably one of the following:
 - The computer is turned off.
 - It cannot be accessed by the network from your location due to network connectivity problems.
 - It cannot be found by name (the NetBIOS computer name cannot be resolved into the address due to DNS/WINS issues).
 - No computer with that name exists.
 - Firewall software is installed and enabled on a computer (such as Windows Firewall in Windows XP Service Pack 2) that prevents connecting to the computer.
 - It is a laptop computer whose owner is currently out of the office.

You need to diagnose each situation separately. Use the `ping <ComputerName>`, `ping <Computer_IP_Address>` and `nslookup` commands to determine whether you have a name-resolution problem or whether the computer is turned off or cannot be accessed by the network from your location.

- **The Server service is not started** – Start the Server service and try to process the computer again or follow the command-line updating procedure to update the computer locally.

Servers

We recommend you update servers during non-business hours to ensure that end users will not be affected by a possible decrease in server performance.

After adding servers to the Resource Updating Manager and starting the updating, wait for all computers to finish even if some have errors. Once the processing is finished, use the sort feature on the column headings to sort computers by problem, fix the problems, and start the updating again for only those computers.

Step 4. Move End-user Workstations and Servers to the Target Domain

We recommend you first move end-user workstations to the target domain after they have been processed and then move the servers.

Move End-user Workstations to the Target Domain

This step is called the user switch. We recommend you move user workstations using Resource Updating Manager. In this case, after a computer is moved to the target domain, the default logon domain name that is displayed in Windows logon dialog is changed to the target domain name. This helps a user log on to the target domain and start working with his or her target account without having to select the target domain manually in the Windows logon dialog.

! **CAUTION:** You should always make sure that user profiles have been processed before you move their workstations to the target domain. If you selected the **Change last logged-in domain to the target domain** option when moving computer accounts, the default logon name will be changed to the target domain name on each computer that has been moved. If user names were not changed during the migration, most probably users will not notice the change and will log on to the target domain. If their profiles have not been updated yet, they will get new profiles.

i **NOTE:** The **ChangeProfile** utility from the Migration Manager Resource Kit associates the profile of users currently being migrated with their previous (source) accounts. It can help to link a target user account to the old profile. Refer to the *Quest Migration Manager for Active Directory Resource Kit - User Guide* for more details.

Laptop Users and Cached Credentials

Many users now work with portable laptop computers. When such users are in the office, they plug their computers into the local area network (LAN) and log on to the domain using their credentials. These credentials are then automatically cached by the system on the laptop computer. These users may also often work out of the office, not connected to the LAN. In these cases they log on to the system with the same username, password, and domain name as they used in the office. The system authenticates the user using the credentials that were cached.

However, if you move a computer to another domain and then reboot it, all *cached credentials are automatically cleared by the system on that computer after reboot*. Thus, a user is no longer able to log on using the cached credentials. The system behaves as follows:

- Until the computer is rebooted (after the move), a user can use cached credentials to log on to the source domain.
- Once the computer is moved and rebooted, attempts to log on using cached credentials will be unsuccessful.
- Accordingly, we advise doing either of the following:
- Reboot laptop computers while the user is in the office, connected to LAN. Then he or she can log on to the target domain, and the new credentials will be cached.
- Hibernate laptop computers rather than switch them off before taking them home. At home the users can restore the laptop from hibernation and successfully log on to the domain using cached credentials. The next day at work, the laptop users can plug the machine in the network and reboot. Once they are logged on to the target domain, their new credentials are cached and can be used offline.

However, if the cached credentials have already been cleaned and the user is at home, the or she can still log on to the target domain by selecting the **Log on using dial-up connection** option in the Windows logon dialog and connecting to LAN via dial-up.

Locked Workstations

When a workstation is moved to another domain, the list of trusted domains (the domains trusted by the domain to which this workstation belongs) that is stored in system registry is automatically cleared by the system.

An issue may occur if a user workstation was locked when it was moved to the target domain and if the user tries to unlock it right after it is moved, without waiting for the domain list to be rebuilt. In this case, when pressing **Alt-Ctrl-Del**, the user will only be able to choose from the local workstation and the target domain to log on. To resolve this, the user should press **Esc** and then press **Alt-Ctrl-Del** again; the trusted domain list is rebuilt and the last logged-in domain (the default domain) setting is returned to the source domain. The user is then able to unlock the workstation.

Move Servers to the Target Domain

Next, move file and print servers and servers running scheduled tasks using Resource Updating Manager. For procedures for moving servers running IIS and other applications and services, refer to the product's documentation. Be sure to test these procedures first in your test environment.

Moving Cluster Servers

To move a cluster server whose nodes are all member servers of some domain to a different domain, select all the nodes in Resource Updating Manager and move them simultaneously. After a couple of minutes all nodes and the virtual cluster server will appear in the new domain.

CAUTION: Always move all cluster nodes to the new domain simultaneously. Do not move a virtual cluster server to the new domain.

NOTE: The Cluster service account is not changed when moving a cluster server to another domain.

Step 5. Process BackOffice Servers

Microsoft BackOffice includes SQL, Exchange, and SMS server products. To update BackOffice server permissions for the target accounts, use the following processing wizards:

BackOffice Server	Wizard	Description
Exchange Server	Exchange Processing Wizard	Updates client and administrative permissions on mailboxes, public folders, and all other Exchange objects.
SQL Servers	SQL Processing Wizard	Automatically detects the SQL Server version and performs the updates accordingly.
SMS/SCCM	SMS Processing Wizard	Updates Microsoft Systems Management Server and Microsoft System Center Configuration Manager permissions to reflect the migration.

NOTE: For details on supported backoffice server versions, refer to *Processed Platforms* section of the *System Requirements and Access Rights* document.

For more information about the processing wizards listed above, refer to the *Quest Migration Manager for Active Directory Resource Processing Guide*.

Centralized BackOffice Server Updating

To update BackOffice servers from a central location, in Migration Manager select Resource Processing from the Tools menu and select the appropriate wizard.

Delegated BackOffice Server Updating

The best practice for delegating BackOffice server update is to create a resource processing task and delegate it to the other administrators. In addition, you can create a setup package for the task and send it to the administrator, who installs the package and runs the task at his or her location. For more information on these techniques, refer to the [Step 3. Process Distributed Resources](#) topic.

Alternatively, you can update BackOffice servers using INI files.

Step 6. Move BackOffice Servers to the Target Domain

Microsoft BackOffice includes SQL, Exchange, and SMS server products. The recommendations for them are different, and are explained in the sections below.

SQL Servers

SQL servers can be moved to the target domain easily after they have been processed. You can do this manually or from Resource Updating Manager.

Exchange 5.5 Servers

Instead of moving Exchange 5.5 servers to the target environment, we recommend using Quest Exchange Migration Wizard to migrate Exchange 5.5 to Exchange. For more details on Exchange migration scenarios, refer to the *Quest Exchange Migration Wizard—User's Guide*.

Exchange Servers

Use Migration Manager for Exchange to migrate messaging system from the source Exchange organization to the target. After the inter-org Exchange migration is completed, you can decommission the old Exchange environment.

SMS Servers

Due to SMS implementation complexity, we recommend you re-deploy SMS in the target environment instead of moving the SMS servers.

If you do choose to move your BackOffice servers to another domain, refer to the procedures described in the Microsoft BackOffice documentation and in the Microsoft Knowledge Base for more information.

Step 7. Stop Directory Synchronization (Optional)

If you established directory synchronization in step 1, stop and uninstall the Directory Synchronization Agents.

Step 8. Disable Source Accounts

It is good practice to disable the source accounts at this stage and thus make sure that all users are logging on with their target accounts. We recommend that before you take this step, you wait some time to make sure that all users are already using their target accounts.

Step 9. Enable SID filtering

After the distributed resources and BackOffice servers have been processed, enable SID filtering between the source and the target domains.

After SID filtering is enabled, wait some time to ensure that all target users can access the resources they used before the migration. If after enabling SID filtering some users cannot access the resources, turn SID filtering off, process any skipped resources, and turn it on again.

We recommend you enable SID filtering prior to SIDHistory cleanup to verify that all resources were re-permissioned correctly and users can access resources with their target accounts not using SIDHistory.

Refer to the [SID Filtering](#) topic for the procedures on how to enable SID filtering.

Step 10. Clean Up SIDHistory Attributes

SIDHistory cleanup is done by Active Directory Processing Wizard, which is started from Migration Manager.

! **CAUTION:** Changes have probably been made to permissions, service accounts, group membership, etc. on resources since resource processing was last executed. We recommend you update distributed resources and BackOffice servers one more time before you clean up SIDHistory to make sure that all permissions, service accounts, and group membership are up to date.

i **NOTE:** If after SIDHistory cleanup some users cannot access the resources, SIDHistory can be re-applied back to the accounts that lost access only by re-migrating and merging the accounts. This will give you time to check whether the resources were processed correctly for these accounts, fix the problem, and clean up SIDHistory again.

However, by turning on SID filtering with users losing access to resources it is much easier to disable SID filtering and process resource skipped than it is to restore SIDHistory.

After SIDHistory is cleaned up, wait some time to ensure that all target users can access the resources they used before the migration.

Step 11. Clean Up Legacy Account Permissions

Cleanup of legacy account permissions is performed by the same wizards used to update resources.

Cleanup is hard to undo, so it is recommended that you clean up permissions only when you are sure that all users are using their target accounts for all applications and have no problems accessing resources.

Step 12. Clean Up Service Attributes Used for Migration

Cleanup of service attributes used by the Directory Synchronization Agent during migration and synchronization can be performed by using the **Active Directory Cleanup Utility for Quest Migration Manager**, which is included in the Migration Manager for Active Directory Resource Kit.

Step 13. Decommission the Migrated Environments

This is the last step of the migration process. If all previous steps were successful, you can switch off your source domain controllers and re-use the freed-up hardware.

! **CAUTION:** Do not switch off or demote your source domain controllers if there are servers (Exchange, SQL, or other servers running specific applications and services) that are still members of the source domain. Perform decommissioning only when you are sure that there are no member servers or workstations currently accessed by users left in the source domain.

Exchange Migration

In multi-forest Active Directory deployments, users from several forests might have mailboxes in one Exchange organization. This deployment type is sometimes referred to as *Exchange Resource Forest* or *Multiple Forests/Single Org*.

! **CAUTION:** If you are performing calendar synchronization on Microsoft Exchange 2007, make sure that the Public Folder database exists in your Exchange 2007 environment.

The main characteristic of such deployments is that users have mailboxes that are not in the forest in which they get authenticated. Thus, security directory is separate from the Exchange directory.

Migration Manager supports migration and deployment of such configurations. The product will migrate the Exchange org in such a way that users get switched to the new messaging system while remaining in their existing forest from a security perspective.

However, this scenario described below can be used also when Active Directory migration has already been completed; that is, Active Directory objects and resources have already been migrated from the source to the target forest by means of the other migration tools, such as Microsoft Active Directory Migration Tool (ADMT), and all users already log on to the target domain. In this case also, only Exchange data must be migrated from the source to the target Exchange organization.

