Quest® Migration Manager for Active Directory 8.15

User Guide

Migration Manager for Active Directory User Guide
Updated - October 2023
Version - 8.15

Contents

Introduction to Migration Manager for Active Directory

We assume that you are familiar with Migration Manager concepts described in the *Migration Manager Installation Guide*, and that you have already installed the product, following the instructions provided in that guide. It is also recommended that you read the release notes for the current version of Migration Manager. The release notes contain information about specific product behavior, limitations, known issues, and workarounds that may be useful for planning and performing your migration.

About Migration Manager for Active Directory

Migration Manager for Active Directory is an efficient, flexible, and comprehensive solution for restructuring your Active Directory. The restructuring is performed by migrating objects between forests or domains. Migration involves moving Active Directory objects (users, groups, and resources) from a source domain to a target domain.

Depending on the environment and the goals of the migration, the migration scenario may vary. The *Migration Manager Tips and Tricks* document contains common migration scenarios and considerations for choosing the one that best fits your requirements.

While the migration is underway (and some large-scale migrations can last for years), the source and target domains must coexist and stay in synch. Migration Manager eases the administrative burden during this period by providing synchronization capabilities, such as synchronization of account properties, group membership, and passwords: administrators simply make necessary changes in one environment and those changes are automatically replicated to the other environment.

Migration Manager also allows delegation of migration tasks to other people. Delegated administrators are able to migrate and process only the resources you specify.

This guide describes most of the migration process, including pre-migration activities, the migration itself, directory synchronization, and delegation. The final step, resource update, is described in the dedicated *Resource Processing* document.

Initial Configuration of Migration Project

The **Open Project Wizard** is the central place for configuring the migration project and all its components: the ADAM/AD LDS database, the SQL/MSDE database (for Exchange migrations only), and Statistics Portal.

The wizard can be used to start a new project, connect to an existing project, or change any settings of the current project.

> **! CAUTION: In most cases you need just one Migration Manager project for your whole migration, no matter how many domains you have. Normally, you would use a separate project only for lab testing before you start migration in the production environment.**

We assume that you became familiar with the Open Project Wizard while carrying out the procedures described in the Opening a Migration Project topic of the Migration Manager Installation Guide. Refer to that document for a full description of the wizard steps.

To briefly review, the Open Project Wizard steps are as follows:

1. On the **Configure ADAM/ADLDS Project** step, you specify the server where ADAM or AD LDS is installed, specify the port number used by ADAM (the default is 389; it might be different if ADAM or AD LDS is installed on a domain controller or this is not the first ADAM or AD LDS instance on the server), and select a project to connect to (you can choose to create a new one).

2. Next, on the **Set Auxiliary Account** step, you supply the user name and password that program components will use to access the ADAM or AD LDS database.

3. The **Configure SQL/MSDE Database** step is displayed if you installed Migration Manager for Exchange. On that step, you configure the database that will be used to store information related to your Exchange migration.

4. On the **Configure Statistics Portal** step, you can optionally specify the portal's URL and connection settings.

5. On the final step, review the settings and finish the wizard. For detailed description, refer to the *Opening a Migration Project* topic of the *Migration Manager Installation Guide*.

! **CAUTION: We recommend that you avoid managing the Properties setting of the same objects within the same migration project from several Migration Manager Console computers simultaneously.**
Note also that having several parallel Remote Desktop connections to the same console computer is not supported.

Managing the Migration Project

Migration of simple environments where accounts and resources are centralized in one major location can be accomplished from that location using a single Migration Manager console.

However, this method is not appropriate when you need to migrate a large distributed environment where accounts and resources reside in different geographic locations (sites), particularly if the sites are connected by slow links that limit the amount of data that can be transferred effectively. In this case, the migration must be performed in each site locally.

Migration Manager provides for effective project implementation and management when migrating large distributed environments.
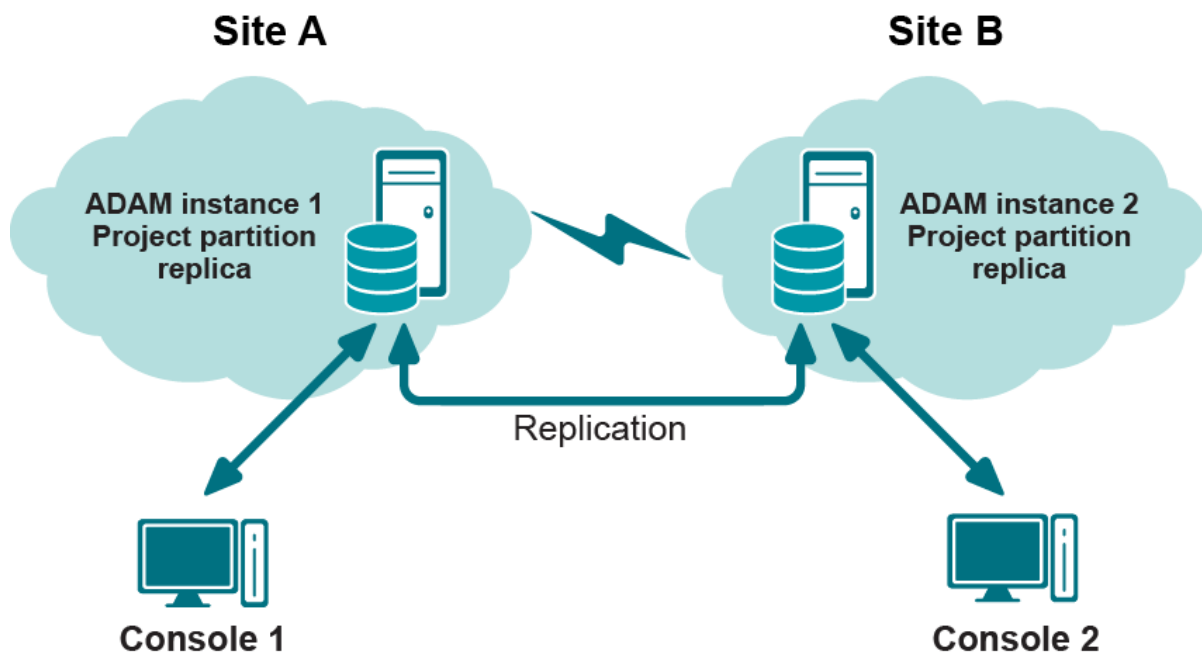
Figure 1: Figure 1. Managing a migration project in a distributed environment

If a number of administrators from multiple locations are going to perform migration, it is recommended to have local ADAM/AD LDS installations in each major location with a replica of the project partition. To avoid excessive traffic over slow links, administrators should work with the nearest ADAM or AD LDS server that contains a replica of the project.

In this case, you will need to delegate responsibility for migration tasks, as explained in the Delegating Migration Tasks topic.

Delegating Migration Tasks

Migration Manager allows administrators to delegate specific tasks to trusted persons responsible for particular stages of migration. For example, you can delegate the rights to manage the migration within a specified pair of domains to the administrators of those domains, or you can delegate the rights to process a specific server to the server administrator.

The trusted person will have the appropriate access to the objects that he or she is granted rights to only and will not have the opportunity to perform any actions with other objects.

The delegation is performed by assigning a role that defines a level of permissions to a person within a migration project. Only accounts that have Full Admin role over an object can delegate the rights over that object to another account.

Table 1: Table 1. Roles that can be assigned to a trusted person within a migration project.

Tree Node Object	Available Roles
Migration project	**Full Admin**—Can create and configure any objects (domain pairs, directory synchronization jobs, migration sessions, delegated migrations, resource processing tasks, etc.) and perform any task within the project.

Tree Node Object	Available Roles
	Reader—Can view objects within the project and their settings, except domain pairs and Exchange stores in the synchronization settings, but cannot perform any migration tasks or change the configuration.
Directory migration	**Domain Pair Creator**—Can create domain pairs. This person automatically gets the Full Admin role over the domain pairs he or she creates and therefore is able to change settings of the domain pairs later.
Domain pair	**Full Admin**—Can change any settings on a domain pair (such as credentials used to connect to each domain), and can also configure the directory synchronization between the domains, create and run migration sessions, and delegate the migration to other trusted persons. **Reader**—Cannot view or change settings of the domain pair, and cannot perform any migration tasks.
Delegated migration	**Migration Admin**—Can run migration sessions within the delegated migration job. This person can migrate objects from and to only those containers specified in the delegated migration job. He or she cannot see the credentials used to connect to the source and target domain or change the object processing custom add-in or the migration agent, but can see which custom add-in and agent are used.
Tasks	**Task Creator**—Can create resource processing tasks. This person is automatically assigned Full Admin role to the tasks he or she creates and therefore can change the task settings later and run the task.
Resource processing task	**Resource Admin**—Can run the resource processing task and change its settings. However, this person cannot change the re-permissioning options or perform undo or cleanup.

ℹ **NOTE:** Only a domain pair's Full Admins are granted access to the domain pair's sub-nodes (synchronization, migration, and other manually created delegated migration jobs) and migration sessions.

To delegate the rights to perform a migration task:

1. In Migration Manager, right-click the tree node object and select **Delegate** on the shortcut menu.

2. Specify the account to which you want to delegate the task.

3. Select the level of permissions for that account from the **Role** list.

4. Click **Add Account**. This will add the account to the **Delegated accounts** list.

5. Click **OK**.
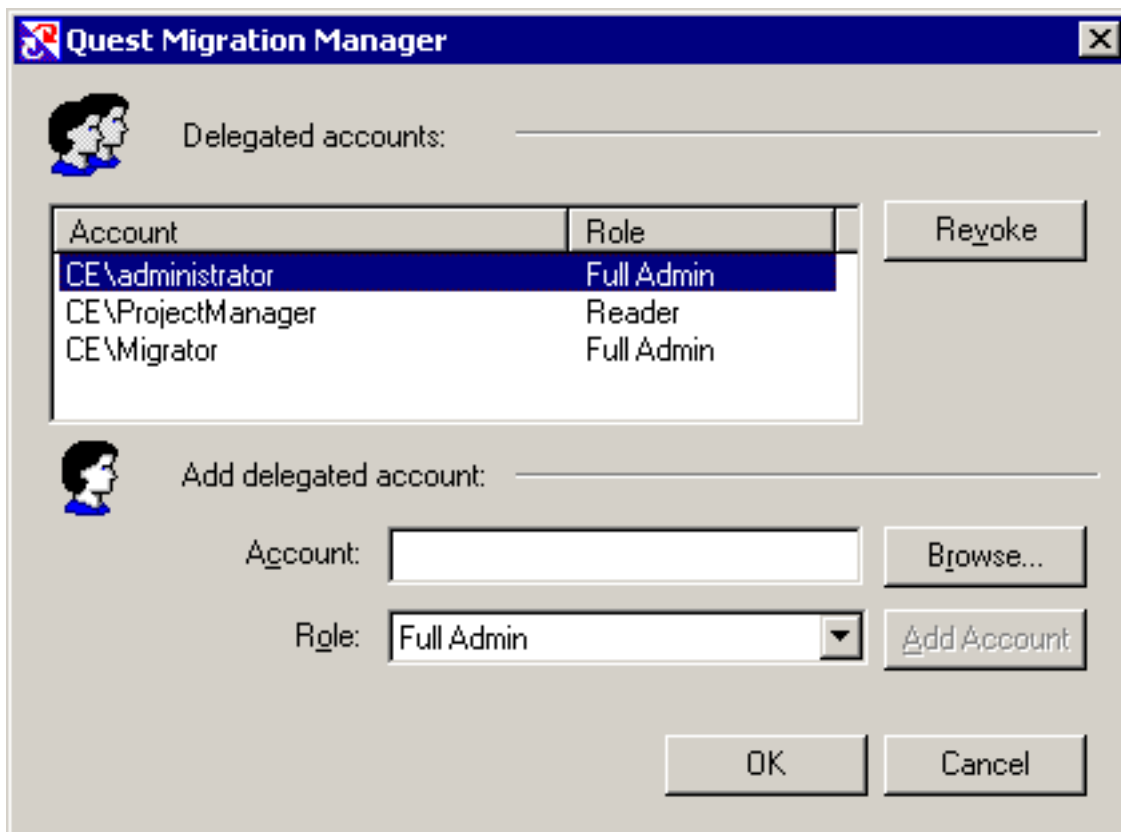
To revoke permissions, complete the following steps:

1. In Migration Manager, right-click the tree node object and select **Delegate** on the shortcut menu.

2. In the **Delegated accounts** list, select the account you want to revoke the rights from and click **Revoke**.

   i | **NOTE:** Refer to the Delegating Account Migration topic of this guide for more details about delegating account migration tasks.

3. Click **OK**.

Pre-migration Activities

Before you start directory migration, analyze the existing directory. This includes identifying required hardware and software upgrades, possible naming conflicts in the case of directory merges, and for an inter-forest migration, comparing and unifying the source and target forest schemas. You can use Quest Reporter to obtain detailed information about the existing Active Directory configuration, hardware and software inventory, etc.

You can optionally set up the target Active Directory forest, administrative accounts, and organizational units before migration. For more information about environment preparation, refer to the *Migration Manager Installation Guide*.

> **i** | **NOTE:** The RC4 encryption (Rivest Cipher 4 or RC4-HMAC) is an element of Microsoft Kerberos authentication that Quest migration products require to sync Active Directory passwords between Source and Target environments. Disabling the use of the RC4 protocol enabled makes password syncing between environments impossible.
>
> Beginning on November 8, 2022 Microsoft recommended an out of band (OOB) patch be employed to set AES as the default encryption type. The enabling and disabling use of the RC4 encryption protocol has potential impact beyond the function of password syncing of Quest migration tooling and should be considered carefully.

Exchange Migration Considerations

The following information is relevant to migration projects that involve both Active Directory and Exchange migration. Due to the tight integration of Exchange and Active Directory, some Exchange-specific decisions need to be made before even starting Active Directory migration.