The Exchange Migration scenario is shown schematically in the figure below:

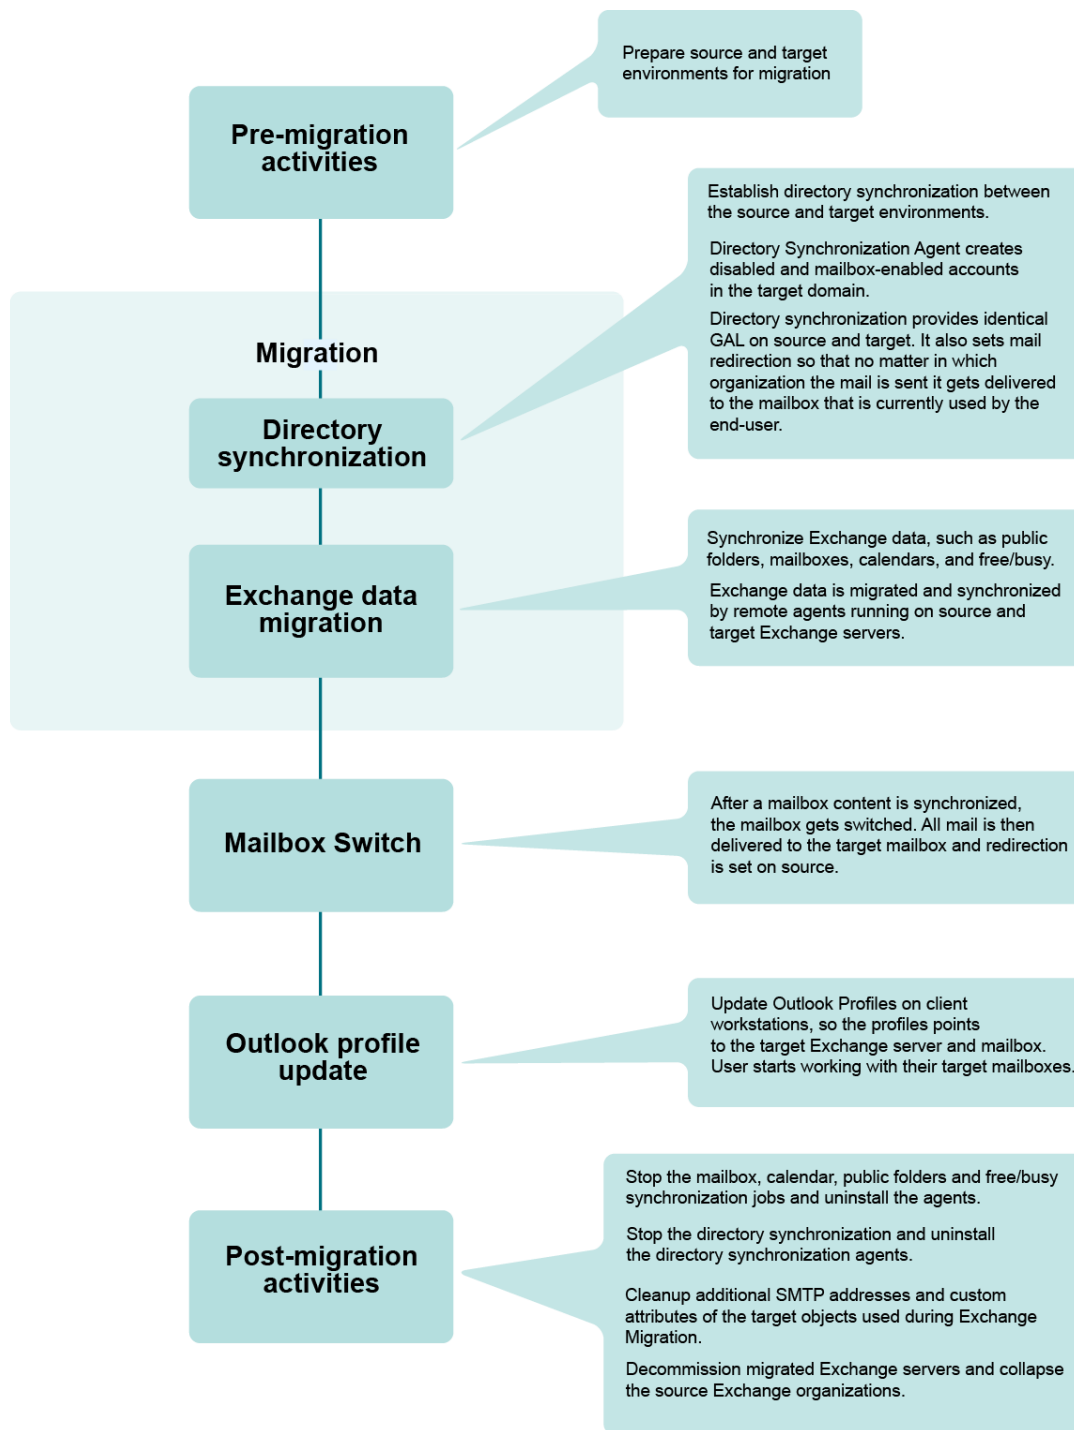


Figure 2: Overview of the Exchange migration process.

To migrate Exchange data, complete the following steps:

1. Establish directory synchronization between the source and target domains. Configure the Directory Synchronization Agent to create disabled and mailbox-enabled user accounts in the target domain.

Directory synchronization ensures that account properties and Global Address Lists (GALs) are identical in the source and target organizations. Directory synchronization also sets mail redirection so that mail is delivered to the mailbox currently used by the end user, regardless of which organization the mail is sent from.

2. **Start public folder and calendar synchronization.** Establishing public folder synchronization ensures that changes made in one organization get replicated to the other, so users can share the same public folder space. Migration Manager also allows you to set calendar synchronization independently of mailbox migration. That way you can ensure that calendar information is also available for any user in any organization.
3. **Establish free/busy synchronization (optional).** Free/busy synchronization enables users to schedule common activities. Migration Manager can synchronize free/busy information independently from other data and thus make sure that the information gets updated as close to real-time as possible.
4. **Synchronize mailbox data.** When mailbox synchronization is launched, Migration Manager starts transferring the source mailboxes' content to the target mailboxes and synchronizing mailbox permissions.
5. **Switch to the new Exchange mailboxes.** When a mailbox is switched, Migration Manager sets redirection to the opposite direction: all new mail sent to the old (source) mailbox is automatically forwarded to the new mailbox in the target organization. Migration Manager also marks the mailbox in a way that initiates the Outlook profile update at the user's next logon.
6. **Update Outlook profiles.** Migration Manager is shipped with the Profile Updating Utility (EMWProf), which handles Outlook profile update. After update the profile points to the target Exchange server and user mailbox. The majority of the properties stored in a profile are also get updated.
7. **Change the mail exchanger or alias records.** Switch incoming SMTP traffic to the target Exchange bridgehead server when about 50 percent of the users have had their mailboxes switched to the target in order to optimize routing.
8. **Stop and uninstall the synchronization agents.** When all mailboxes are migrated and switched, you can stop the synchronization. The following agents should be stopped and uninstalled:
 - Directory Synchronization Agents
 - Mailbox Synchronization Agents
 - Synchronization Agents
 - Public Folder Synchronization Agents
 - Free/Busy Synchronization Agents (if they were used)
9. **Clean up the additional SMTP addresses and service attributes.** Clean up the additional SMTP addresses set for redirection purposes and the custom attributes of the target objects used during Exchange migration.
10. **Decommission the migrated environments.**

Step 1. Establish Directory Synchronization_EX

You always need to establish directory synchronization when you migrate user mailboxes from one Exchange organization to another. Configure the Directory Synchronization Agent to create disabled and mailbox-enabled user accounts in the target domain.

The initial directory synchronization creates new user accounts in the target domain and a mailbox for each user corresponding to a source mailbox. This should be completed for all source mailboxes you want to migrate to the new Exchange organization before any other activity is started, for the following reasons:

- Directory synchronization is required to maintain a common Global Address List (GAL) for source and target Exchange organizations.
- Synchronization of client permissions for mailbox folders and public folders depends upon the existence of mailbox-enabled users in the target domain.
- Mailbox and calendar synchronization require that each source user has a corresponding mailbox in the target Exchange organization.

When the Directory Synchronization Agent creates a disabled account in the target domain that corresponds to the source user account, it automatically sets the source user account as the Associated External Account (i.e., the SID of the source user is added to the `msExchMasterAccountSID` property of the target user). This ensures that source users will be able to access all target Exchange resources with the old (source) accounts.

If security accounts have been created in the target domain prior to Exchange migration (Active Directory migration has been completed previously), you should configure the Directory Synchronization Agent to search for the matching objects in the target domain for each source object within the specified synchronization scope. The following matching rules can be used:

- **Account name** – If the `sAMAccountName` attributes of the source and target object are the same, the objects will be matched.
- **SIDHistory** – If the `SIDHistory` attribute of an object from one directory contains the SID of an object from another, the objects will be matched.
- **E-mail** – This matching rule can be used if target objects were created mail-enabled. This is, for example, if Quest Collaboration Services was used and stub mail-enabled accounts were created by that product in the target domain. For mail-enabled objects, if source and target object have the same primary SMTP address, the objects will be matched.

All three matching rules are turned on by default. We recommend you select only those rules that are relevant for your previous migration and switch off the other rules that do not apply to your situation. For example, if you have migrated accounts and added `SIDHistory` to the target accounts, use the `SIDHistory` matching rule. If you have migrated accounts without `SIDHistory` but did not change the account names (source and target accounts have the same names), use matching by account name.

Directory synchronization sets mail redirection so that mail is delivered to the mailbox currently used by the end user, regardless of which organization the mail is sent from. The additional SMTP addresses are used for redirection. These addresses are generated upon the template you specify when you configure directory synchronization, and are automatically added to the source and target mailboxes.

Directory synchronization also ensures that account properties and Global Address Lists are identical in the source and target organizations.

Once the initial synchronization is completed, you can proceed to Step 2. However, directory synchronization should continue to run until the last user is migrated to the target Exchange organization. This ensures that changes made by the administrators in the source or target directory are automatically propagated to the other directory.

For more information about directory synchronization, see the [Directory Synchronization](#) topic.

Step 2. Start Public Folder and Calendar Synchronization

Synchronize both public folder and calendar data to the target organization before the first user logs on to the target mailbox. This will make the transition transparent.

Perform initial public folder synchronization before starting mailbox migration (Step 3). It may take a few days to complete if there is a large amount of public folder data. By starting this before mailbox migration, you can ensure that all public folder data is available before the first user logs on to the target mailbox.

! CAUTION: In Exchange 2003, the size limit for public folder items is set by default to 10 MB. It is recommended that you increase or remove the size limits before you start the synchronization. Otherwise, if you set a two-way synchronization, you may lose the attachments larger than 10 MB on the source side.

It is best to configure mailbox synchronization but not start it, and then use mailbox synchronization collection members to populate a calendar synchronization collection in a calendar synchronization job. Then start calendar synchronization. Calendar data synchronizes much faster than mailbox data, so if calendar synchronization starts before or at the same time as the first mailbox synchronization, it normally completes for all mailboxes before the first mailbox is switched. That way you can ensure that calendar information is available for any user in any organization, regardless of whether the user's mailbox is switched.

i NOTE: If a mailbox is synchronized in a Remote Users Collection, the whole mailbox content, including the Calendar folder, is transferred to the target server at once by the Mail Source Agent. The mailbox is automatically switched as soon as its content is transferred.

The Calendar Synchronization Agent logs on to all mailboxes in the calendar synchronization job and does not log off until it finishes processing. The Mail Source Agent skips the logged on mailboxes during processing of a Remote Users Collection. Thus, the calendars of the mailboxes migrated in the Remote Users Collections should not be synchronized until these mailboxes are migrated and switched by the Mail Source Agent.

Permissions are synchronized together with public folder and calendar data, so the users have uninterrupted access to the data they need. The ongoing two-way synchronization of public folders and calendars lets users in different organizations share the same public folder data and resource mailboxes.

For more information about public folder synchronization, see the [Directory Synchronization](#) topic.

For more information about calendar synchronization, see the [Calendar Synchronization](#) topic.

! CAUTION: Deleted public folders cannot be migrated again. If you want to do a pilot public folder migration, create a test public folder and use this folder for your experiments. Avoid doing pilots on the production folders. That way you ensure that the pilot project will not interfere with the production migration in the future.

Step 3. Establish Free/Busy Synchronization (Optional)

Establish free/busy synchronization to enable users to schedule common activities. Migration Manager can synchronize free/busy information independently from other data and thus make sure that the information gets updated as closely to real-time as possible.