There are four choices regarding the kind of users that will exist in the target domain after Active Directory migration:

- Users without mail options
- Mail-enabled users
- Mail-enabled users for Native Move
- Mailbox-enabled users

For Exchange migration, you need to have mail-enabled users, mail-enabled users for Native Move or mailbox-enabled users. For a smooth Exchange migration experience down the road, decide in advance which option will work for you. This choice is important for selecting the mailbox migration method.

Mail-enabled users are required if there are plans to include the users in the GAL without performing fully-fledged Exchange migration for them

Mail-enabled users for Native Move are required if you are plan to move mailboxes using the Native Move job.

Mailbox-enabled users are required if there are plans to:

- Move mailboxes using the Mailbox Synchronization job
- Move mailboxes using the Legacy Mailbox Synchronization job
- Set up calendar coexistence (with or without mailbox migration) using Calendar Synchronization jobs
- Set up calendar coexistence (with or without mailbox migration) using Legacy Calendar Synchronization jobs

To make an informed decision, discuss this with your Exchange migration operator. For details and guidelines, refer to the Mailbox Migration Process topic in the *Migration Manager for Exchange User Guide.*

The relevant Exchange options are specified during Active Directory Synchronization configuration. For details, see the *Step 5: Specify Exchange Options*section of the Configuring the Synchronization Job topic.

Domain Pairs

All migration activities are performed between source and target domain pairs. You should configure the pairs of source and target domains that will be involved in the migration and directory synchronization processes. The subsequent sections discuss how to create and configure the domain pairs:

- Creating a Domain Pair
- Configuring a Domain Pair

Creating a Domain Pair

This section explains how to create a new domain pair in the migration project.

i | **NOTE:** Before you create a domain pair, at least one Directory Synchronization Agent should be installed in your environment. Refer to the Directory Synchronization Agent topic for more details.

To create a domain pair, right-click the **Directory Migration** node and select **New Source and Target Domain Pair** from the shortcut menu. This will start the **New Domain Pair Wizard**, which will guide you through the process:

Step 1. Select Source Domain

Specify or browse to the domain controller of the domain that you want to make a source for the migration.

Specify the credentials for accessing the domain controller.
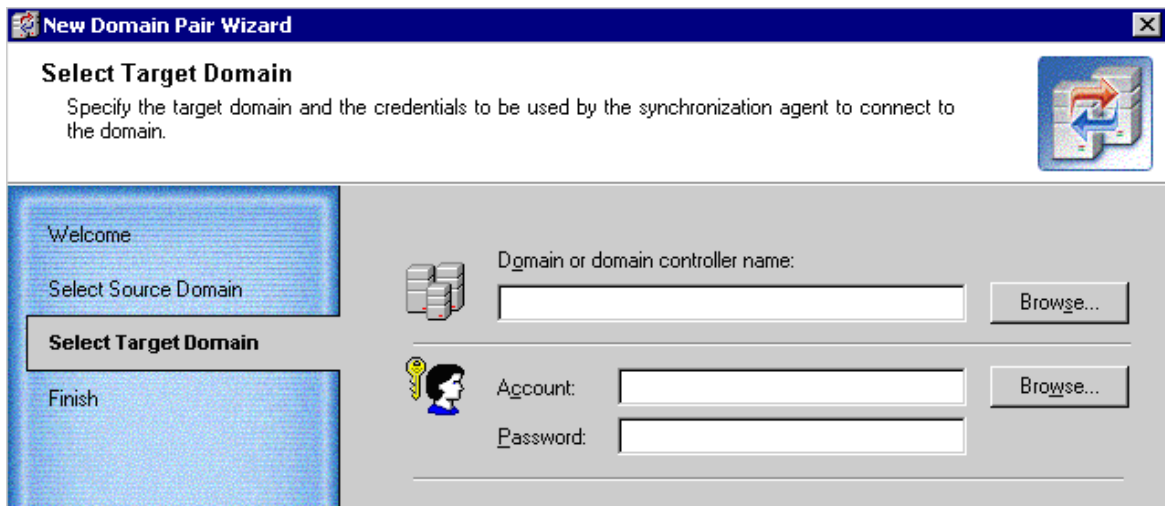
i | **NOTE:** Please note that the account you specify will be used by the Directory Synchronization Agent to access the source domain objects and to perform directory migration. Specified account should be a member of **Administrators** group in the source domain, or in accordance with the least privilege principle a Preinstalled Service feature should be configured and enabled.

Step 2. Select Target Domain

Specify or browse to the domain controller of the domain that you want to make a target of migration.

Specify the credentials for accessing the domain controller.

> **i** | **NOTE:** Please note that the account you specify on this page will be used by the Directory Synchronization Agent for accessing the target domain objects and to perform directory migration. Specified account should be a member of **Administrators** group in the target domain, or in accordance with the least privilege principle a Preinstalled Service feature should be configured and enabled.

Step 3. Complete the New Domain Pair Wizard

The wizard displays the names of the source and target domains and the accounts you specified for connecting to domains.

As soon as a domain pair is created, it will be displayed in the Migration Manager console management tree as a node having two sub-nodes, **Migration** and **Synchronization**.

Configuring a Domain Pair

After you have created a domain pair, you can specify configuration parameters for the domain pair. To do this, right-click the domain pair and select **Properties**. The parameters you can specify are described in the related topics.

> **i** | **NOTE:** Modifying these parameters requires full directory resynchronization. You must stop the synchronization job for the domain pair and then restart it using the **Start and Re-sync** option.

Skip Objects

This step allows you to specify the categories of objects that will be skipped during processing for all migration and synchronization tasks. You can select to skip the following types of objects:

- Active Directory default objects (objects present in Active Directory by default, such as built-in accounts and accounts like **Domain Admins** and **Domain Users**)

- Disabled accounts

- Expired accounts

If you select to skip any of these objects, you will not see them and therefore will not be able to select them when you browse the source or target domain of the domain pair.
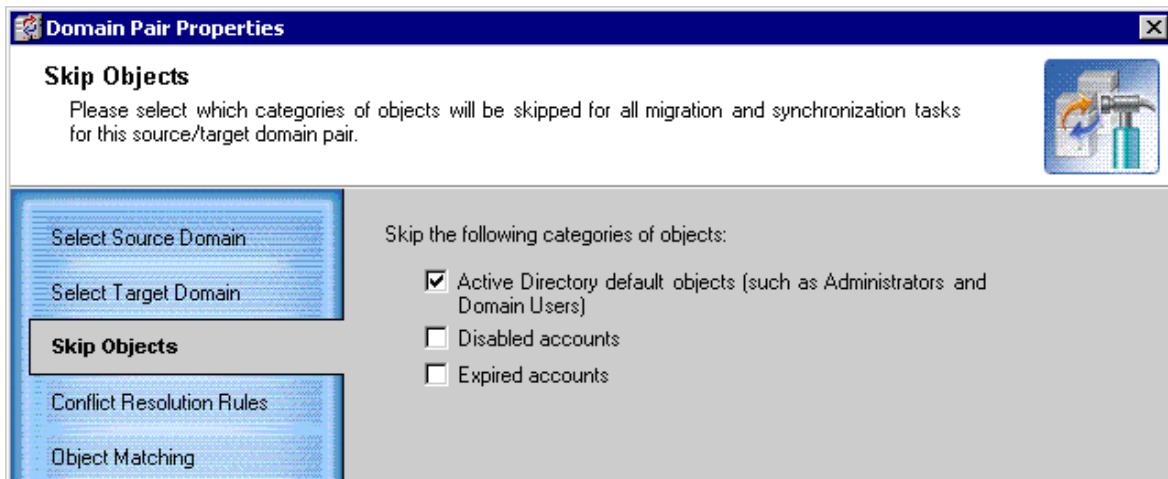
Specify Conflict Resolution Rules

You can specify the attributes that are to be unique within the given scope (forest, domain, or container) on source and target, and the action to be performed if these attributes are not unique (i.e., two or more objects exist with the same value for a specified attribute). This is done by setting conflict resolution rules.

i | **NOTE:** The conflict resolution rules you specify affect both migration and synchronization.
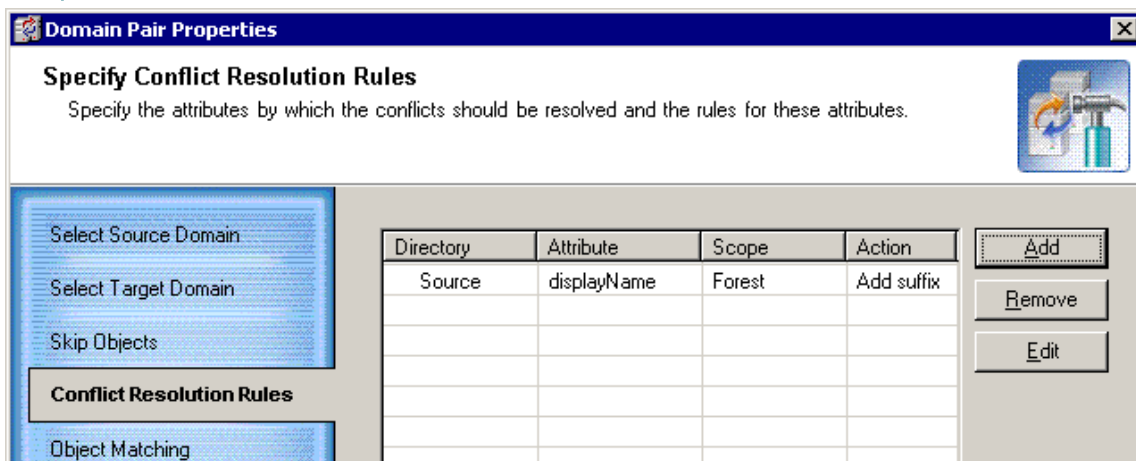
Click **Add** to set a new rule for automatic conflict resolution and make the appropriate settings in the **New Conflict Resolution Rule** dialog box, described below.

Click **Edit** to edit an existing conflict resolution rule.

Click **Remove** to remove the selected conflict resolution rule from the list.
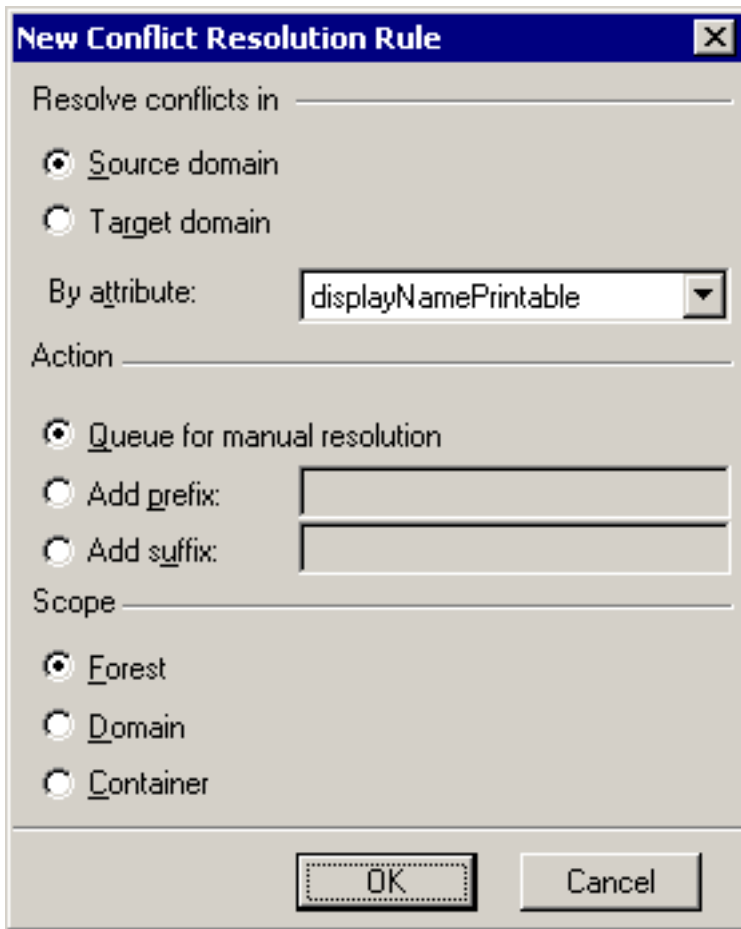
The rule consists of the following settings:

- **Source domain**—Specifies that the current conflict resolution rule will be applied on source.

- **Target domain**—Specifies that the current conflict resolution rule will be applied on target.

- **By attribute**—Select the attribute that you want to resolve the conflicts by.

- **Queue for manual resolution**—If this option is selected, conflicts in the selected attribute will not be resolved automatically but instead will be queued for later manual resolution. If directory synchronization is established between the domains in a domain pair, you will see objects that were queued for manual conflict resolution in the **Conflicts queue** of the directory synchronization job for the domain pair. Refer to the Directory Synchronization topic for more details. In the case of migration, you can see the conflicting objects by inspecting the migration log. Refer to the Viewing Migration Session Details topic for more details.

- **Add prefix**—If this option is selected, the specified prefix will be added to the attribute value if the attribute is not unique within the specified scope.

- **Add suffix**—If this option is selected, the specified suffix will be added to the attribute value if the attribute is not unique within the specified scope.

- **Forest**—Specifies that conflicts should be resolved within the whole forest.

- **Domain**—Specifies that conflicts should be resolved within the whole domain.

- **Container**—Specifies that conflicts should be resolved within each container.