Step 4. Synchronize Mailbox Data

Mailbox synchronization can be started at any time after the initial directory synchronization has been completed. However if you want users to have full collaboration capabilities and a consistent view in source and target Exchange organization, public folder and calendar data should exist on the target before either the first mailbox is switched to the target Exchange server or a new user and mailbox are created in the target.

To make migration transparent for all users, Migration Manager provides the Remote Users Collection feature, which allows you to preserve the offline folder (OST) files for remote and laptop users who typically work offline and occasionally connect to their Exchange mailboxes. For the mailboxes processed within a Remote Users Collection, Migration Manager keeps the OST file of the source mailbox and makes the target mailbox use this file after migration. For more information about Remote Users Collections, see the [Offline Folder \(OST\) Files and Remote Users Collections](#) topic.

Step 5. Switch to the New Exchange Mailboxes

When a mailbox is switched, Migration Manager sets redirection to the opposite direction: all new mail sent to the old (source) mailbox is automatically forwarded to the new mailbox in the target organization. Migration Manager also marks the mailbox in a way that initiates the Outlook profile update at the user's next logon.

Migration Manager can automatically switch mailboxes as they become synchronized. Alternatively, administrator can monitor the synchronization status and schedule the switch for a particular date and time.

Step 6. Update Outlook Profiles

Migration Manager can automatically switch a mailbox once its content is fully synchronized, or a user can be switched manually by the administrator. Once a user's mailbox is switched, all new mail arrives at the Exchange mailbox only.

Public folders and calendars continue to be synchronized throughout the migration until all users are using their target mailboxes. However, it is recommended that the initial public folder and calendar synchronization complete before the first user mailbox is switched. This will ensure that users in the source and target organizations see the same calendar and public folder data.

Once a user's mailbox is switched, the Outlook profile of that user should be updated by the Client Profile Updating Utility (EMWProf). Normally, administrators should include this utility in user logon scripts for the migration period to ensure that Outlook profiles for the switched mailboxes are updated automatically. In that case, EMWProf will start at each user logon, scan for Outlook profiles, and check whether any of the profiles point to switched mailboxes. If a switched mailbox is found, EMWProf updates the profile so that now it points to the new mailbox on the target server. If no profiles require update, EMWProf quits.

If a user has mail profiles on multiple machines, such as a laptop and a desktop, the profile on each machine will be updated when the user logs on to the machine.

For more information about switching mailboxes and updating client profiles, see the [Mailbox Switch and Profile Update](#) topic.

Step 7. Change the Mail Exchanger or Alias Records

Initially, all incoming SMTP mail is received by the source Exchange bridgehead server. For mailboxes that have been switched, it is automatically redirected to the target Exchange bridgehead server.

In order to optimize routing, you should switch incoming SMTP traffic to the target Exchange bridgehead server when about 50 percent of the users have had their mailboxes switched to the target. To switch incoming SMTP traffic, change the name server Mail Exchange (MX) and/or Alias (A) records for Internet mail.

After this change, all SMTP mail will arrive in the target mailboxes. For mailboxes that have not yet been switched, SMTP mail is automatically redirected to the source Exchange bridgehead server since redirection for such mailboxes is set on target.

Step 8. Stop and Uninstall the Synchronization Agents

After all the data is transferred to the target servers and all mailboxes are switched, you can stop and uninstall the synchronization agents. The agents are:

- Directory Synchronization Agents
- Mailbox Synchronization Agents
- Calendar Synchronization Agents
- Public Folder Synchronization Agents
- Free/Busy Synchronization Agents (if they were used)

Step 9. Clean Up the Additional SMTP Addresses and Service Attributes Used for Migration

Clean up the additional SMTP addresses set for redirection purposes and the custom attributes of the target objects used during Exchange migration.

You can use standard Windows utilities for bulk LDAP operations, such as CSVDE, or third-party tools to clean up service attributes used during migration and directory synchronization.

Cleanup of service attributes can also be done by the **Active Directory Cleanup Utility for Quest Migration Manager**, which is included in the Migration Manager for Active Directory Resource Kit.

Step 10. Decommission the Migrated Environments

Decommission the migrated source Exchange servers.

Active Directory and Exchange Migration

Starting with Exchange 2000, Microsoft effectively merged Exchange directory and Active Directory. In the majority of Active Directory and Exchange deployments, users get authenticated in the same directory that hosts their Exchange system.

Therefore, moving users to another forest includes moving both their security accounts and mailboxes. That is why you need both Migration Manager for Active Directory and Migration Manager for Exchange for this migration scenario.

There are two possible ways to migrate Active Directory and Exchange:

- Migrate Active Directory and Exchange data in parallel.
- Migrate Exchange first (implement the Exchange Resource Forest scenario), and then migrate Active Directory.

Parallel Active Directory and Exchange Migration

This scenario is a combination of the Active Directory Migration and Exchange Migration scenarios described above. It allows you migrate Active Directory accounts, resources, mailboxes, and public folders in parallel.

The Active Directory Migration and Exchange scenario is shown schematically in the figure below:

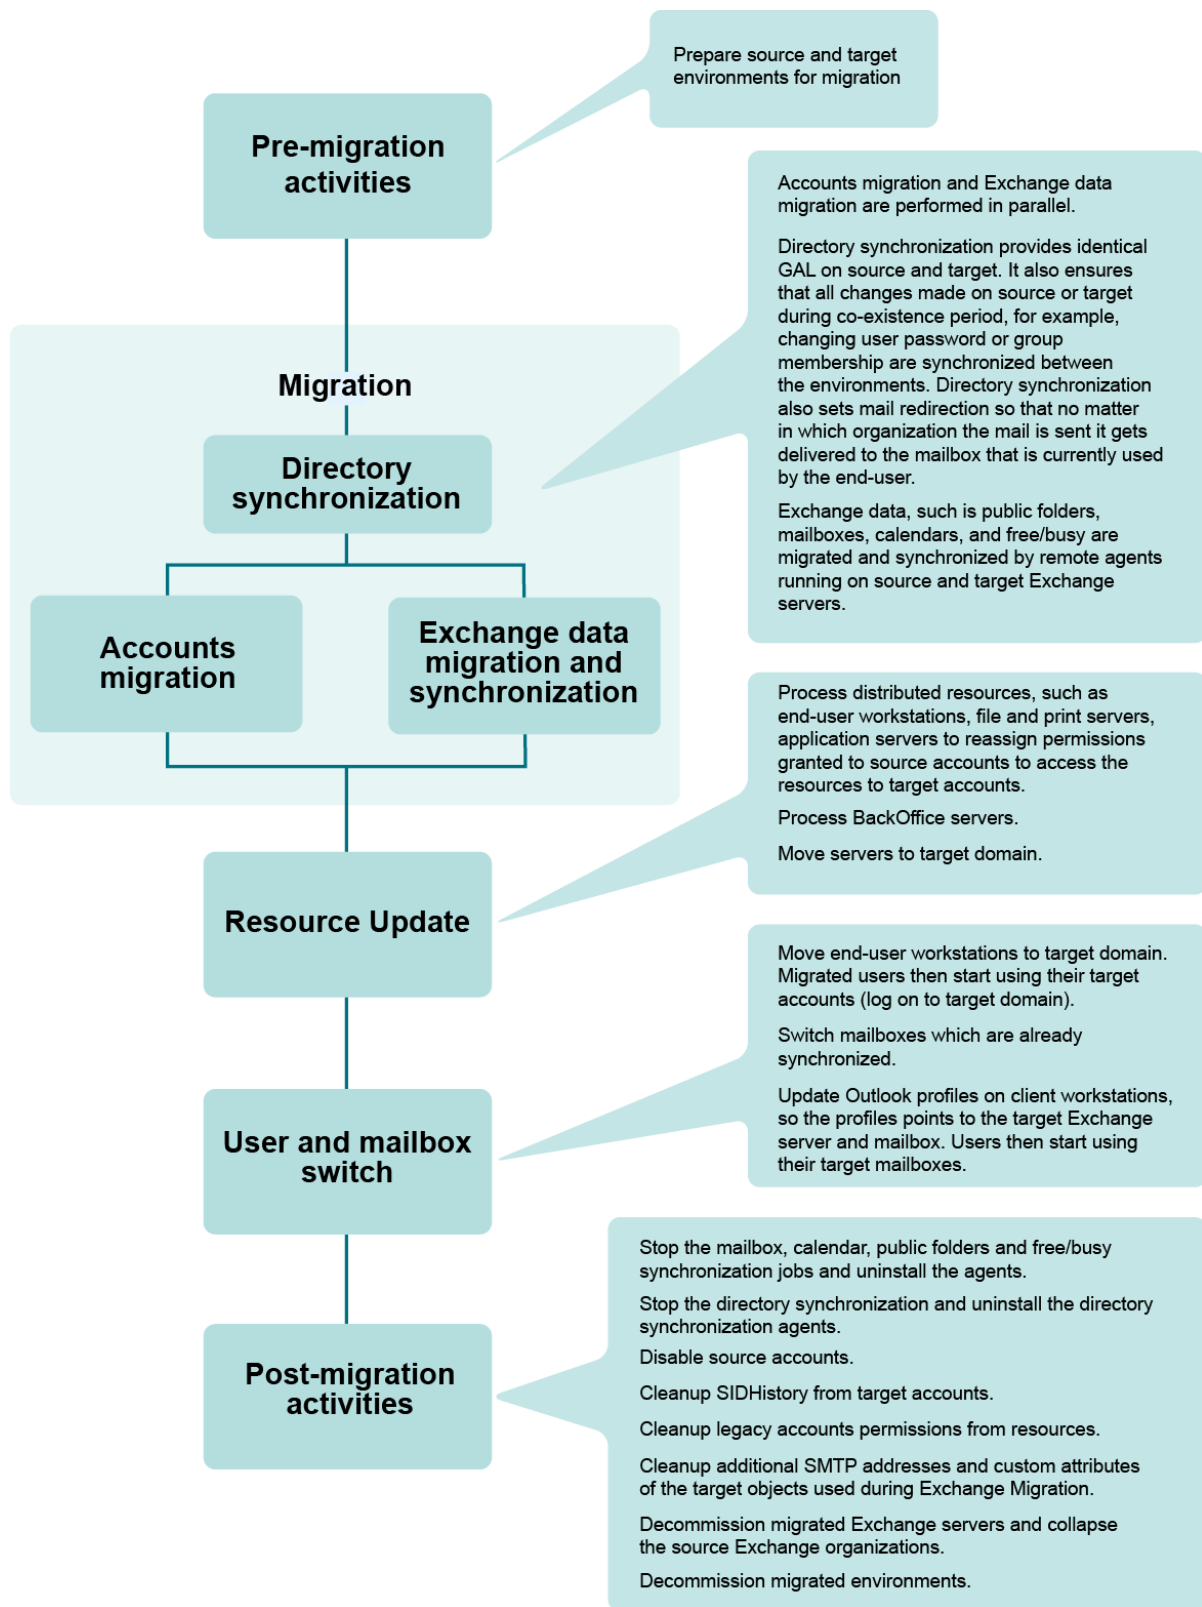


Figure 3: Overview of the parallel Active Directory and Exchange migration process.

To migrate Active Directory and Exchange in parallel, complete the following steps:

To migrate Active Directory and Exchange in parallel, complete the following steps:

1. Establish directory synchronization between the source and target domains.

- Configure the Directory Synchronization Agent to synchronize source and target objects and mailbox-enable target user accounts that do not have mailboxes yet.
- Allow the Directory Synchronization Agent to create new objects in the staging OU if the matching object cannot be found in the target domain.
- Configure the Directory Synchronization Agent to create disabled user accounts in that OU.

Directory synchronization ensures that account properties and Global Address Lists (GALs) are identical in the source and target organizations. Directory synchronization also sets mail redirection so that mail is delivered to the mailbox currently used by the end user, regardless of which organization the mail is sent from.

- 2. Start public folder and calendar synchronization.** Establishing public folder synchronization ensures that changes made in one organization get replicated to the other, so users can share the same public folder space. Migration Manager also allows you to set calendar synchronization independently of mailbox migration. That way you can ensure that calendar information is also available for any user in any organization.
- 3. Establish free/busy synchronization (optional).** Free/busy synchronization enables users to schedule common activities. Migration Manager can synchronize free/busy information independently from other data and thus make sure that the information gets updated as closely to real-time as possible.
- 4. Synchronize mailbox data.** When mailbox synchronization is launched, Migration Manager starts transferring the source mailboxes' content to the target mailboxes and synchronizing mailbox permissions.
- 5. Migrate the directory.** Migrate accounts from the source to the target domain in migration sessions. This is really a re-migration of accounts but effectively it moves the disabled account created during the directory synchronization to the permanent destination OU in the target directory. You can also delegate rights to migrate accounts to other administrators in your environment.
- 6. Process distributed resources.** Update distributed resources (such as user workstations, file and print servers, and application servers) using Resource Updating Manager, adding permissions to the resources for the target users. When user workstations are updated, user profiles should also be updated, so the migrated users will get the same profile as the corresponding source users when they log on to the target domain for the first time. You can delegate rights to perform resource update to other administrators in your environment.
- 7. Move user workstations to the target domain.** This step is actually the user switch, because when you move a workstation to a target domain using Resource Updating Manager, the last logged-in domain on users' workstations is changed from the source to the target domain and thus users start to log in to the target domain. Then also move file and print servers and other application servers that have been processed by Resource Updating Manager to the target domain.
- 8. Switch to the new Exchange mailboxes.** When a mailbox is switched, Migration Manager sets redirection to the opposite direction: all new mail sent to the old (source) mailbox is automatically forwarded to the new mailbox in the target organization. Migration Manager also marks the mailbox in a way that initiates the Outlook profile update at the user's next logon.

9. **Update Outlook profiles.** Migration Manager is shipped with the Profile Updating Utility (EMWProf), which handles Outlook profile update. After update the profile points to the target Exchange server and user mailbox. The majority of the properties stored in a profile also get updated.

i **NOTE:** Migration of users can be performed in batches. If you choose this approach, repeat steps 5 through 9 for each group of users you migrate.

! **CAUTION:** Note that user and mailbox switch (steps 7–9) should be completed in a closed time period to make sure that users will not access their source mailboxes when they start to log on to the target domain (with target accounts). This is to ensure that target users do not create new items in their old (source) mailboxes or public folders. If they do create such items, the owner of all those newly-created messages will be set to 'unknown' and the creators will not be able to modify the items later.

10. **Change the mail exchanger or alias records.** Switch incoming SMTP traffic to the target Exchange bridgehead server when about 50 percent of the users have had their mailboxes switched to the target in order to optimize routing.
11. **Stop and uninstall the synchronization agents.** When all mailboxes are migrated and switched, you can stop the synchronization. The following agents should be stopped and uninstalled:
 - Directory Synchronization Agents
 - Mailbox Synchronization Agents
 - Synchronization Agents
 - Folder Synchronization Agents
 - Free/Busy Synchronization Agents (if they were used)
12. **Process BackOffice servers and move them to the target domain.** Update Microsoft BackOffice servers, such as Exchange, SQL, and SMS Server, using the corresponding processing wizards. You can delegate rights to perform BackOffice servers update to other administrators in your environment. You can move the server to the target domain before or after it is updated. Note that you can start to update BackOffice servers right after all accounts have been migrated to the target domain (step 5).
13. **Disable the source accounts.** If the source accounts were not disabled before user and mailbox switch in order to set the Associated External Account to the source accounts (see the note above), disable them in this step. We recommend that you wait some time after disabling the source accounts before proceeding with the next step to make sure that all users are using their target accounts.
14. **Enable SID filtering.** After SID filtering is enabled, wait some time to ensure that all target users can access the resources they used before the migration.
15. **Clean up SIDHistory attributes from the target accounts.** After SIDHistory is cleaned up, wait some time to ensure that all target users can access the resources they used before the migration.
16. **Clean up legacy account permissions from resources.** Note that cleanup is hard to undo. It is recommended that you clean up permissions only when you are sure that all users are using their target accounts for all applications and have no problems accessing resources.
17. **Clean up the additional SMTP addresses and custom attributes.** Clean up the additional SMTP addresses set for redirection purposes and the custom attributes of the target objects used during Exchange migration.
18. **Decommission the migrated environments.**

Exchange First, Then Active Directory Migration

The migration according to this scenario consists of two separate phases:

1. Exchange data migration, as described in the [Exchange Migration](#) topic.
2. Active Directory migration, as described in the [Active Directory Migration](#) topic. One additional step is required: process all target Exchange servers with the Exchange Processing Wizard.

The main idea of this scenario is to complete the Exchange migration first; that is, implement the Exchange Resource Forest, and then complete Active Directory migration to that forest, merging source user with their stub mailboxes in the Exchange Resource Forest. These migration phases are independent and can be separated in time.

After the first phase is complete, source users access their target mailboxes. The Associated External Account of each target mailbox-enabled user (stub) contains a SID of the corresponding source user. Thus, the ACLs of all items created in the target mailboxes by source users contain source users' SIDs.

When you complete the second phase, the Active Directory migration, the only additional step you need to take is to process all target Exchange servers with the Exchange Processing Wizard. This is needed to translate permissions on items in mailboxes and public folders granted to source users through the Associated External Account from source to target accounts.

Migration Steps

Both of the preceding migration scenarios are a combination of two other migration scenarios: Active Directory Migration and Exchange Migration.

All descriptions and recommendations for the migration steps in the Active Directory Migration and Exchange Migration scenarios are fully applicable to the combined Active Directory and Exchange migration scenario. For descriptions of the migration steps, refer to the [Active Directory Migration](#) and [Exchange Migration](#) topics.

Additional best practices and recommendations for Active Directory migration, resource update, and Exchange migration are described in the following topics: [Considerations for Active Directory Migration and Resource Update](#) and [Considerations for Exchange Migration](#).

Rollback

If only one-way directory, public folder, calendar, and free/busy synchronization is established, Migration Manager does not make changes to the source environment.

During the account migration phase, Migration Manager just reads data from source and applies it to target. That means that until you disable source accounts, if migration issues arise, users can log in back to source. We recommend that you keep source accounts for a period of time after the migration is finished.

Migration Manager is designed so you can roll back any changes made to the environment at any step of the migration process, as follows:

- [Roll Back Active Directory Changes](#)
- [Roll Back Changes Made to Exchange Organization](#)

Roll Back Active Directory Changes

Return to the Source Accounts

The easiest way to roll back is to start using source accounts again. You can start using source accounts at any stage of the migration process as long as you leave the source accounts' permissions while processing the resources.

Session Undo

You can also undo the corresponding session in Migration Manager to roll back the account migration. This will delete the accounts created by the session on target. Merged accounts will be returned to the states they had before the migration.

! **CAUTION:** When you undo a session from Migration Manager, all migrated accounts are removed from the target domain. But if you have already re-assigned permissions on the resources to the target accounts (that is, performed steps 3 and 5 described in the [Active Directory Migration](#) topic), each resource's ACL after the session undo will contain unresolved SIDs of the deleted objects. Therefore, you should always return the target environment to its original state by performing permissions revert before doing a migration session rollback.

Permissions Revert

Permissions revert is done by the same wizards that were used to re-assign permissions to the target accounts.

i **NOTE:** Migration Manager does not keep the information about permissions originally set to each object before the migration and resource updating. Revert procedures simply substitute target SIDs with the corresponding source SIDs.

Restore from Backup

If changes were made to the source environment that cannot be undone, use standard procedures to restore from backup. That's why it is recommended that source and target environments be backed up with standard backup procedures. Ensure that you have the latest valid backup of your servers in both source and target.

We also encourage using Recovery Manager for Active Directory during all migration and post-migration activities to back up Active Directory. This allows online granular restoration of objects down to the attribute level without a domain controller reboot if they are accidentally deleted or corrupted.

Roll Back Changes Made to Exchange Organization

Return to the Source Mailboxes

Users work with their source mailboxes until the mailbox is switched and the client profile is updated on the user workstation. After the mailbox is switched and client profile is updated, a user starts using the target mailbox. However, if for some reason you want the user to work with the source mailbox again, at any point of migration you can do the following:

- Undo the mailbox switch from Migration Manager for Exchange
- Undo the client profile update using the Client Profile Updating utility (EMWProf)

For more information about the mailbox switch and the client profile undo, refer to the *Quest Migration Manager for Exchange—User Guide*.

Restore from Backup

If changes were made to the source environment that cannot be undone, use standard procedures to restore from backup.

We also encourage using Recovery Manager for Exchange during all migration and post-migration activities to back up source Exchange servers data. This allows online granular restoration of objects down to the message item level if they are accidentally deleted or corrupted.

Considerations for Active Directory Migration and Resource Update

[Using Import Lists to Modify Attributes or Merge and Rename Accounts](#)
[Linked Attributes and Group Migration](#)
[Subsequent Migrations](#)
[Skipping Attributes](#)
[Pre-installing the Resource Updating Agents](#)
[Migrating Services](#)
[Interoperability with Quest Collaboration Services](#)

Using Import Lists to Modify Attributes or Merge and Rename Accounts

An *import* list is a text file listing accounts to be migrated with Migration Manager and the attribute values for each account. If you are going to perform a one-time migration and plan to migrate all users and groups at once, an import list is not necessary. However, import lists should be used when migrating large environments (over 2,500 users). It makes the migration process more manageable, more accurate, and much more efficient. In addition, using import lists can help you easily track migration activities by site, region, location, or any other criteria set up during the planning phase.

You can use import lists to modify target object attributes. To easily create an import list, when creating a new migration session in Migration Manager, use the Import/Export functionality. Migration Manager allows you export information about the accounts from the selected OU and save it in a plain text, tab-delimited format. You also can select what attributes to export. Then you can open the file in Excel, modify the attribute values, and import the list back. The modified attribute values will be applied to the target objects during migration.

Import lists can also be used for account renaming. For example, you may want the names in the target to meet new corporate standards. To assign the target account different name on the fly (this can be a name, or sAMAccountName), you can simply add a new column to the exported list right after the first column and populate it with the new object names (name or sAMAccountName). The new names are applied to the newly-created target objects during migration.

i **NOTE:** Microsoft operates with a virtual attribute 'name,' which is the common name (attribute: cn) for some object classes, such as user, group, and computer. For other object classes, the name may be different. For example, for the organizational unit object class, the virtual attribute 'name' takes the value of the attribute ou.

The following is an example of a file for renaming:

```
sAMAccountName sAMAccountName description
JSmith JohnSmith Sales_Department
TBroun TomBroun IT_Department
```

Another approach to account renaming is to determine which names in the source domain are non-standard and change them before migration. To list account names, estimate the number of non-standard names and

decide which ones must be changed, you can use the *General User Information* report in Quest Reporter or export the information from the source domain into a tab-delimited file using third-party tools.

You can also use this method to merge source and target accounts that have different names. To do this, create the import list containing at least two columns, source SAMAccountName and target SAMAccountName, and specify the actual names for the source and target accounts in the columns. Using such an import list in Migration Manager later will cause source and target accounts be matched by their SAMAccountNames and merged.

Linked Attributes and Group Migration

When you migrate a group from source to target, the target group membership should also be updated. This means that the user and group objects in the target domain corresponding to the user and group objects in the source domain that are members of the source group should also be added to the target group membership list. Since group membership is a list of so-called *linked attributes*, this process is called link resolving.