Configure Object Matching

This step allows you to specify attributes for object matching during migration and synchronization.
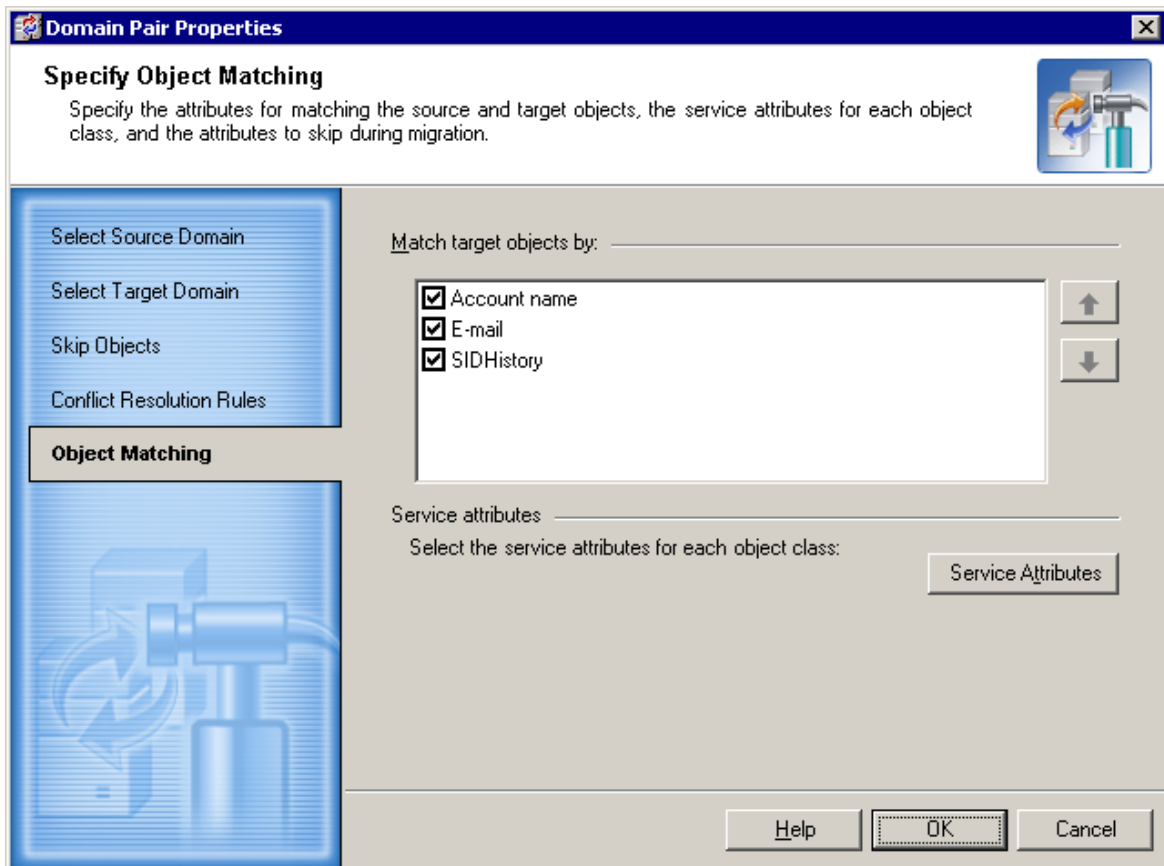
The Directory Synchronization Agent will match the source and target objects according to the attributes selected for object matching. If the agent cannot find a matching object in the target directory, a new object is created in this directory and its attributes are populated with the values of the corresponding source object.

i | **NOTE:** The object matching rules you specify affect both migration and synchronization.

The **Match target objects by** dialog box displays the fixed list of attributes that can be used to match source and target objects. These attributes are:

- **Account name**—If the account names of the source and target object are the same, the objects will be matched.

- **E-mail**—For mail-enabled objects, if a source and target object have the same primary SMTP address, the objects will be matched.

- **SIDHistory**—If the SIDHistory attribute of an object from one directory contains the security identifier (SID) of an object from another, the objects will be matched.

Select the attributes you want to use for matching the source and target objects. Matching will be performed in top-down priority. To change the attribute priority, use the up and down arrow buttons.

When two objects are matched, the information about matching is written to the service attribute specified to store the matching information. See the *Service Attributes* section for more details.

Service Attributes

The service attributes are used by the agent to store its internal information. Click **Service Attributes** on the **Match target objects by** dialog box to specify the attributes for each of the source and target object classes.

Object class—Specifies the object class for which you want to specify service attributes. In most cases you can use the default attributes.

Auxiliary—Specifies the attribute for storing information about conflict resolution and other service information. The attributes selected by default to store the auxiliary information for different object classes are listed below.

If Exchange is deployed in the forest, the following attributes are used:

Object Class	Auxiliary Attribute
contact	extensionAttribute14
container	adminDescription
group	extensionAttribute14
inetOrgPerson	extensionAttribute14
organizationalUnit	adminDescription
printQueue	adminDescription
user	extensionAttribute14
volume	adminDescription

If there is no Exchange in the forest, then by default the **adminDescription** attribute is used to store auxiliary information for all object classes.

Matching—When two objects are matched, the Global Unique Identifier (GUID) of the corresponding object is written to the service attribute selected to store information about matching. The attributes selected by default to store the matching information for different object classes are listed below.

If Exchange is deployed in the forest, the following attributes are used:

Object Class	Auxiliary Attribute
contact	extensionAttribute15

Object Class	Auxiliary Attribute
container	adminDisplayName
group	extensionAttribute15
inetOrgPerson	extensionAttribute15
organizationalUnit	adminDisplayName
printQueue	adminDisplayName
user	extensionAttribute15
volume	adminDisplayName

If there is no Exchange in the forest, then by default the **adminDisplayName** attribute is used to store matching information for all object classes.

> **!** **CAUTION: The service attributes must meet the following criteria:**
>
> - **The attributes configured as service attributes should not be used for other settings.**
> - **Only Unicode String attributes should be used as service attributes.**
> - **If you have multiple target domains in your environment make sure that the Matching attributes you select are replicated to Global Catalog.**
> - **Matching attributes should be indexed in Active Directory.**

Account Migration

Before you start your migration activities, be sure to install at least one Directory Synchronization Agent in your environment. You should also create in the migration project the pairs of source and target domains that will be involved in the migration and directory synchronization processes. Refer to the Directory Synchronization Agent and Creating a Domain Pair topics of this document for more details.

As soon as a domain pair is created, it will be displayed in the Migration Manager console management tree as a node having two sub-nodes, **Migration** and **Synchronization**.

The **Synchronization** node shows the template created for the directory synchronization job. The job can be further configured and started to perform object synchronization between the two domains. For more information about directory synchronization, refer to the Directory Synchronization topic.

Migration Session

Using the **Migration** node, you can create migration sessions. Each migration session is basically a group of accounts to be migrated from the source domain to the target domain. The **Migration Wizard** will help you migrate the selected objects within the source-target domain pair along with the objects' properties and security settings. Also, during migration you can change, in bulk, any of the attributes of the accounts being migrated by simply importing a tab-separated list with the new values. As soon as the selected objects have been migrated to the target domain, the session is considered to be complete.

The history and configuration of the session is stored in the database and can be viewed when necessary.

Considerations for Migration Sessions

When splitting up domain migration into sessions you should consider the way that linked attributes (such as group membership) get resolved:

- Linked attributes always get resolved in the scope of the session. For example, if you migrate a group and its members within the same session, the membership will get migrated.

- Linked attributes are also resolved for previously migrated objects. For example, if you first migrate users and then migrate a group, the group will be migrated with its membership.

- Backlinks (such as "member of") are not updated across sessions. For example, if you first migrate a group and then in other sessions migrate its members, the newly migrated accounts will not get added to the target group. If you have to migrate a group before its members, you can restore the membership by either re-migrating the group or doing full re-synchronization.
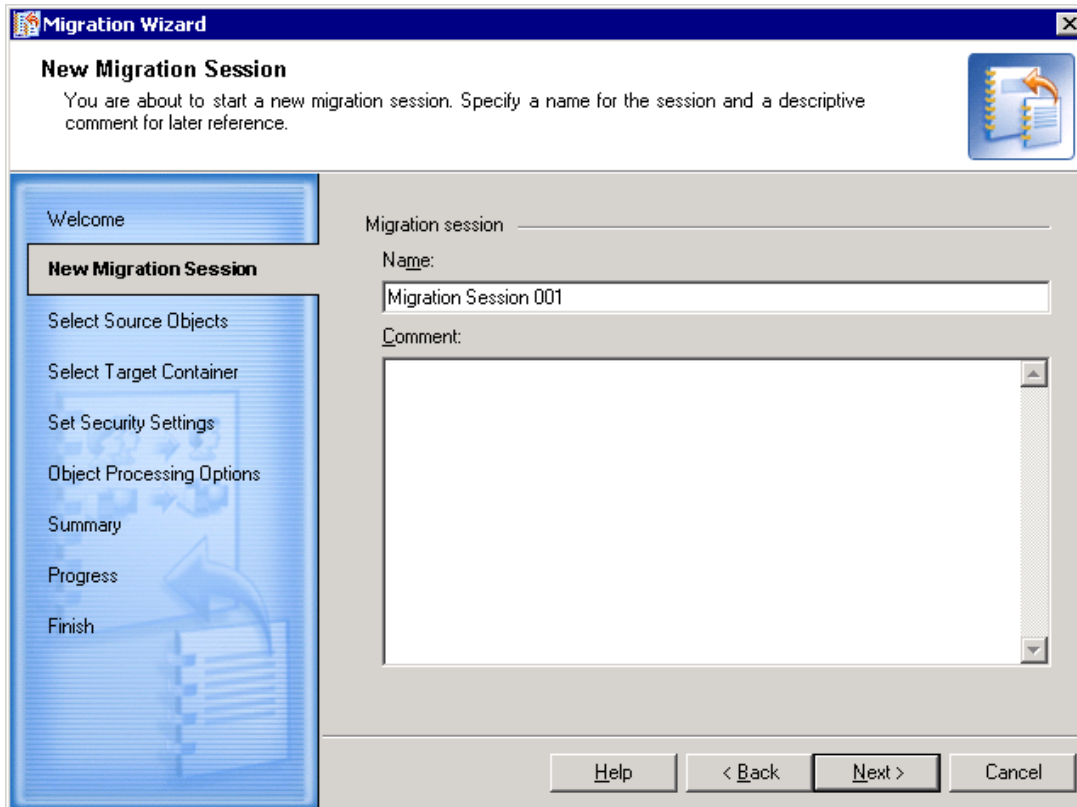
See the *Migration Manager Tips and Tricks* document for additional considerations and recommendations for setting up migration sessions.

Creating a Migration Session

To create a new migration session, right-click the **Migration** node under the appropriate domain pair and select **New Session** from the shortcut menu. This will start the **Migration Wizard**.

Step 1. New Migration Session

Specify a name for the migration session and optionally provide a comment for it.

Step 2. Select Source Objects

This step allows you to select objects for migration. The list of objects contains the objects currently selected for migration.
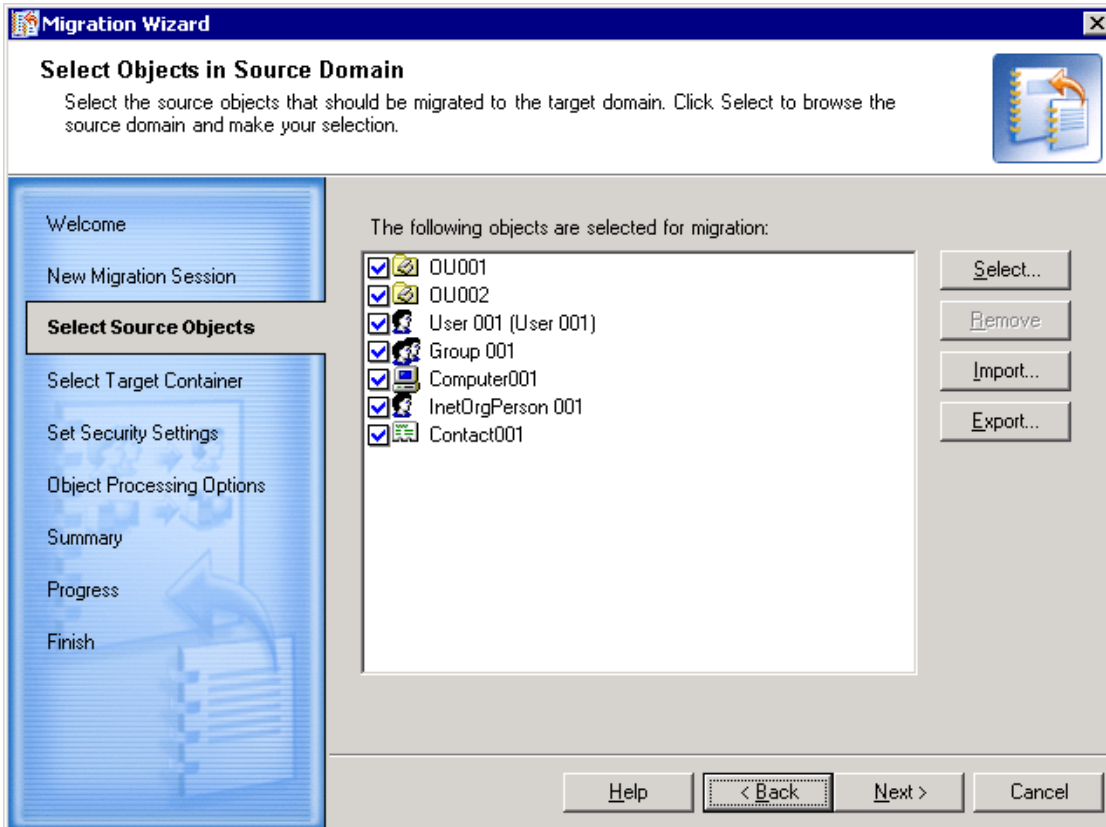
Click **Select** to add to or modify the selection. In the Browse Source Domain window, select the containers and/or individual objects you want to migrate.
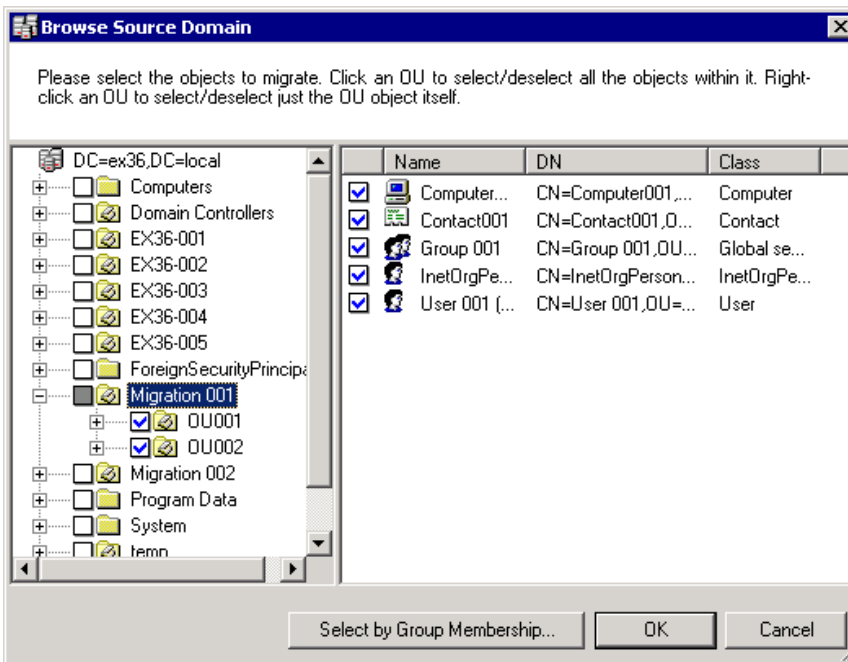
When you select or clear a container, all its sub-objects are displayed in the right-hand pane and are automatically selected or cleared. If you want to select or clear individual objects, perform the selection in the list of objects in the right-hand pane.

To select or deselect containers, right-click the container and click **Select** or **Deselect**. Depending on your selection, the checkboxes will be marked as follows:

- A blue check mark indicates that the container and all objects and sub-containers within the container will be migrated.

- A white check mark indicates that the container will be created on the target and some objects and sub-containers within the container will be migrated as well.

- A grey check box without any check mark indicates that the container will not be created on the target but some objects or sub-containers within the container will be migrated.

Clicking the **Select by group membership** button offers you two options:

- **Select groups for selected objects**—Select the groups in which the selected objects are members. You can also specify the scope (global, local, or universal) and type (security or distribution) of groups for which to perform the select operation.

- **Select objects within selected groups**—Select objects that are members of the selected groups. You can also specify the object class (users, contacts, or groups) for which to perform the select operation.
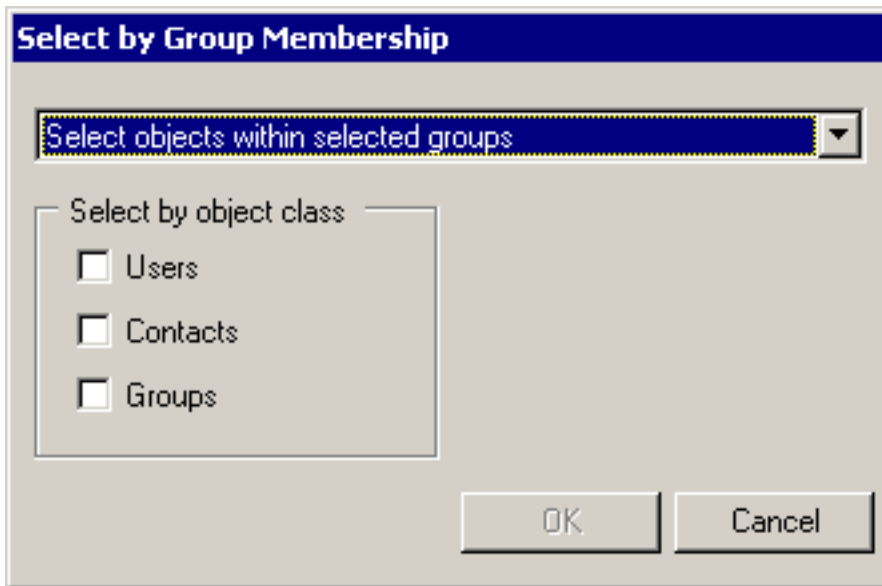
After the selection is made, click **OK** to save it and close the window.

To remove an object from the list of selected objects, select the object and click **Remove**.

You can also export the current selection to an external text file for later use. The external list of objects is usually used for mass object renaming and populating target object attributes with different values.

To create an export file, click **Export**. In the **Export Selection to File** window, select the attributes you want to export for the selected users. This creates a tab-delimited list: the first column is the source object's DN, and the remaining columns are the selected attributes. The administrator can later modify the attribute values in this file and import it back by clicking **Import**. The modified attribute values will be applied to the target objects during migration.
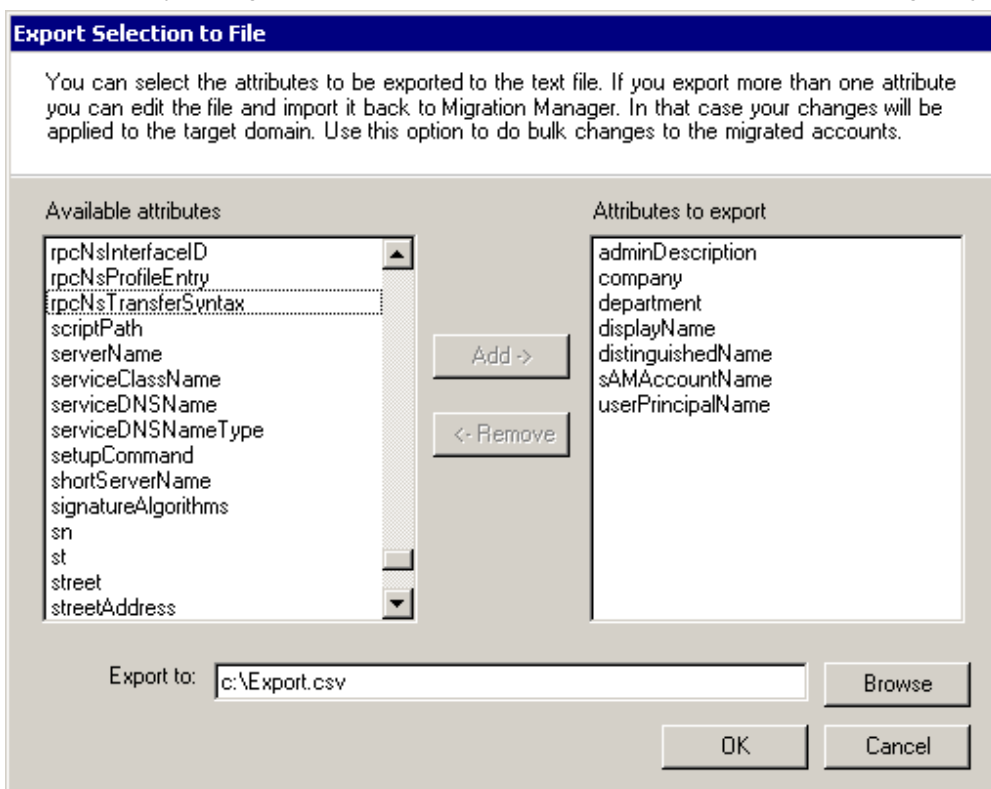
One common use for such import files is to make Migration Manager rename user accounts and groups as part of the migration. For details about how to edit the import files in this scenario, see the Configuring User and Group Renaming topic.

Step 3. Select Target Container

This step allows you to select the container where migrated objects will be created. Click **Browse** to select the container where the migrated objects should be placed during migration.
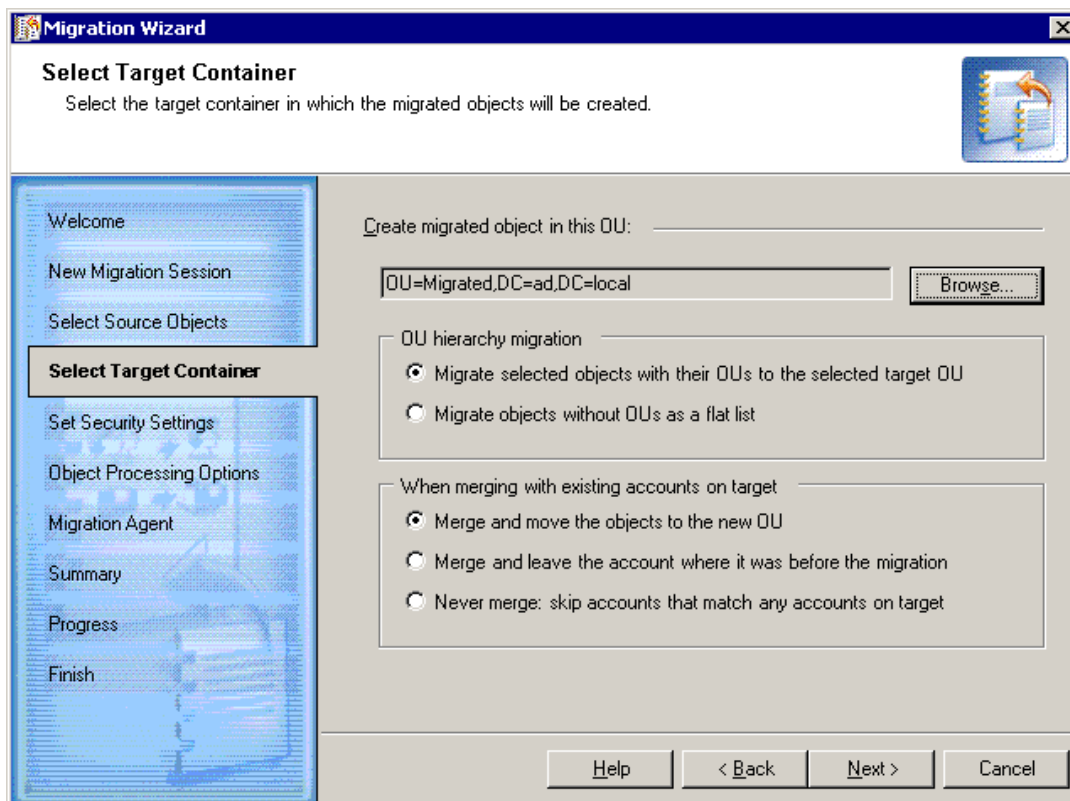
You also can specify whether the OU hierarchy will be migrated and whether the accounts should be merged with the existing target accounts.

In the **OU hierarchy migration** section, choose one of the following:

- **Migrate selected objects with their OUs to the selected target OU**—If this option is selected, all selected objects and containers will be created on the target.

- **Migrate objects without OUs as a flat list**—If this option is selected, only selected objects (not containers) will be created on the target in the specified target OU.

In the When merging with existing account on target section, choose one of the following:

- **Merge and move the objects to the new OU**—If this option is selected, matched objects will be merged and moved to the specified target OU.

- **Merge and leave the account where it was before the migration**—If this option is selected, matched objects will be merged and left in the original target OU (not moved to the specified target OU).

- **Never merge: skip accounts that match any accounts on target**—If this option is selected, matched objects will be skipped from migration.

Step 4. Set Security Settings

This step allows you to specify the security settings for the migration.
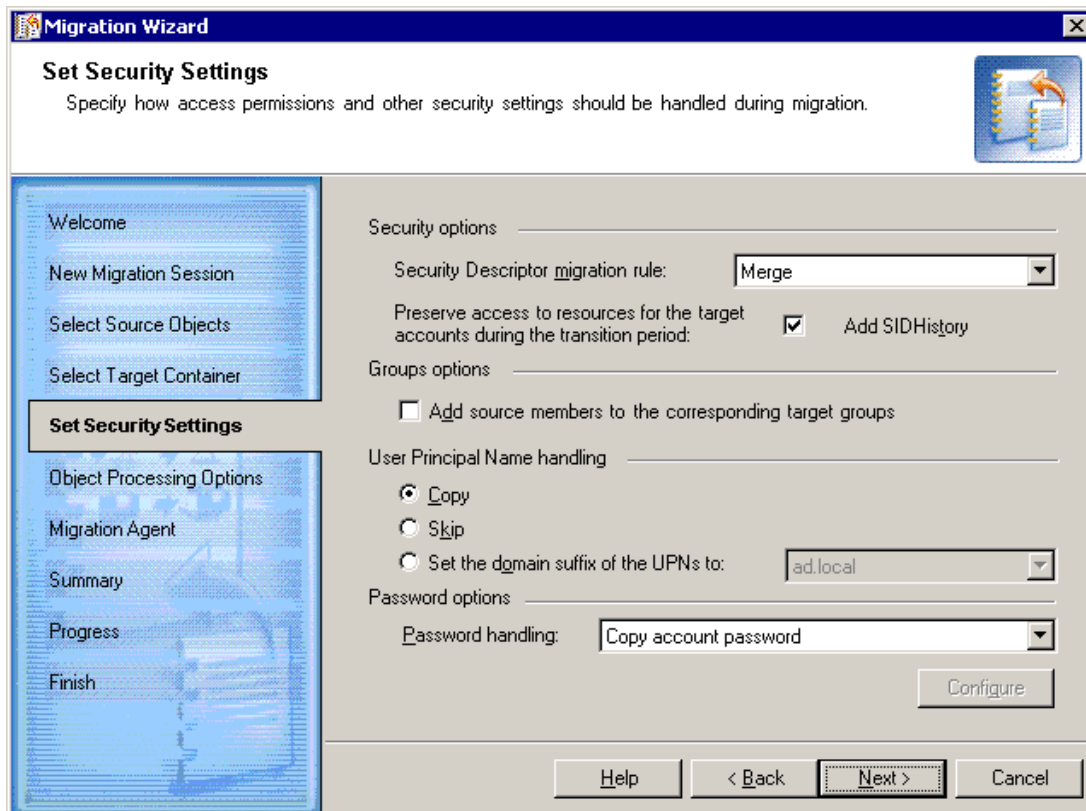
Security Descriptor migration rule—Select the way security descriptors of the matched source and target objects will be handled.

All objects in Active Directory are securable objects. Each securable object has a security descriptor (SD) that identifies the object's owner and can also contain the following access control lists:

- A discretionary access control list (DACL) that identifies the users and groups allowed or denied access to the object

- A system access control list (SACL) that controls how the system audits attempts to access the object

An ACL contains a list of access control entries (ACEs). Each ACE in an ACL identifies a trustee (a user account or group account) by its SID and specifies the access rights allowed, denied, or audited for that trustee.