Examples of linked attributes are:

- member/memberOf
- manager/directReports
- managedBy/managedObjects

Migration Manager resolves links when you migrate or synchronize objects. However, if the objects that should be added to the target group membership (attribute: member) or set as a manager (attribute: manager) for other objects do not exist in the target domain (have not been migrated yet), group membership and manager information will not be updated for such objects. Such objects will be placed into the unresolved links queue. To resolve the links in the future when all group members are migrated to the target, you can either re-migrate and merge groups or run full resynchronization.

To make sure that all the links (group membership, Manager, Managed by, etc.) will be successfully resolved from the beginning without the need to run full resynchronization later, we recommend you migrate group members first and then migrate a group or migrate a group and its members together in the same session.

i **NOTE:** DSA cannot synchronize groups with more than 5000 members. Even if you are using Microsoft Windows Server 2003 Forest mode or higher, as recommended by Microsoft, use primary group membership to ensure correct synchronization of large groups.

Subsequent Migrations

If you are running a number of subsequent migrations with ongoing directory synchronization (for example, migrating domain A to domain B and then migrating domain B to domain C), you should use different service attributes for the domain pairs A–B and B–C.

The attributes that can be used as the service attributes must meet the following criteria:

- Are not used for other purposes
- Have “Unicode String” syntax
- Are replicated to Global Catalog (only for matching attribute and only for projects with multiple target domains)
- Are indexed in Active Directory

Normally, you can use any extension attributes that are not used by Migration Manager for Exchange and not used to store any other information in the enterprise.

Skipping Attributes

You can skip certain attributes from directory synchronization. However, although Migration Manager allows you skip the majority of object attributes from its interface, you should never skip the following attributes from synchronization:

The Directory Synchronization Agent service attributes (`extensionAttribute15`, `extensionAttribute14`, `adminDescription`, `AdminDisplayName`, by default)

- `objectClass`
- `objectGUID`
- `userAccountControl`
- `objectCategory`
- `objectSID`
- `msExchMasterAccountSid`
- `msExchMailboxSecurityDescriptor`

Otherwise, you may experience problems during directory synchronization.

Pre-installing the Resource Updating Agents

The Resource Updating Agents are used for resource update. The agent is normally installed to the computer when you start processing from Resource Updating Manager.

If desired, Resource Updating Agents can be pre-installed using Group Policy or SMS.

This also allows you make sure that you have enough rights over the workstation or server to perform update and Windows Firewall is turned off on these computers.

Migrating Services

To minimize service downtime when you move your production servers to the target domain, you need to make sure that the service processing operation happens immediately after the computer move operation for those servers.

Suppose your source domain has a server where a critical service runs under **`sourcedomain\serviceaccount`**. If you process services along with other resources prior to moving the server, this will cause the service to stop working. The service will switch to using **`targetdomain\serviceaccount`** and will not be available until you have moved the server.

The recommended steps are as follows:

1. Process non-critical resources.
2. Move the computer.
3. Process the critical services using a task that immediately follows the move task.

Interoperability with Quest Collaboration Services

If you have previously established Global Address List and free/busy synchronization between the source and the target Exchange organizations using Quest Collaboration Services, the recommended migration procedure is the following:

1. Stop and disable Collaboration Services in both the source and the target domains.
2. Configure directory synchronization to merge source objects with the stub objects created in the target directory by Collaboration Services. Use the matching rule by e-mail address to match source objects with their stubs in the target domain.

The Directory Synchronization Agent enables the disabled stub account and creates a mailbox for it (if configured).

Also note that you need to migrate legacyExchangeDN attribute values of all stub objects created by Collaboration Services in the source directory to the original objects in the target directory. Otherwise, migrated users will get a non-delivery report (NDR) if they reply to an older email for which one of the recipients was Collaboration Services stub object. There are special custom add-ins for migrating these values.

Considerations for Exchange Migration

- Synchronization Agents
 - Directory Synchronization
 - Public Folder Synchronization
 - Configure Public Folder Jobs
 - What Accounts to Use for Public Folder Synchronization
 - How to Configure Public Folder Synchronization Collections
 - Whether Hierarchy Reorganization is Required
 - How Public Folder Deletions Are Handled
 - When Public Folder Resynchronization is Required
 - Mailbox Synchronization
 - Mail Synchronization Agents and Jobs
 - Using Collections to Map Exchange Migration Wizard
 - Offline Folder (OST) Files and Remote Users Collections
 - Restructuring Exchange Servers
 - Calendar Synchronization
 - One-Way vs. Two-Way Synchronization
 - Mailbox Switch and Profile Update
 - Interoperability with Other Quest Products

Synchronization Agents

Agent File Locations and Disk Space

Agents should usually be installed on the drive with the maximum free disk space available. Each server can have only one location for all Exchange Migration Wizard agents. Changing this location requires agent reinstallation and job reconfiguration, so plan carefully before installing agents.

Each server allows you to specify the amount of free disk space required for the agents to run. You can specify a percentage of the total disk space or a particular size. Use whatever setting you are comfortable with, but it is best to stay on the safe side and make sure at least 500 MB of free space exists on all servers.

The mail and public folder agents use temporary PST files to transfer data and a database for a local cache. Substantial temporary disk space may be required, depending on the settings used.

Each agent has a log file associated with it. Use each agent's option to specify the size at which the log will be archived and compressed.

Pre-installing the Agents on Exchange Servers

In large distributed networks with sites connected by slow links, Exchange agent deployment from Migration Manager for Exchange console may take up all available bandwidth. This happens because the Migration Manager for Exchange has to transfer the agent setup packages to the remote server before it can start installation. The shared components setup is about 15 MB and can take considerable time and bandwidth to be transferred.

Migration Manager for Exchange includes the files necessary for creating the installation package to install the Exchange agents on remote Exchange servers. The package can be distributed to remote sites on any

removable media and local site administrators can perform installation prior to starting the migration process. The package setup creates all necessary folders and shares on servers and copies files to required locations. Before installing the agents, make sure that all servers in your environment meet the software requirements listed in the *Quest Migration Manager—System Requirements and Access Rights* document.

The following files are located in the Migration Manager installation folder on the console computer:

- `makepack.cmd` – Creates the installation package on the console computer
- `clnsetup.cmd` – Setup file for cluster servers
- `srvsetup.cmd` – Setup file for non-cluster servers
- `ccrsetup.cmd` – Setup file for Exchange 2007 CCR cluster servers

These batch files allow you to create an installation package that can be then distributed to remote locations without consuming network bandwidth.

These files do not eliminate the need to run the agent installation procedure from the Migration Manager console. They simply allow the copying of the setup files to the required location in advance. After the files are copied, you need to finish the agent installation from the console.

Using Agent Hosts for Synchronization Agents

By default, the best migration performance is achieved when Migration Manager for Exchange agents are installed on the same Exchange servers as the mailboxes and public folders they process. However, in some cases, the only way to perform data migration successfully is to use an agent host for the Exchange server.

An *agent host* is a server that performs a role of substituted Exchange server in migration process. An agent host can be either an Exchange server or any other server. Single instances of all Migration Manager for Exchange agents are installed on that server (including the Public Folder Source Agent, Public Folder Target Agent, Mail Source Agent, Mail Target Agent, Mail Transmission Agent, Calendar Synchronization Agent, and Free/Busy Synchronization Agent).

Using agent hosts can be recommended in the following cases:

- If no additional software (i.e., Migration Manager for Exchange agents) should be installed on the production Exchange servers due to specific reasons
- If Migration Manager for Exchange agents installed on a particular Exchange server do not work correctly due to conflicts with the third-part software installed on that server

Before you start using agent hosts, consider the following:

- The following computers cannot be used as agent hosts:
 - Microsoft Cluster Servers (MSCS)
 - Microsoft Cluster Server nodes
 - Exchange Virtual Server
 - Exchange 2000/2003/2007 SCC clusters
 - Exchange 2007 CCR clusters
- If agent host is used, network traffic increases and migration slows down if compared to usual migration scenario. This happens due to additional mail transfer (from source Exchange server to agent host and from agent host to target Exchange server).
- To increase the performance and speed up migration process, it is recommended to use the agent host located in the same network segment as the corresponding Exchange server.

- One agent host can work with several Exchange servers, in particular, be an agent host for both a source and a target Exchange server at a time. However, in this case, the agent will process all assigned jobs sequentially, one after another, which may lead to performance degradation.
- Using agent host can be recommended for Exchange 2007 CCR to prevent problems with unexpected failover, as well as to solve problems with replication of data from the active to the passive node of the CCR cluster server (if they exist).

Directory Synchronization

When you configure the directory synchronization job, you specify the target administrative group and the information store in which the Directory Synchronization Agent will create mailboxes. These mailboxes can later be moved by the Mail Target Agent (MTA) or Calendar Synchronization Agent (CSA) to another information store or administrative group, as specified in the mailbox or calendar synchronization job.

However, it is recommended that you plan the initial mailbox creation during directory synchronization in a way that minimizes unnecessary mailbox moves across administrative groups by the mail and calendar synchronization agents.

Note that mailbox move from one administrative group to another is only possible if the Exchange organization is running in native mode.

Public Folder Synchronization

The following are the important issues to consider when synchronizing public folders:

- Whether to use one-way or two-way public folder synchronization
- How many public folder synchronization jobs to configure
- What accounts and administrative mailboxes to use for public folder synchronization
- How to configure public folder synchronization collections
- Whether hierarchy reorganization is required
- How public folder deletions are handled
- When public folder resynchronization is required

For more details, see the related topics.

However, note that public folder synchronization is optional and you may skip this procedure if desired.

i | **NOTE:** To learn how to migrate public folder data to Microsoft SharePoint, please refer to the Quest Public Folder Migrator for SharePoint documentation.

Configure Public Folder Jobs

A public folder job is configured between a source and target server pair to synchronize public folders. The source and target servers for a public folder synchronization job must have local replicas for all folders to be synchronized.

It is highly recommended that you set up and configure all the public folder synchronization jobs before actually starting the synchronization process. If any job requires reconfiguration after it has been started, replication may not run for period of time and some data might be lost because of incorrect reconfiguration.

Whether to Choose One-Way or Two-Way Synchronization

The main decision when setting up a job is whether to choose one-way or two-way synchronization:

- Two-way synchronization should be set for those folders in which data may change in either source or target environment. Two-way synchronization of public folder data and permissions will allow all users to have access to the up-to-date information they need.
- One-way synchronization should be used if a folder is updated only in the source, or for a one-time migration of the folders that change very rarely and do not require ongoing synchronization. One-way synchronization can be configured from source to target only.

! CAUTION: If after one-way public folder synchronization is started, you decide to switch it to two-way, the whole public folder content that was transferred to the target will be transferred back to the source. This may cause performance degradation in the agents. Therefore, it is recommended that you do not change the public folder synchronization direction after synchronization starts.

How Many Public Folder Synchronization Jobs are Needed

If you want to migrate your public folders “as is,” you can simply set up one public folder job with All Public Folders as the root of synchronization. However, it may not be possible to have replicas of all folders on a single source Exchange server. In this case, you will have to create several synchronization jobs with different servers as their sources.

Having multiple jobs running in parallel improves the synchronization speed but complicates troubleshooting. It is recommended that you minimize the number of jobs by replicating all source public folders to selected servers. For example, you can use one source Exchange server in each physical location to have replicas of all public folders used in this location. These servers will serve as bridgeheads for public folder synchronization with target Exchange servers.

What Accounts to Use for Public Folder Synchronization

Public folder agents are installed on each source and target server involved in a public folder synchronization job. A public folder agent can synchronize only folders for which it has a local replica. For this reason, public folder jobs should span a set of servers that have at least one local replica of each source public folder in the source organizations.

When planning public folder synchronization, you should consider that public folders created by the agents on the target Exchange servers will be owned by the account that the Public Folder Target Agent uses. It is recommended that this account be mailbox enabled.

Also, a mailbox must exist on each server participating in public folder synchronization. The public folder agents use a local mailbox to access the public folder data. You should create dedicated mailboxes on the servers that host public folders. These mailboxes should not be included in the mailbox synchronization.

How to Configure Public Folder Synchronization Collections

Public folder collections are used to prioritize public folder synchronization. Normally, each job has one collection configured and the public folders in these collections will be synchronized until the end of the migration.

Multiple collections may be associated with each job. Each collection has a priority number starting from 1, which is the highest priority. The collections are processed by the agent in sessions according to their priority number. The agents start processing with the collection that has the highest priority. The collection priority can be changed.

A collection is not processed until all public folders in higher priority collections are in sync. Therefore, multiple collections should be used with care. If data will be continuously synchronized and a high-priority collection has its source folders change frequently, collections with lower priority may not be processed.

Whether Hierarchy Reorganization is Required

Only the default MAPI top level hierarchy in each Exchange organization can participate in public folder synchronization.

You may want to restructure your target Exchange public folder hierarchy so that it is different from the existing source Exchange hierarchy, particularly if you are consolidating Exchange organizations. Migration Manager for Exchange fully supports migration and restructuring even if your source Exchange public folder hierarchies have different structures and in the new environment they will be standardized.

You can specify that a source public folder (and, optionally, subfolders) synchronize with any existing public folder in the target Exchange organization, or you can specify a new folder in the target Exchange public folder hierarchy with which to synchronize data. Migration Manager for Exchange will create new public folders in the target Exchange organization as required.

You can also exclude any folder from a synchronization job (and include it in a job elsewhere if desired).

Merging Folders

There may be cases where multiple source public folders (from the same or different Exchange organizations) will be merged into one target public folder. There are limitations in doing this and the result is that the source public folders will have different data.

When an item is added, moved, or deleted in a source public folder, the corresponding change will take place in the target public folder. However, none of the other source public folders will be notified of this change because when the change is made to the item in the target public folder, it is marked in such a way that it is considered to be up to date. Similarly, if an item is changed in the target public folder, this change will be propagated to only one of the source public folders. For this reason, when multiple folders are merged to one target folder, the synchronization direction should be from source to target only.

Even if only one job is configured to participate in two-way synchronization, that source public folder will not contain an exact copy of the target public folder. The agent running on the target is not designed to synchronize back to multiple folders although it can receive updates from multiple source public folders.

Synchronization with Multiple Organizations

Target public folders may be synchronized with folders from multiple source Exchange organizations with the limitations stated above. If there are no source public folders to be merged into the same target folder, then two-way synchronization can be set up between all organizations without issues. The setup of the synchronization jobs must be done carefully so that each public folder in the target unambiguously corresponds to one source public folder.

How Public Folder Deletions Are Handled

Migration Manager for Exchange synchronizes deletions of public folder items and folders themselves. Messages deleted on one side are deleted from the corresponding public folder on the other side.

In order to avoid data loss, public folder synchronization agents do not delete public folders. Instead, they move deleted folders to a special public folder called Aelita EMW Recycle Bin. When a public folder is deleted in one of the environments, the public folder synchronization agents move the corresponding folder in the other environment to the Recycle Bin. Thus, administrators can verify that no important information has been deleted and restore any data deleted by mistake.

! CAUTION: Although the folders are moved to the Recycle Bin and an administrator can move them back, folders that have been deleted once cannot be migrated again. This is because public folder synchronization agents preserve the folders' SOURCE_KEY property. On the side where the folders are actually deleted, Exchange tombstones the deleted folders and does not allow creation of folders with the same SOURCE_KEY property again.

Therefore, you should never delete production public folders on either source or on target. If you want to experiment with a folder and delete it later, create a test public folder specifically for that purpose. Never experiment with live folders. If you want to do a pilot public folder migration, create a test public folder and use it for your experiments. Avoid doing pilots on the production folders.

Public folder synchronization agents can create the Recycle Bin folder, but it is recommended that an administrator create and configure this folder manually prior to starting public folder synchronization. The reasons for this are as follows:

- **Replication** – The folder should be replicated among all the Exchange servers within the organization that participate in public folder synchronization. If this is not configured, the agents will create a separate Recycle Bin folder on each of these servers, which will result in several folders with the same name in the public folder tree.
- **Access to the deleted folders** – If the agents create the Recycle Bin folder, it will have the default settings for **Default** and **Anonymous** access. This may not be desired, as usually the default permissions are set to the Author role.
- **Folder properties** – If the agents create the Aelita EMW Recycle Bin folder, the other folder settings will have their default values as well. For example, you might want to change folder limits or replication schedule.

Aelita EMW Recycle Bin must be a top-level folder right under the All Public Folders node. The agents recognize this folder by name, so if you create it manually, be sure to type its name correctly.

When Public Folder Resynchronization is Required

In certain situations you may want to resynchronize public folders. However, because Migration Manager for Exchange uses Exchange replication APIs for public folder synchronization, resynchronization is not an easy task. Resynchronization that is not performed properly may interfere with the Exchange public folder replication. Always contact your migration consultant or technical support specialist before choosing to do a full resynchronization.

Mailbox Synchronization

The following are the important issues to consider when synchronizing mailboxes:

- [How many jobs have to be configured and how many agents should be installed](#)
- [How to use collections to map Exchange Migration Wizard configuration to your migration plans](#)
- [What users typically work with offline folder \(OST\) files](#)
- [How to restructure Exchange servers during migration](#)

Mail Synchronization Agents and Jobs

A mailbox synchronization job is configured between each source/target server pair. The jobs are performed by mailbox synchronization agents. The Mail Source Agents (MSA) and Transmission Agents work on the source Exchange servers, while the Mail Target Agent (MTA) works on the target Exchange servers.

i **NOTE:** Mailboxes included in Remote Users Collections are synchronized by the Mail Source Agents only. The whole mailbox content, including the Calendar folder, is transferred to the target Exchange 2003 server at once by the MSA. The Mail Target Agents and Transmission Agents are not used.

Synchronization agents work with the Private Information Store locally, so they will be installed on each server involved in mailbox synchronization (on all source and target servers). Each agent can participate in multiple synchronization jobs. This allows for messaging system restructuring, such as splitting or merging servers, as described in the Restructuring Exchange Servers topic.

Using Collections to Map Exchange Migration Wizard

A *collection* is a group of mailboxes that should be migrated within the same time period and to the same target mailbox store. Mailboxes can be placed in a collection directly or through a recipient container or distribution list (DL). Where possible, use containers to populate collections. This will ensure that the newly-created mailboxes are included in synchronization automatically.

However, if the mailboxes you plan to synchronize in a Remote Users Collection are in the same container as normal users' mailboxes, it is better to use distribution lists to populate both the Remote Users Collections and

normal collections. This will ensure that mailboxes of the remote users will not be synchronized in normal collections and vice versa.

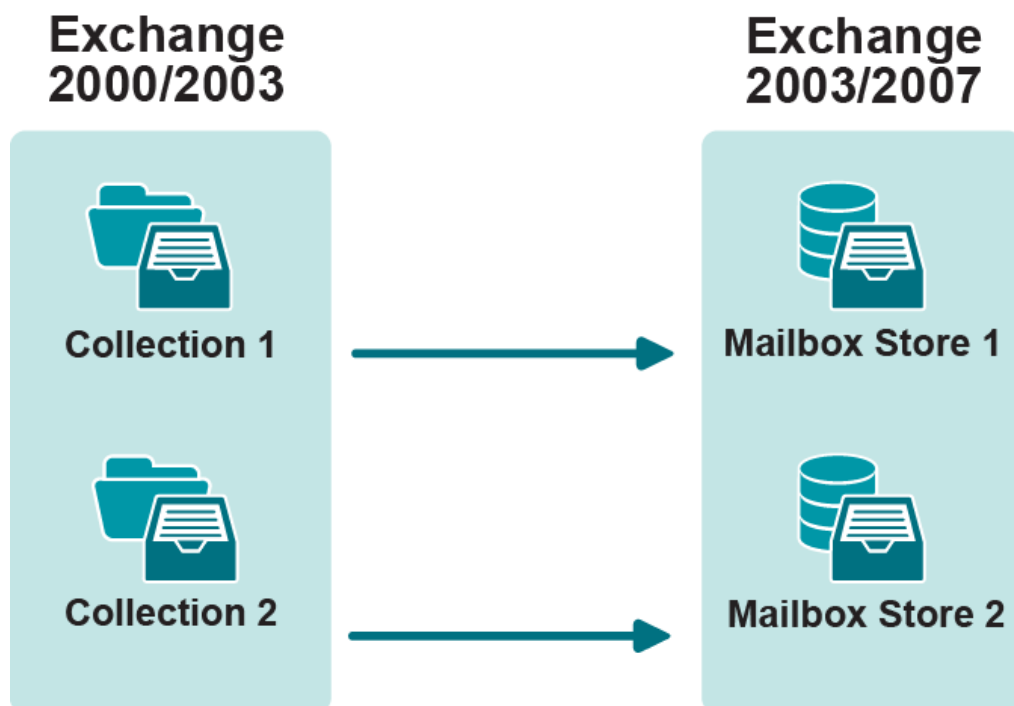
Multiple collections may be associated with each job. Each collection has a priority number starting from 1, which is the highest priority. The collections are processed by the agent in sessions according to their priority number. The agents start processing with the collection that has the highest priority. The collection priority can be changed.

Mailbox synchronization jobs and collections for the entire migration should be set up at the start of the project. Each collection has settings that control when the mailboxes that belong to it are synchronized.

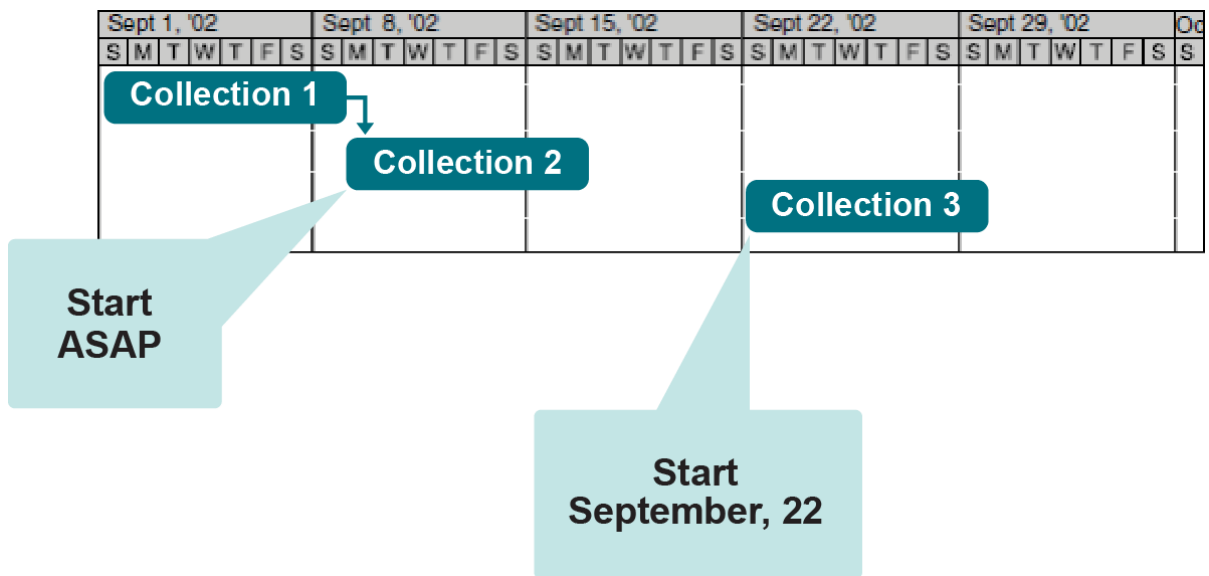
If you plan to use one collection for a source server, create and populate mailbox collections to organize the migration of mailbox content as appropriate for your needs—for example, by department, by projected migration date, and so on.