You have the opportunity to **Merge**, **Replace**, or **Skip** the security descriptors:

- **Merge**—The security descriptor entries of the source object will be added to the security descriptor of the target object.

- **Skip**—The security descriptor of the target object will be left intact.

- **Replace**—All entries of the target object's security descriptor will be deleted. The entries of the source object' security descriptor will be copied to the target object's security descriptor.

The DACL and SACL security descriptor entries of the source objects are assigned to the newly-created target objects during migration.

Regardless of the option you select to migrate security descriptors (**Merge**, **Skip**, or **Replace**) for each newly created target object, the default security descriptor defined for that object class will also be applied.

> ℹ **NOTE:** Only ACEs explicitly added to the source security descriptor are migrated.
>
> The inheritance flag (the **Allow inheritable permissions from parent to propagate to this object** option on the **Security** tab of the object **Properties**) is migrated as well. That is, if the inheritance flag is set for the source object, it will be set for the corresponding target object; if the inheritance flag is not set for the source object, it will be cleared from the corresponding target object.

During migration, the ACEs of the source security descriptor referencing the source objects (source SIDs) are not translated to the target objects (target SIDs). To translate or clean up the source objects' SIDs migrated to the target object's security descriptor, use the **Active Directory Processing Wizard**.

Add SIDHistory—Select this checkbox if you want to allow the target accounts to access the source domain resources using the SIDHistory mechanism during the coexistence period. For more information on SIDHistory adding, see Adding SID History.

When you migrate accounts and groups, target group membership is automatically updated for the target users. In other words, the target group will have target user accounts as members corresponding to the source user accounts (members of the source group) migrated by that time. If you also want to add source accounts (the members of the source groups) to the corresponding target groups, select the **Add source members to the corresponding target groups** check box.

The **User Principal Name handling** section allows you specify how the User Principal Name (UPN) will be formed for each target user:

- **Copy**—If this option is selected, UPNs of source users will be assigned to the target users. This option is available only if the source and target domains belong to different forests.

- **Switch**— If this option is selected, the UPN is switched from the source user to the target user. This option is available only if the source and target domains belong to the same forest.

- **Skip**—If this option is selected, the target user UPN will be left intact.

- **Set the domain suffix of the UPNs to**—This option allows you to set the domain suffix of the UPNs of the target users to the value you specify.

The **Password handling** option allows you to also specify how user passwords will be handled:

- **Copy account password**—Passwords will be copied from the source to the target accounts.

- **Skip account password**—Passwords will not be copied for merged objects. The newly-created target accounts will get blank passwords.

- **Set password to username with**—Sets the target user password to its username with the specified prefix and/or suffix. To set the prefix or suffix, click Configure.

- **Set password to**—Sets the target user password to the specified value. To specify the common password value, click Configure.

- **Set random password**—Sets the password to a random value generated by some criteria. The passwords are stored in ADAM or AD LDS. You can select to generate strong or custom passwords. For custom passwords, you can set the range for the password length and the allowed characters.
  To configure the complexity of the random password, click **Configure** and use the **Random length between** and **Allow characters** controls:
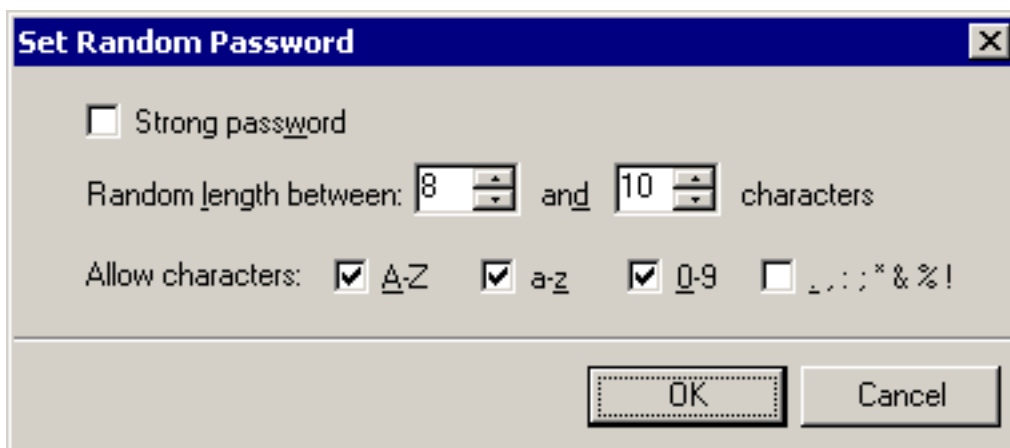
Selecting **Strong password** will disable the **Random length between** and **Allow** characters controls and make the generated passwords comply with the password requirements specified in Microsoft Knowledge Base article Q161990. According to the article, passwords must satisfy the following requirements:

- Passwords must be at least six (6) characters long. You can set the password length limit by editing the **Random length between** boxes.

- Passwords must contain characters from at least three (3) of the following four (4) classes:

  - English uppercase letters (A, B, C, ... Z)

  - English lowercase letters (a, b, c, ... z)

  - Westernized Arabic numerals (0, 1, 2, ... 9)

  - Non-alphanumeric or 'special characters', such as punctuation symbols

  - Passwords may not contain a user name or any part of a full name.

ℹ **NOTE:** To let users know their temporary passwords so that they can change them, you first need to retrieve the passwords from ADAM or AD LDS. For that, use the utility provided by Quest Support in solution 32124, available at https://support.quest.com/kb/SOL32124.

Step 5. Specify Object Processing Options

Specify whether the target objects should be enabled after the migration session is completed. This setting makes sense if the users start using their target accounts immediately after the migration is completed. You may also want to disable source accounts after migration has been completed. Use the **Enable target accounts** and **Disable source accounts** options.
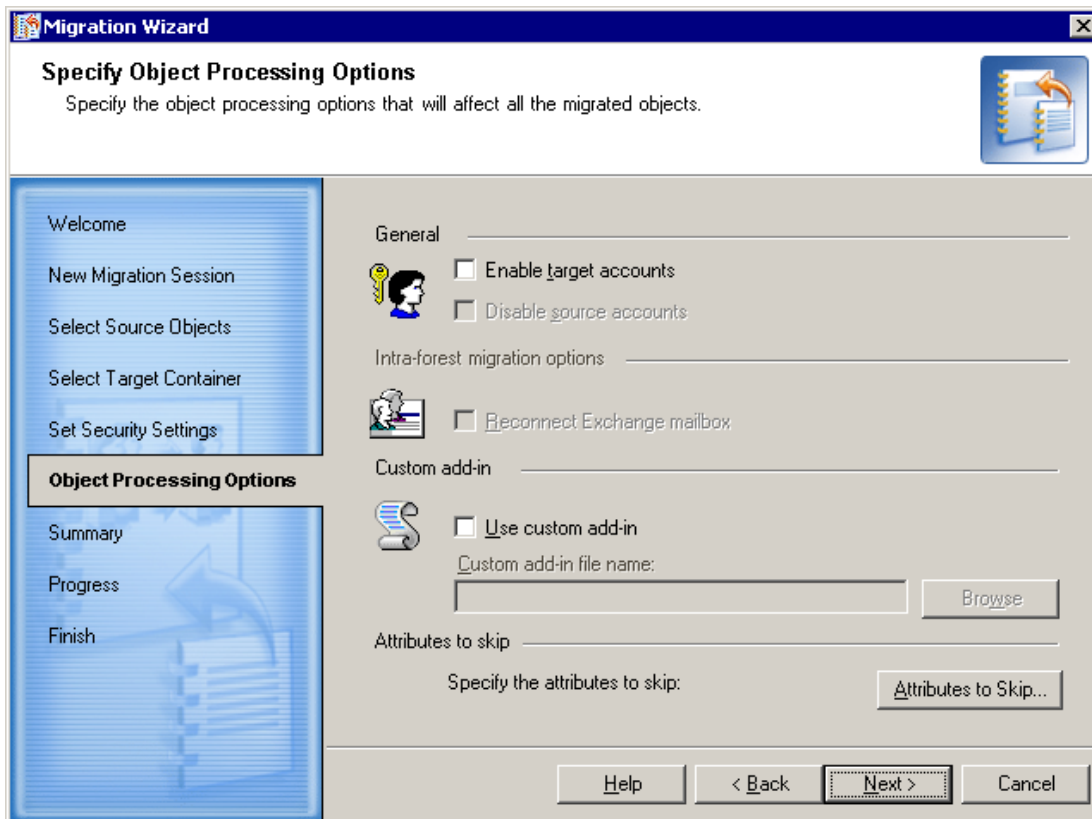
For intra-forest migrations, you can select to reconnect the source Exchange mailboxes to the target users so that users logged on to the target environment can use their source mailboxes until the Exchange migration is performed. Selecting **the Reconnect Exchange mailbox** check box will reconnect the source user mailboxes to the corresponding target users.

If you are planning to use a custom add-in to process the selected objects after they have been retrieved from the source domain and before they are copied to the target domain, select the **Use custom add-in** checkbox and browse for the **.xml** custom add-in file.

Attributes to Skip

If you want to skip particular attributes, click the **Attributes to Skip** button. Then select the check boxes next to the attributes you do not want to migrate to the target domain.

Use the **Save Settings** button if you want to use the same settings when you create another migration session. Click **Load Settings** to apply a set of attribute skipping options you saved earlier.
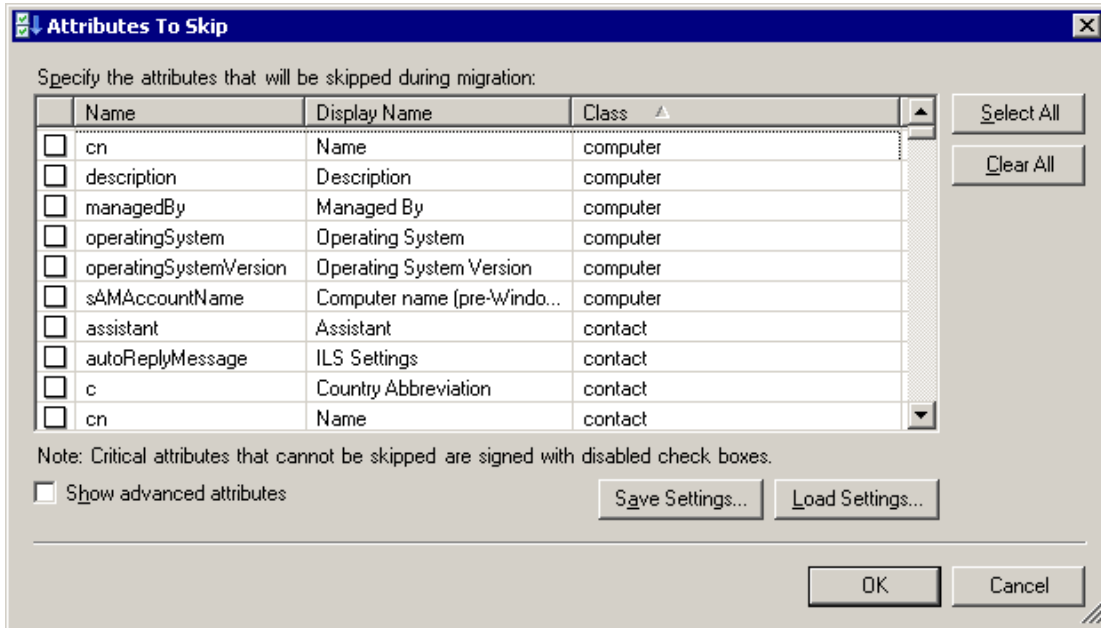
Select the **Show advanced attributes** check box to choose from the complete list of attributes you can skip.

> ⚠ **CAUTION: The Directory Synchronization Agent service attributes used by Migration Manager should never be skipped. Otherwise, Migration Manager will not be able to migrate and synchronize objects.**

Step 6. Select Migration Agent

This step allows you to select the migration agent to perform the migration.

> ℹ **NOTE:** If you have only one agent installed in your environment, you will not be presented with this step.
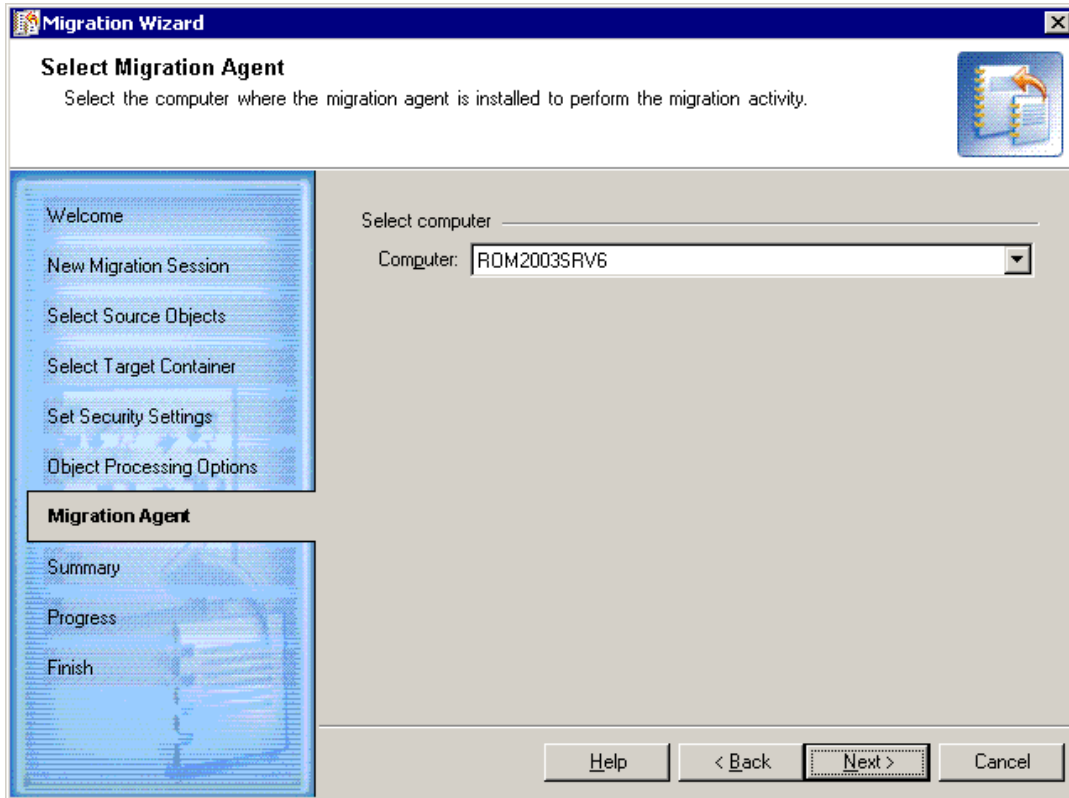
Step 7. Summary

The wizard allows you to view the settings you made for the session. You can click **Back** to modify the settings if needed.
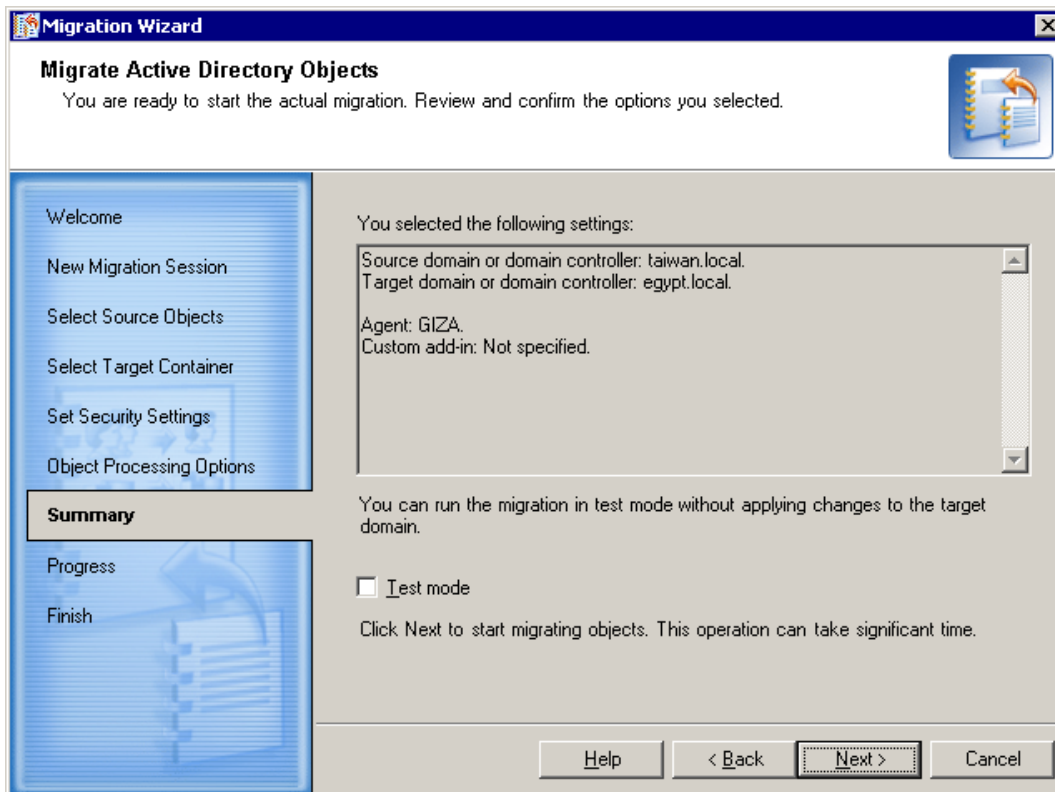
Select the **Test mode** check box if you want to run the migration in test mode, which does not apply any changes to the real target environment. Running a migration session in test mode allows you to check how the settings you made for the session will be applied to the target directory and decide whether these settings suit your needs. You can review the results by clicking **View log** after the migration is completed.

> **i** | **NOTE:** You can click **Back** and clear the **Test mode** check box to perform the actual migration only if you did not close the Migration Wizard dialog after the migration completed. You cannot clear the **Test mode** check box when you open the properties of the already-completed session. However, you can use this session as a template when you create a new session to perform the account migration. Refer to the Viewing Migration Session Details topic for more details.

Step 8. Migrating Active Directory Objects

The wizard now migrates the selected objects to the target domain. All the activity takes place in the target domain only. The wizard displays the target domain directory update progress. Please wait while the wizard completes.

Step 9. Complete the Wizard

The migration session has been completed. The completed session configuration is now stored in the project database. Click the **View log** button to see if any errors or conflicts occurred during migration.

You can view the completed session configuration and use it later as a template for other migration sessions. For more information about migration session details, refer to the Viewing Migration Session Details topic.

All changes made to the target environment during a migration session can be rolled back. For more information about undoing a migration, refer to the Undo Account Migration topic.

Viewing Migration Session Details

To view the migration session details, select the **Migration** node in the Migration Manager console management tree. The list of completed sessions is displayed in the right-hand pane. Right-click the session whose details you want to view and select **Properties** from the shortcut menu.

You can view the log of a completed session to see the results of migration by selecting the **Summary** tab and clicking **View log**.

Using a Completed Session as a Template

You can use a completed session as a template to create a new session with the same or similar configuration settings or similar objects to be migrated. For example, you might want to select a session that was previously run in test mode and perform the actual migration of objects using the settings specified in that session.

To use a completed session as a template, select the **Migration** node in the Migration Manager console management tree. The list of completed sessions will be displayed in the right-hand pane. Right-click the session you want to use as a template and select **New Session** from the shortcut menu. This will start the **Migration Wizard**. All the settings you made for the completed session, including objects selected for migration, are preserved and you do not need to specify these settings again.

Configuring User and Group Renaming

You have the option of renaming users and groups as part of the migration. This includes merging multiple source users or groups into a single target user or group.

Renaming of objects is performed using a specially formatted plain-text file, which should be imported on the Select Objects in Source Domain step of the Migration Wizard (see the *Creating a Migration Session* procedure).

All entries in the file should be tab-separated. The header is mandatory and should contain at least these two entries (note that the white space is a tab character):

```
SAMAccountName SAMAccountName
```

One object per line should be specified. Here is a sample import file:

```
SAMAccountName SAMAccountName

sourceuser01 targetuser1

sourceuser02 targetuser2

sourceuser03 targetuser3

sourceuser04 targetuser4
```

> **! CAUTION:**
>
> - **If you are using this file to create a new user account and there is already a target domain user with the same name as the source user, the file above will not work. It is absolutely necessary to add at least the name attribute; then the file should look like this:**
>
> ```
> SAMAccountName SAMAccountName name
>
> sourceuser01 targetuser01 targetuser1
> ```
>
> - **If object names contain spaces, there is no need to use quotation marks (spaces are not treated as delimiters). For example, to rename the Executive Users group to Target Executive Users, use this syntax:**
>
> ```
> SAMAccountName SAMAccountName name
>
> Executive Users Target Executive Users Target Executive Users
> ```

The same file format can be used to populate other target attributes; here is a syntax sample:

```
SAMAccountName SAMAccountName name displayname userprincipalname

sourceuser1 targetuser1 targetuser1 Target User1 targetuser1@dom.com

sourceuser2 targetuser2 targetuser2 Target User2 targetuser2@dom.com

sourceuser3 targetuser3 targetuser3 Target User3 targetuser3@dom.com
```

The same import file can be used for merging with existing objects. When using such a file, you can merge two completely different accounts, for example Peter M. from the source domain can be merged with James T. from the target domain. The same applies to groups. In rare cases, the same approach is used to merge a built-in group with another (regular) group. When merging users or groups using **sAMAccountName<tab>sAMAccountName**, make sure that matching by **sAMAccountName** is enabled in the domain pair properties; otherwise DSA will attempt to create a new user with the same **sAMAccountName** and post a conflict instead of doing a merge.

> **! CAUTION: The CN attribute cannot be changed using the CN as syntax. In order to modify the CN value, the Name attribute must be used. This will change the CN name on target.**

Finally, note the following specifics:

- The maximum user or group name in the file is 64 characters. If this number is exceeded, you will receive an error message during the migration. Here is a sample log entry of the error:

```
LDAP error 0x13. Constraint Violation (00002082: AtrErr: DSID-03050B04, #1:0:
00002082: DSID-03050B04, problem 1005 (CONSTRAINT_ATT_TYPE, data 0, Att 3
(cn):len 130).
```

- Import files work only for migration and are not applicable to synchronization.
- If synchronization is turned on and the attributes being changed during migration are not excluded, synchronization will overwrite the changes.
- When you migrate accounts using an import file, they will always be migrated as a flat list to the location you specify during the migration session no matter what is specified by the migration session settings. If you want to retain the OU structure of source domain for the listed objects, one possible option is to migrate the objects normally first, then rename them later using the import file and a merging migration session. When performing this, make sure the **Merge and leave the account where it was before the migration** option is selected.

Delegating Account Migration

If you want to re-assign the directory migration task to some other person without revealing key credentials (like administrative accounts and passwords to access source and target domains), you can delegate migration to that person and even limit the scope of the delegated migration to certain organizational units (OUs) belonging to the source and target domains.

The results and status information on the delegated migrations will be added to the project, so no matter how many delegated administrators are involved, you can keep track of the overall project.
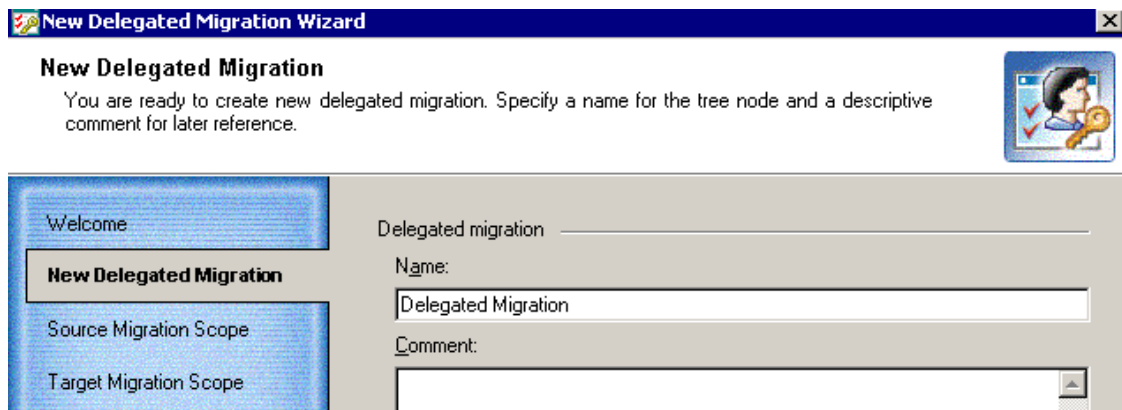
Since the delegated administrators get access to only the tasks to which you grant them access, you can be sure that they do not interfere with other tasks.

A person to whom the migration task is delegated is called a **Migration Admin**. This person will be able to create migration sessions using the Migration Wizard. However, the objects available for migration will be limited to the scope you specified when you delegated the task. The Migration Admin to whom the task was delegated will not be able to select the agent to perform migration or the custom add-in to process the objects being migrated.

To delegate a migration task, right-click the domain pair node and select **New Delegated Migration** from the shortcut menu. The **New Delegated Migration Wizard** will help you delegate a part of migration to another person.

Step 1. New Delegated Migration

Specify a name for the delegated migration and enter comments.

Step 2. Restrict Source Migration Scope

This step allows you to select the OUs from which the objects can be migrated within the delegated migration scope.
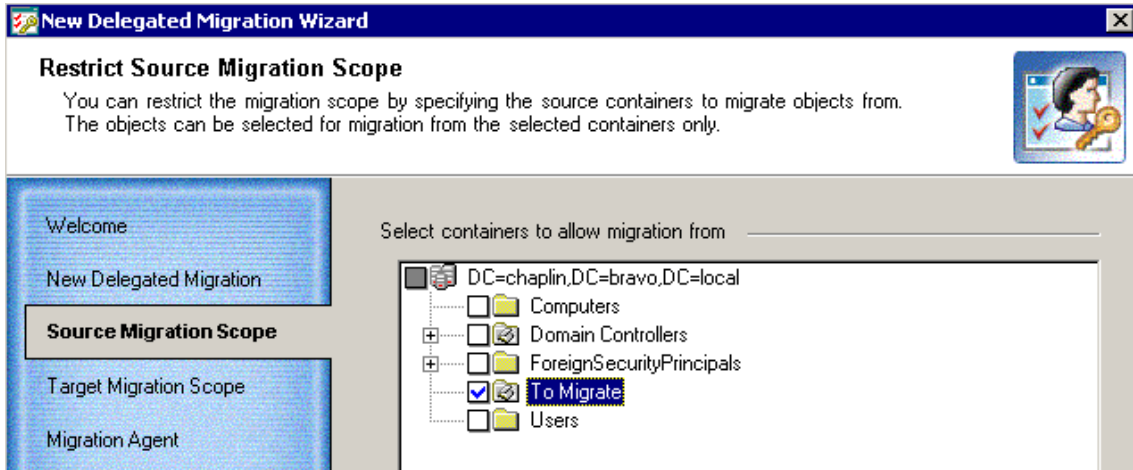
Step 3. Restrict Target Migration Scope

This step allows you to select the OUs to which the objects can be migrated within the delegated migration scope.
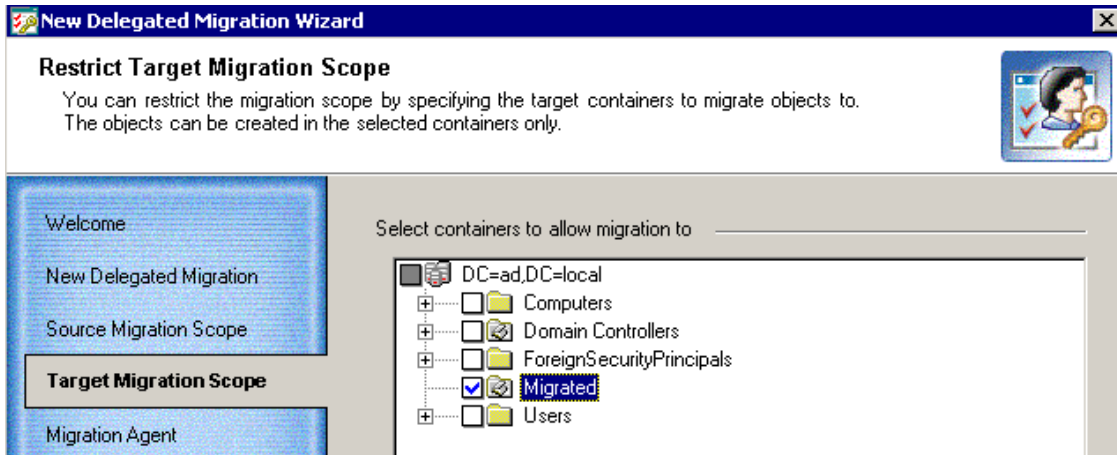
Step 4. Select Migration Agent

A person to whom the migration task is delegated will be able neither to install the Directory Synchronization Agent nor to select the agent's instance for performing migration. Therefore, in this step you should select the instance of the Directory Synchronization Agent which will perform the delegated migration.