The **All mailboxes** collection cannot be used as a Remote Users Collection. Usually jobs are split into collections in order to achieve the following goals:

- **Reorganize stores** – Each collection puts its mailboxes into a separate mailbox store, as illustrated in the following diagram:



- **Project management and scheduling** – When a server has a large number of mailboxes to be migrated, collections help keep track of the migration process:



- **Migration of remote users' mailboxes** – When a server has a number of remote users who work with the local offline folders (OST), the mailboxes of these should be grouped in Remote Users Collections, which are processed separately from other mailbox collections.

Offline Folder (OST) Files and Remote Users Collections

Populating Remote Users Collections

You should add to Remote Users Collections those remote and laptop users who:

- Strongly rely upon their offline folder (OST) files and need them preserved
- Occasionally connect to their Exchange mailboxes using slow links, such as a dial-up connection, to synchronize mail
- Would be severely affected if the OST file must be rebuilt

You should not add to the Remote Users Collection any user who:

- Normally works online and either does not have an offline folder (OST) file or does not need it to be preserved
- Has a fast and reliable connection to the Exchange server
- Should not be switched to the target Exchange mailbox automatically

Considerations

You should also consider the following before adding users to a Remote Users Collection:

- Each mailbox in a Remote Users Collection and its corresponding target mailbox are unavailable for the user while the Mail Source Agent is processing them.
- Each mailbox in a Remote Users Collection is automatically switched as soon as the Mail Source Agent completes copying its data to the target mailbox. These mailboxes cannot be switched manually.
- It is recommended that the Remote Users Collection be synchronized after the directory synchronization has been completed and before the migration of other mailboxes is started.
- The mailboxes of a Remote Users Collection are resynchronized only when you manually undo the switch for them. A Remote Users Collection cannot be configured to be resynchronized using Project Manager.

Migrating Remote Users Collection

To determine the optimal number of mailboxes to be added to each Remote Users Collection, you can evaluate how many mailboxes can be migrated per night (in, say, 8 hours). It is recommended that you build a lab similar to your production environment and test how many mailboxes the Mail Source Agent can process during the time interval you specify.

Our testing has shown that the Mail Source Agent can process 26 mailboxes, 150 MB each, of a Remote Users Collection per hour. However, please note that agent performance tightly depends on the server and network capacity, mailbox size, and many other factors, and therefore these figures are for approximation only.

To add the mailboxes of remote and laptop users to a Remote Users Collection and start synchronizing the collection:

1. Determine which users should be added to the Remote Users Collection.
2. Group the mailboxes of these users in one or several Remote Users Collections, depending on the optimal collection size you determined through testing.
3. Make sure the target environment is ready for the new users to log on.
4. Schedule the Mail Source Agent to process the Remote Users Collections for a time when the users normally do not use their mailboxes. Note that this schedule will affect the Remote Users Collections only; normal collections are processed all the time when the agent is scheduled to run.
5. If the MSA is configured to create mailboxes in a mailbox store different from the store specified for the Directory Synchronization Agent, the MSA will re-home the mailboxes by modifying some attributes in Active Directory and process the re-homed mailboxes in its next session. In this case, the changes made to the attributes in Active Directory must be replicated before the MSA starts the next session. If the replication latency between Active Directory domains in the forest is large, it is recommended that you leave the default sleep interval between sessions for the Mail Source Agent (30 minutes). This is to ensure that at least one Active Directory replication cycle between the MSA sessions is completed. Otherwise, the MSA may transfer the mailbox content to the store that was initially set by the Directory Synchronization Agent rather than to the store specified in the Remote Users Collection.

i **NOTE:** During synchronization of a remote user's mailbox, the Mail Source Agent removes the corresponding target mailbox (if any) and recreates it. When the Mail Source Agent logs on to the mailbox in order to process it, the mailbox becomes unavailable for the user for the time it is being processed by the agent.

6. Add EMWProf to the logon script of the users whose mailboxes are included in the Remote Users Collections.

When the time interval specified for the agent to process the Remote Users Collections is over, check whether the mailboxes of the Remote Users Collections have been synchronized and switched by the Mail Source

Agent. When the users log on, EMWProf will start updating their profiles, and after update the target mailboxes will become the active users' mailboxes.

Synchronization of the Calendar folders of the remote users' mailboxes will be performed by the Calendar Synchronization Agent during the entire period of migration, but you should start calendar synchronization for such mailboxes only after the switch. This is to prevent the CSA from creating mailboxes on the target before they are processed by the Mail Source Agent. The calendar agent will transfer calendar data from the source mailbox to the target mailbox.

Restructuring Exchange Servers

In inter-org Exchange migration you may want to split mailboxes across multiple servers or storage groups or stores (databases), or consolidate mailboxes onto fewer Exchange servers.

Splitting Servers or Stores

Each job specifies a source-target server pair. This facilitates splitting mailboxes across multiple servers. Mail agents process jobs one at a time. To migrate from one server to multiple servers, create separate jobs.

To migrate mailboxes to multiple stores, create separate collections and specify the appropriate target store in the collection's properties.

Consolidating Servers

It is simple to consolidate multiple source Exchange servers into one target Exchange server. Each source-target server pair will have a job. Data received from multiple jobs is handled in serial by the agent running on the target server, while source mailboxes are processed in parallel by the multiple agents running on the source servers.

Calendar Synchronization

The following are the important issues to consider when synchronizing calendars and free/busy data:

- [How many jobs have to be configured and where the agents should be installed](#)
- [Whether you need one-way or two-way calendar synchronization](#)

Calendar Synchronization Agents and Jobs

One Calendar Synchronization Agent (CSA) performs calendar synchronization for the mailboxes. There is one calendar synchronization job per source-target server pair. While setting up a calendar synchronization job, you are prompted to install the Calendar Synchronization Agent on either the source or target Exchange server. The Calendar Synchronization Agent synchronizes the personal Calendar folders of the mailboxes hosted on the source Exchange server with the corresponding Calendar folders of the mailboxes hosted on the target Exchange server.

A calendar synchronization job must be populated with the members it will synchronize. A mailbox can be part of a job directly or through recipient container or distribution list membership. Whenever possible, use containers to populate jobs. This will ensure that a newly created mailbox will be included in a calendar synchronization job even if it is not explicitly added.

However, if the mailboxes you plan to synchronize in a Remote Users Collection are in the same container as normal users' mailboxes, it is better to use distribution lists to populate both the Remote Users Collections and normal collections. This will ensure that mailboxes of the remote users will not be synchronized in normal collections and vice versa.

One-Way vs. Two-Way Synchronization

Migration Manager for Exchange allows you to specify whether a calendar synchronization job should be one-way or two-way.

One-way calendar synchronization means that the calendar entries for each source-target mailbox pair are synchronized one way only, from the currently active mailbox to the other mailbox. An active mailbox is the source user's Exchange mailbox until the user is switched. Once switched, the active mailbox is the target user's Exchange mailbox. Choose one-way synchronization if mailboxes that collaborate are migrated in groups such that for each mailbox, all mailboxes it accesses will be switched at the same time.

Two-way synchronization should usually be selected, since it is difficult to migrate mailboxes in closed sets, especially when resource mailboxes are used (for example, if you represent conference rooms as mailboxes).

! CAUTION: The CSA synchronizes permissions granted for the Calendar folder depending on the mailbox status.

If the mailbox is already switched, permissions are synchronized from target to source.

If the mailbox is not switched yet, permissions are synchronized from source to target.

Free/Busy Synchronization

Free/busy synchronization is performed by the Free/Busy Synchronization Agent.

While setting up a free/busy synchronization job, you are prompted to install the Free/Busy Synchronization Agent on either the source or target Exchange server.

Free/busy data resides in a public folder. In this public folder there is one message for each user that contains that user's data. When the Calendar Synchronization Agent synchronizes appointments, the free/busy data is automatically updated.

It is often not necessary to synchronize free/busy separately from calendar data. Instead, you can increase the calendar synchronization frequency by reducing the time period during which the Calendar Synchronization Agent sleeps.

However, if you migrate users in Remote Users Collections and do not include them in calendar synchronization, which is recommended, free/busy synchronization will allow the already migrated users to see free/busy information of the users that are not migrated yet.

You need to synchronize free/busy separately only if either of the following applies:

- The time taken for a Calendar Synchronization Agent session to complete (scan and synchronize all changes) does not meet your requirement.
- You migrate mailboxes in Remote Users Collections.

When set to two-way synchronization, the CSA processes all mailboxes of a server for about 15–45 minutes, assuming the number of mailboxes is close to the Microsoft recommendations (approximately 1000 mailboxes

per server). If you need to provide users with more up-to-date free/busy data, use one of the agents per source site to synchronize free/busy data.

While processing the free/busy synchronization job, the agent synchronizes all the free/busy messages for all the mailboxes it finds in the organizations selected for free/busy synchronization. Thus, you normally need only one Free/Busy Synchronization Agent and on free/busy synchronization job for the pair of source and target Exchange organizations.

It is important to remember that Migration Manager for Exchange does not enable multiple source Exchange organizations to have a view of the entire directory (and associated mailboxes and calendar data). Calendar synchronization will enable all target Exchange users to view and modify all calendar data. The users connecting to mailboxes in the source Exchange organizations will continue to have access to calendar data for users in their organization, whether those users are switched or not. However, users connecting to the source Exchange organization will not be able to see calendar detail data for users that were migrated from other source Exchange organizations.

Mailbox Switch and Profile Update

The following are the important items to consider regarding mailbox switching and profile updating:

- [When and how to switch mailboxes](#)
- [When and how to update client profiles](#)

Mailbox Switch

Switching a mailbox means that all new mail comes to the target Exchange mailbox. A mailbox should be switched once all data from the source mailbox is migrated to the target mailbox. You can switch mailboxes automatically or manually, as follows:

- Mailboxes can be switched automatically (which is recommended) by the Mail Source Agent when the mailbox is in sync. If you want to control when mailboxes are switched automatically, specify the **Start switching mailboxes** date and time in the collection properties. The agent will start switching mailboxes that are in sync (or almost in sync) only after the date and time you specify.

i **NOTE:** When Mail Source Agent automatically switches a mailbox, it checks whether the PRV file with 'ready to switch' flag is already processed by the Mail Target Agent - this is used to minimize the risk of not all data yet present in the target mailbox at the instant that the mailbox is switched. This default behavior, however, requires additional logon to target mailbox, which may decrease the performance if the logon is performed slowly.

To switch this additional logon to target mailbox OFF, set LastPSTAppliedSwitch=0 for Mail Source Agent.

As the MTA inserts the mail into the target mailbox, and this is run independently of the MSA, in this case you may want to run the MTA continuously to minimize the time that the data does not exist in the target after the switch.

In addition, run the agents during off hours so that the MTA will run before the users run EMWProf and connect to their target Exchange mailboxes.

- You can switch particular mailboxes manually from the Migration Manager for Exchange interface. Note that no further synchronization will occur after you switch a mailbox manually.

CAUTION: Manual switch is not available for the mailboxes of Remote Users Collections. These mailboxes can only be switched automatically by the Mail Source Agent.

Client Profile Update

Profile update usually happens during user logon. If users remain logged on during the mailbox switching, use Migration Manager for Exchange to send them an e-mail message (switch notification message) instructing them to log off and back on so that their profiles will be updated. You can customize the message sent to the users for each source Exchange server. For more details on switch notification capabilities, refer to the *Quest Migration Manager for Exchange—User Guide*.

The EMWProf utility can also be configured to send a notification message when it finishes processing the Outlook profile on a client workstation. You can specify the administrator's mailbox to which the messages will be sent. The notification message that the EMWProf sends includes the update status (succeeded or failed), the computer name on which the EMWProf was executed, and the EMWProf log file. You can filter these messages by their status and take action for those with the failed status as soon as you receive them. For more details on EMWProf's notification capabilities, refer to the Client Profiles Updating Utility documentation.