ℹ NOTE: If you have only one agent installed in your environment, you will not be presented with this step.
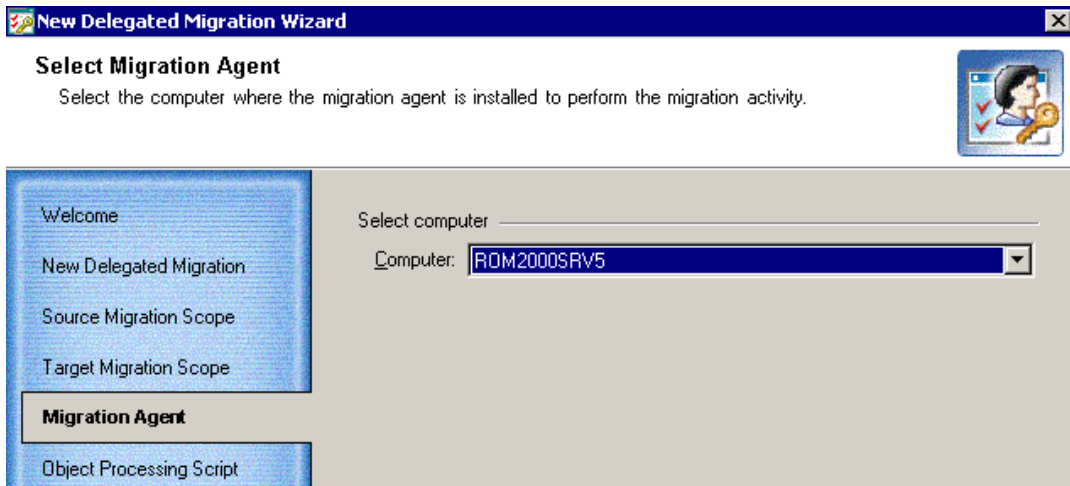
Step 5. Specify the Object Processing Custom Add-in File

A person to whom the migration task is delegated will not be able to select a custom add-in file for processing objects being migrated. Therefore, on this step you should select the custom add-in file if the migrated objects should be processed with the script before copying to the target server. Select the **Use custom add-in** checkbox and browse for the .**xml** custom add-in file.
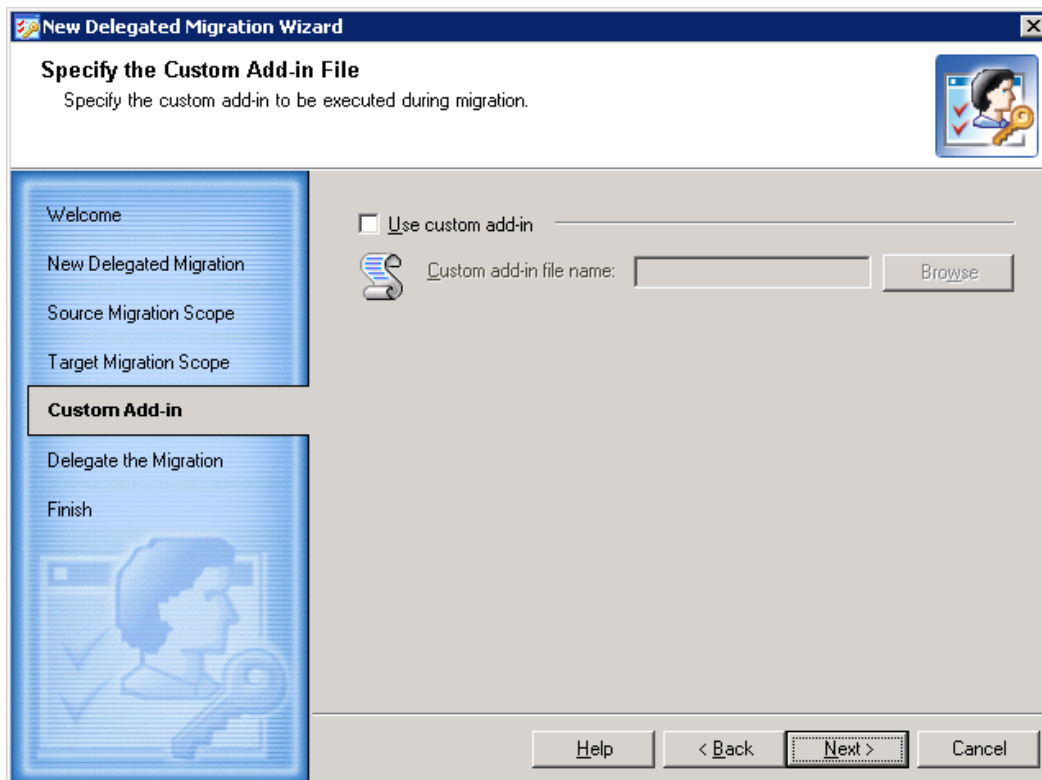
Step 6. Delegate the Migration

This step allows you to assign the Migration Admin role to the account you select. The Migration Admin will be able to create migration sessions and make migration settings within the scope you just selected.
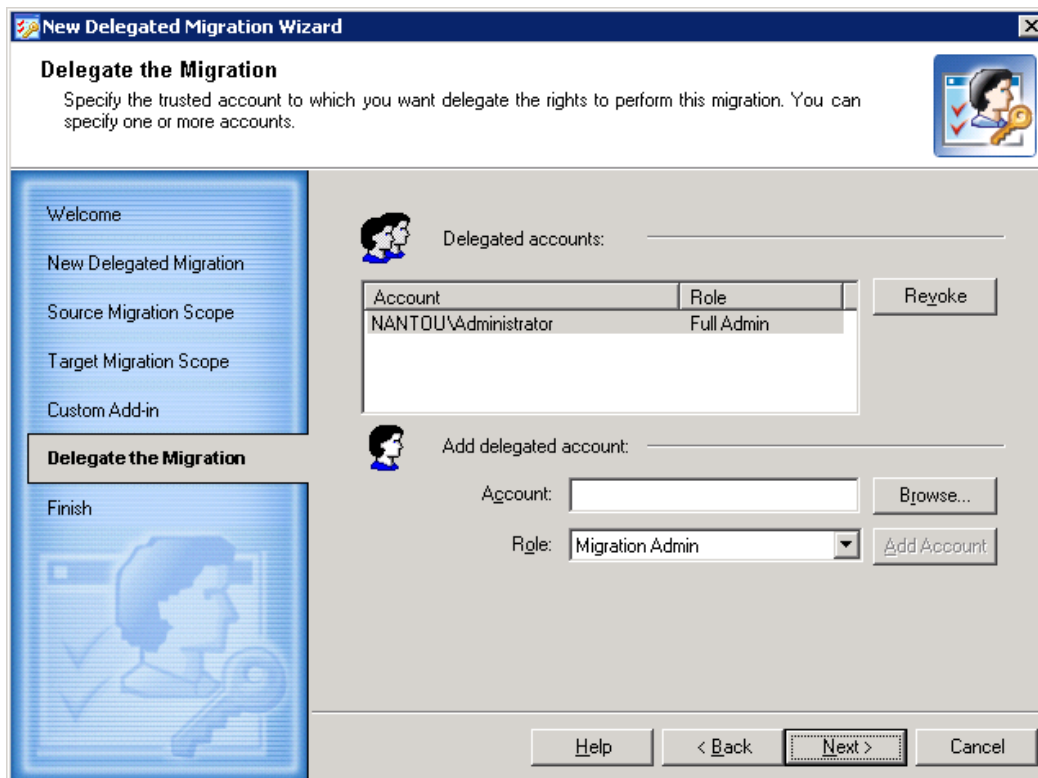
Step 7. Complete the Wizard

The wizard displays the delegation settings you made.

Undo Account Migration

You can roll back the changes made to the target directory by each migration session independently.

All the changes made to the target domain by the session will be rolled back exactly to the state before the session started. The following example illustrates this behavior:

1. Suppose that during migration session 1, you migrated a number of objects from the source domain and specified to merge them with the existing target objects and move them to organizational unit A on the target.

2. Then during migration session 2, you re-migrated the same set of objects and specified to merge them with the existing objects in the target and move them to target organizational unit B.

3. If you undo migration session 2, the target objects that were merged with the source objects will be moved back to organizational unit A, not to their original location as it was before migration session 1.

In other words, when you undo a session, only changes made by that session are rolled back. Another example of such behavior is the following:

1. Assume a user object was migrated to the target three times in three different sessions. In the last session the Reconnect Exchange mailbox option was selected to reconnect the Exchange mailbox from the source to the target user account.

2. If you undo the first or second session for the migrated user, the mailbox will not be reconnected back to the source account.

> **! CAUTION:**
>
> - **Account passwords changed during a migration session cannot be rolled back.**
>
> - **If you migrated an object to the target domain during a migration session and after that deleted the object along with the container where that object was stored in the source domain, then the object will not be restored in the source domain if you undo the migration session.**

To undo the results of a migration session, select the **Migration** node in the migration project tree, right-click the session displayed in the right-hand pane and select **Undo** from the shortcut menu. The **Undo Wizard** will start.

Step 1. Select Objects to Revert the Changes for

This step allows you to select the objects for which the migration should be undone.
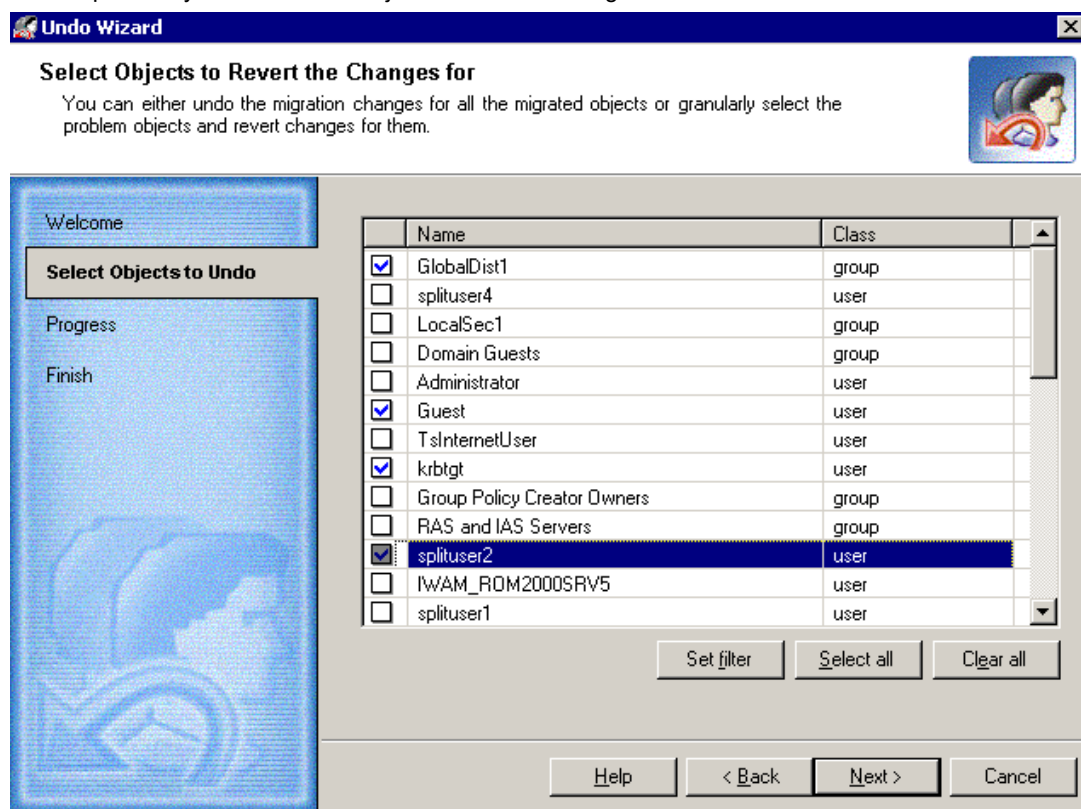
To select the objects, you can either click **Select all** or select particular objects by selecting the check boxes next to those objects. Click **Set filter** to specify the object classes for which you want to undo changes. All object classes are selected by default.

Step 2. Reverting Migration Changes

The wizard now removes the migrated objects you selected from the target domain and reverts changes made during the migration session. Wait while the wizard completes.

> **i** | **NOTE:** If the source object was merged with the existing target object during migration, the target object will not be deleted from the target directory during the undo. All changes made to the target object during the merge operation (including moving to another organizational unit and changing attribute values) will be rolled back to the state the object had before the migration and merge.

If errors occurred during the rollback process, you can view more detailed information by clicking **View log** and inspecting the log file for the failed objects.

Directory Synchronization

Before you start your migration, you should create the domain pairs that will be involved in the migration or synchronization process. Refer to the Creating a Domain Pair topic for more details.

If only Active Directory objects are being migrated from one forest to another, then directory synchronization is recommended but not required. If inter-org Exchange messaging system migration is performed, then directory synchronization is mandatory.