Interoperability with Other Quest Products

Exchange Migration Wizard

If Exchange Migration Wizard is used to migrate from Exchange 5.5 in parallel with Migration Manager for Exchange to migrate from Exchange to the same target Exchange environment, the following restrictions apply:

- Different computers must be used to install Exchange Migration Wizard and Migration Manager for Exchange Console.
- Different target Exchange servers must be used to migrate data from the source environments by Exchange Migration Wizard and Migration Manager for Exchange. That is, the Exchange agents (such as mail, public folder, calendar, and free/busy synchronization agents) from different products must not be installed and run on the same target Exchange servers.
- Different custom attribute sets must be used as the service attributes by the two products.
- Different SMTP address templates must be used for redirection purposes by the products.

Additional planning should also be done to coordinate Exchange 5.5 and Exchange 2000/2003/2007 migration activities, especially if separate migration teams are responsible for each migration.

Preferred Settings for the Directory Synchronization Agent

It is required that you always specify the preferred domain controller and the preferred Global Catalog server for each Directory Synchronization Agent, and that these be located in the same site as the agent. In addition, we recommend that you set the Directory Synchronization Agent to always use the same domain controller and Global Catalog in the domain.

This ensures, first of all, that the agents will not work with domain controllers and Global Catalog servers located in remote sites across WAN links.

Second, the Directory Synchronization Agent uses Microsoft's Dirsync control to query for directory changes, and Dirsync's behavior is not consistent when different domain controllers are used to query for changed objects. This issue with Dirsync may have an undesired effect on the Directory Synchronization Agent performing delta sync: if a domain controller from which the agent used to retrieve information about the objects modified since the last session becomes unavailable and no preferred domain controller is set, the agent will be automatically redirected to another domain controller. The new domain controller may return many more objects (or even all objects from the directory) as modified, causing the agent to perform unnecessary jobs.

Microsoft describes the issue as follows:

For incremental searches, the best practice is to bind to the same domain controller (DC) used in the previous search, that is, the DC that generated the cookie. If the same DC is unavailable, either wait until it is, or bind to a new DC and perform a full synchronization. Store the DNS name of the DC in the secondary storage with the cookie.

You can pass a cookie generated by one DC to a different DC hosting a replica of the same directory partition. There is no chance that a client will miss changes by using a cookie from one DC on another DC. However, it is possible that the search results from the new DC may include reported changes by the old DC; in some cases, the new DC may return all objects and attributes, as with a full synchronization. (<http://msdn2.microsoft.com/en-us/library/ms677626.aspx>)

Refer to the *Quest Migration Manager for Active Directory—User Guide* for procedures on how to configure the preferred settings for the agent.

Directory Synchronization Agent Placement

It is not recommended to install the Directory Synchronization Agent on the domain controller, Exchange Server, or console machine to avoid possible server and agent performance degradation. The best practice is to install the Directory Synchronization Agent on a dedicated server that has good connectivity to both the source and the target domains being synchronized. For the list of supported operating systems, see the *System Requirements and Access Rights document*.

Indexing Service Attributes

When synchronizing objects, the Directory Synchronization Agent populates the objects' attributes specified as the service attributes in Migration Manager with matching and auxiliary information. It is required to index the service attributes in Active Directory to improve the Directory Synchronization Agent performance.

For more information about the service attributes, refer to the *Quest Migration Manager for Active Directory—User Guide*.

For more information about Directory Synchronization Agent performance, refer to the *Migration Manager for Active Directory - Improving Directory Synchronization Performance* document available at <https://support.quest.com/>.

Full Directory Resynchronization

Normally, no resynchronizations are required after the initial synchronization is complete. The ongoing synchronization process brings only changes made since the last agent's session to the opposite directory. Changes are identified by changed Update Sequence Numbers (USNs).

However, in certain cases you might consider full resynchronization for a job. These are described in further detail below.

- Group membership for the groups may not be updated correctly if you have migrated groups before user accounts are migrated to the target domain (see the [Linked Attributes and Group Migration](#) topic.) Running full resynchronization after all objects are migrated will help resolve all the links.
- If you have modified a directory synchronization job configuration after the job has already been started, full re-synchronization is required. You will be prompted by Migration Manager to restart the Directory Synchronization Agent processing the job. The initial synchronization (full source and target directories enumeration and synchronization) will take place.

Note that full directory resynchronization is a time-consuming operation. During the operation, the source and target directories are fully re-enumerated, and then the objects within the specified synchronization scope are matched and synchronized. The process may take up to 10 hours for 100,000 objects, depending on your hardware. You should plan for full re-synchronization and avoid it until it is actually required.

Conclusion

Active Directory and Exchange migrations are monumental tasks. This is especially true for large distributed and complex environments. It is essential that a solid discovery and analysis be completed on the entire enterprise prior to migration. All testing should be performed in an environment that mirrors the production environment as exactly as possible.

To make migration easier, first perform Active Directory migration, ensure that Active Directory is stabilized, and then go forward with the Exchange migration.

Although no two projects are exactly the same, this document outlines some of the key factors for ensuring a successful Active Directory and Exchange migration.

Environment Preparation Checklist

The table below lists the preparation steps for the source and target environment that must be accomplished before you start the migration process.

You can find the detailed explanations and procedures for each step listed below in the *Quest Migration Manager Installation Guide*. The *Quest Migration Manager System Requirements and Access Rights* document will help you to set the required permissions for Active Directory and Exchange migration with Migration Manager.

Task	Comments	Refer to
Check the software requirements	<p>Before installing Migration Manager, make sure that the following servers meet the software requirements:</p> <ul style="list-style-type: none"> • Source and target servers • Console • ADAM or AD LDS server • Configuration database for Exchange migration • Directory Synchronization Agent server • Migration Manager Statistics server <p>Any computer that does not meet these requirements should be upgraded before installation of Migration Manager.</p>	<i>System Requirements and Access Rights</i>
Prepare the source and target environments for Active Directory migration	<ul style="list-style-type: none"> • Establish trusts • Disable SID filtering • Check host name resolution • Verify that the required ports are open on servers, routers, and firewalls. 	<i>Installation Guide</i>
Prepare the source and target environment for Exchange migration	<ul style="list-style-type: none"> • Implement Exchange backup strategy • Create "Aelita EMW Recycle Bin" public folder. • Create administrative mailboxes 	<i>Installation Guide</i>
Connect the source and target Exchange Organizations using the SMTP connector	<ul style="list-style-type: none"> • Set up the SMTP connectors • Force Exchange to accept the redirected mail as inbound mail • Modify name resolution parameters on the bridgehead servers in Exchange • Configure DNS for mail forwarding 	<i>Installation Guide</i>

Task	Comments	Refer to
	<ul style="list-style-type: none"> • Test the SMTP connectors. 	
Set the required permissions	<ul style="list-style-type: none"> • Set the required permissions for Active Directory migration • Set the required permissions for Exchange migration 	<i>System Requirements and Access Rights</i>
Install Migration Manager	No comments	<i>Installation Guide</i>

Exchange Migration without Trusts

Planning migration without trust established

If you are planning an Exchange migration using Migration Manager for Exchange without having trusts established, you will need to perform the following actions:

- In the Migration Manager Console click **Tools | Options**, and on the **License** property page in the **Net use connections** area verify the specified account has access rights to the License Server.
- Users will have to specify the target service account using Client Profile Updating Utility when they are switched to the target server. Because there are no trusts, their source accounts will not have permissions for the target mailboxes.

To configure the Client Profile Updating Utility for Microsoft Outlook profile processing if trust relationships are not established between the source and the target domains, refer to the [Client Profile Updating Utility Administration Guide](#).

When trusts are required

The above specified procedure is valid for the most every cases, but there are some exclusions in specific environment:

- In case the computer on which Migration Manager is installed is a member of the domain in which the Windows 2000 or Windows Server 2003-based Exchange cluster servers reside and such cluster servers are in both the source and target domains, you will need trusts to be established between the source and target domains.

Active Directory Migration without Trusts

This section describes the common issues you may face if you perform an inter-forest migration using Migration Manager for Active Directory without having trusts established. In such an environment, the source and target directories are isolated one from another and there might be no transition period when both source and target accounts can be used. But the migration can still be performed without critical restrictions if the accounts are switched correctly.

Take the following steps to perform the migration:

1. Migrate the directory.
2. Process the distributed resources.
3. Process Microsoft Exchange and SMS/SCCM servers.

These steps can be performed during working hours.

The following steps should be performed one right after another or during non-business hours:

1. Process the Microsoft SQL servers.
2. Update the services and scheduled tasks.
3. Move the computers to the target domain.
4. Restart the computers.

The last two steps make up the user switch.

i **NOTE:** Microsoft SQL servers should be processed during non-business hours because the SQL Processing Wizard replace the source accounts with the target accounts specified in permissions when processing such SQL servers.

Each of these steps is described in more detail below.

Prerequisites

For a successful migration, the computer on which Migration Manager is installed must belong to the target domain.

Migrate the Directory

Migrate accounts from the source to the target domain in migration sessions.

The only two restrictions on this step are as follows:

- SIDHistory does not provide the target accounts with access to resources due to the absence of trust between the source and target domains. It is recommended to clear the Add SIDHistory check box on the Set Security Settings step when configuring a migration session to avoid unresolved SIDs in the target objects' SIDHistory attribute.
- Copying local/global/universal group membership during migration does not provide the members of the target local groups with access to the resources granted to these groups. It is not recommended to select the Add source members to the corresponding target groups option in the Set Security Settings step when configuring a migration session to avoid unresolved SIDs in the opposite domain groups.

Update Distributed Resources and Microsoft Exchange and SMS/SCM Servers

In Resource Updating Manager, add the computers to be processed and specify the administrative accounts for the processed computers (see the Resource Updating Manager online help for details).

When reassigning permissions and group membership to target users and groups with Resource Updating Manager or one of the processing wizards, select the Leave source accounts permissions check box to preserve resource access for the source accounts. This is required since the users continue using their source accounts until they are completely switched to the target accounts.

Switch to the Target Domain

The following operations should be performed one right after another or during non-business hours:

- Process Microsoft SQL servers using SQL Processing Wizard.
- Process services and scheduled tasks using Resource Updating Manager. Since there is no transition period, the update of services and scheduled tasks and the switching of user accounts should be performed one right after another. Otherwise, services and scheduled tasks will continue running under the source accounts until the users are switched.
- Move the workstations and servers to the target domain.
- Restart the computers to make the updated services and scheduled tasks still running under the source accounts run under the target accounts. You can restart the computers after moving, as described in the Resource Updating Manager help.

After that, users begin to use the target accounts.

After resources are processed and users can access their resources using target accounts, you can clean up the source accounts' permissions and disable source accounts.

Quest provides software solutions for the rapidly-changing world of enterprise IT. We help simplify the challenges caused by data explosion, cloud expansion, hybrid datacenters, security threats, and regulatory requirements. We are a global provider to 130,000 companies across 100 countries, including 95% of the Fortune 500 and 90% of the Global 1000. Since 1987, we have built a portfolio of solutions that now includes database management, data protection, identity and access management, Microsoft platform management, and unified endpoint management. With Quest, organizations spend less time on IT administration and more time on business innovation. For more information, visit www.quest.com.

Technical support resources

Technical support is available to Quest customers with a valid maintenance contract and customers who have trial versions. You can access the Quest Support Portal at <https://support.quest.com>.

The Support Portal provides self-help tools you can use to solve problems quickly and independently, 24 hours a day, 365 days a year. The Support Portal enables you to:

- Submit and manage a Service Request
- View Knowledge Base articles
- Sign up for product notifications
- Download software and technical documentation
- View how-to-videos
- Engage in community discussions
- Chat with support engineers online
- View services to assist you with your product