Since large migration projects can last for a long time, there is a period when the source and target environments have to coexist and be kept in sync. Migration Manager not only allows you to migrate accounts, but also to continuously synchronize these accounts and groups.

Considerations for Directory Synchronization

- If your migration scenario includes Microsoft Exchange messaging system migration, it is recommended that you set the directory synchronization to create and synchronize objects for which no match can be found in the target directory. This allows you to maintain a unified global address list (GAL) for source and target Exchange organizations, which is important for Exchange migration. See the Migration Manager for Exchange User Guide for more details.

- If you want to use Migration Manager to implement Exchange Resource Forest, configure the Directory Synchronization Agent to create disabled and mailbox-enabled accounts in the target domain. In this case, the Directory Synchronization Agent will automatically add the **Associated External Account** (msExchMasterAccountSID) attribute to the target accounts. This will let users who still log on to the source domain access mailboxes in the target domain.

- Ongoing synchronization brings over the changes detected in the directory after its last run.

Full re-synchronization deletes the last state information and resynchronizes the directories. It overwrites single-valued attributes and merges multi-valued attributes. For example, if a group has member User1 on the source and the corresponding group has User2 on the target, after you re-sync from source to target the target group will have both User1 and User2.

Directory Synchronization Agent

Both migration and directory synchronization tasks are handled by the synchronization engine called the Directory Synchronization Agent (DSA).

Please refer to the *System Requirements and Access Rights* document for details about operating systems supported by the Directory Synchronization Agent. It can reside locally on the administrator's workstation or can be installed on multiple computers in the network.

Agent Manager

Agent Manager allows you to install and uninstall the Directory Synchronization Agents in your network and specify configuration parameters for them. You can also see which computers already have agents installed and the migration and synchronization jobs processed by each agent.

To start Agent Manager, in the Migration Manager console. select the **Agent Manager** command from the Tools menu.

Installing the Directory Synchronization Agent

To install the agent, in Agent Manager select **Action | Install**, or click the **Install** button on the toolbar. In the Install Agent dialog, type the name of the server to which you want to install the agent, and click **OK**.
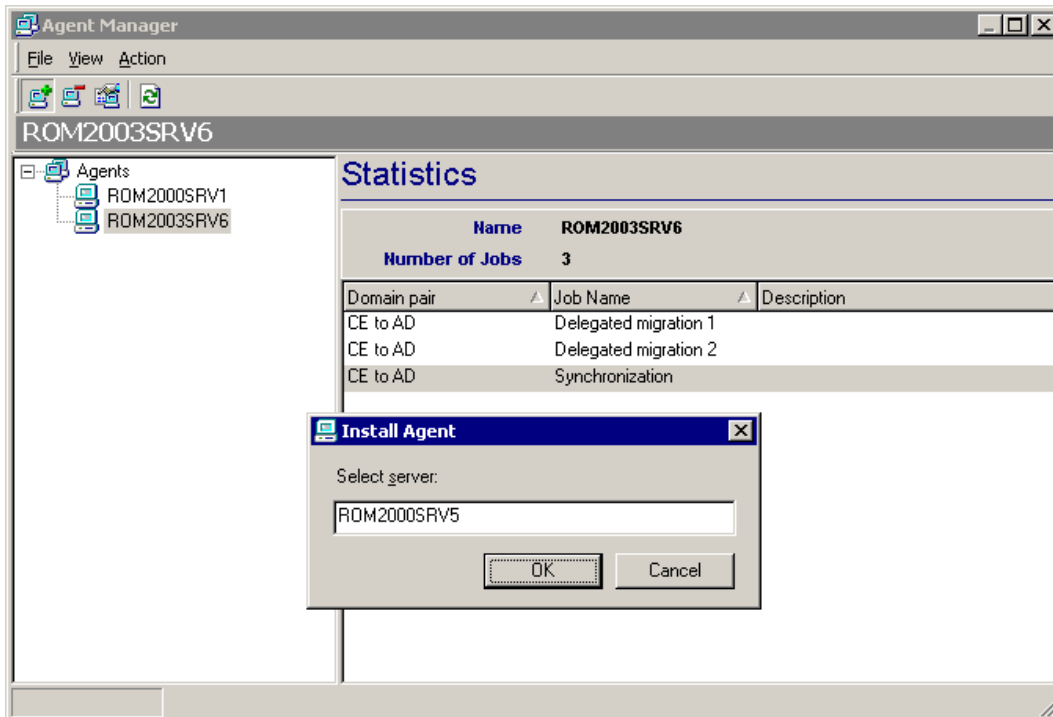
If you specify a remote server, a remote desktop connection is started and installation is performed in a terminal session. During agent installation, you are prompted for the ADAM or AD LDS instance, the credentials to access the ADAM/AD LDS instance and the project the agent will work with.

> **!** **CAUTION: Before installing the agent, make sure that the ADAM/AD LDS port is opened on firewall for both inbound and outbound connections and the user you specify has the Log on as a service right on the server where you plan to install the agent.**

> **i** **NOTE:** Terminal Services must be running on the remote server to which you want to install the agent. Trust relationships must also be established between the domains where the console machine and the remote server are members.
>
> If for some reason you cannot install remotely, you can install the DSA by running the agent setup package (.msi) locally on that server. The agent setup package is located by default in the \Program **Files\Quest Software\Migration Manager\Common\BIN\DeployDistr** folder on the console and is also accessible through automatically created share **\\<ConsoleComputer>\DSASetup** share.

On the specified server the **Directory Synchronization Agent Installation Wizard** will start. Complete the wizard to install the DSA.

To uninstall a selected agent, select **Uninstall** from the **Action** menu or click the **Uninstall** button on the toolbar. You also can uninstall an agent by selecting the **Uninstall** option from the agent's shortcut menu. Note that if you remove the last agent in a project, you cannot perform any migration activity until a new agent is installed.

Configuring the Directory Synchronization Agent

For each Directory Synchronization Agent, you can set a number of parameters. To configure the agent, in Agent Manager, right-click the agent and select **Properties**.

The **Preferences** tab of the agent **Properties** dialog box displays the list of domains that take part in the current project and allows you to specify the preferred domain controllers and Global Catalog servers for each. To decrease possible delays and network traffic during migration and synchronization, for each Directory Synchronization Agent specify the domain controller and Global Catalog server located in the same site as the agent.
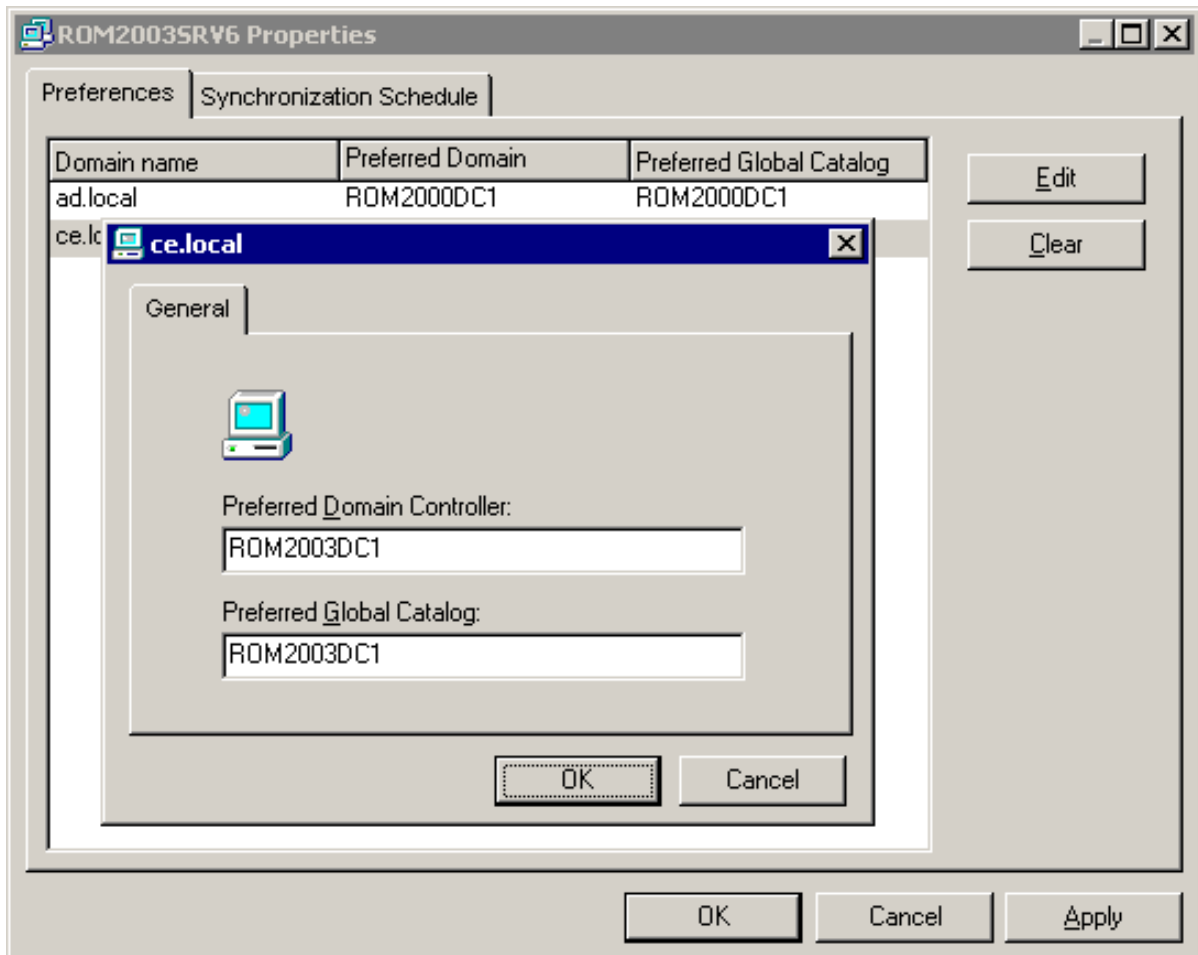
Click **Edit** on the selected domain in the list on the **Preferences** tab to specify the preferred DC and Global Catalog server for the domain.

Click **Clear** to remove the preferred DC and Global Catalog server settings for the selected domain.

You can also specify the time periods when the agent is allowed to perform its synchronization jobs. To do this, in the agent **Properties** dialog box, click the **Synchronization Schedule** tab.
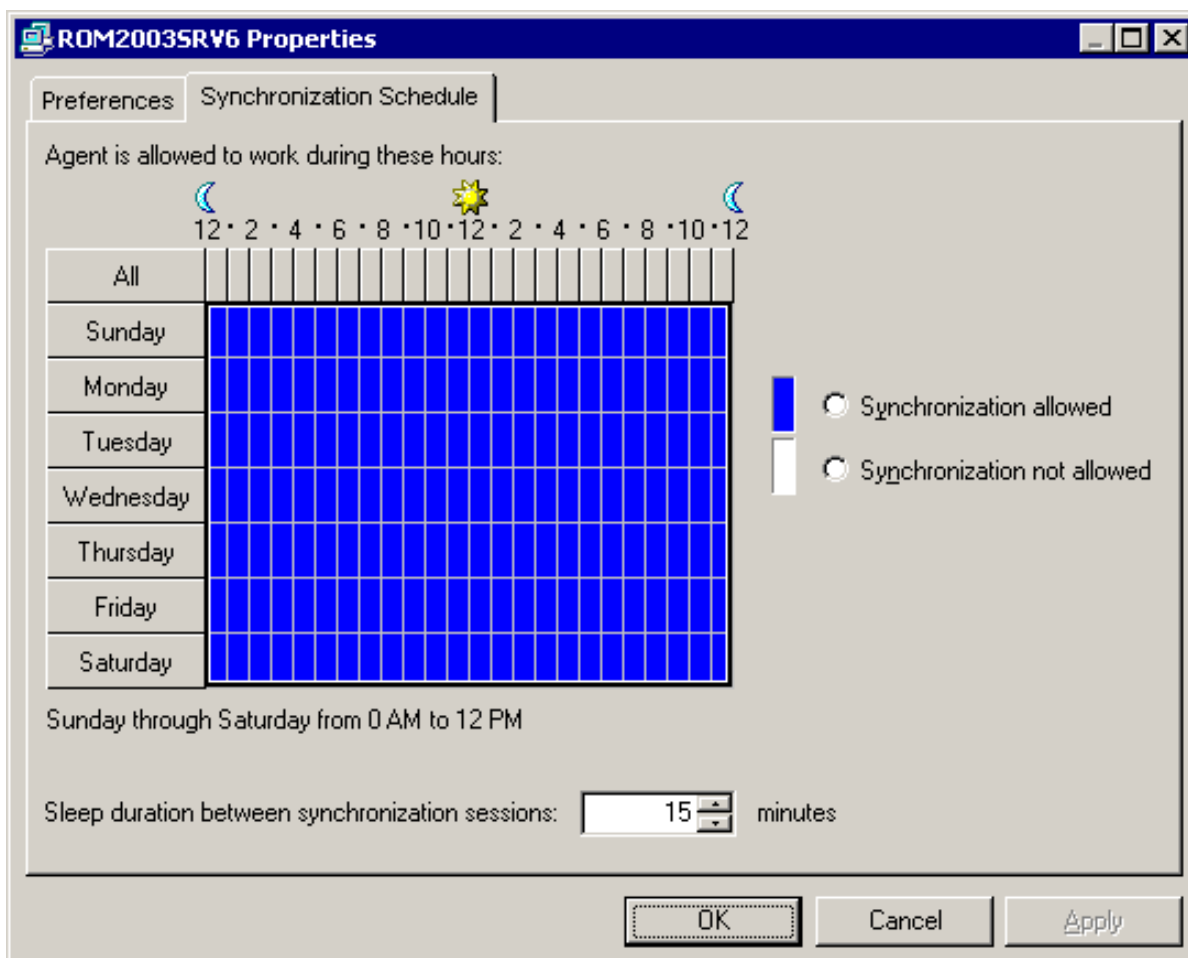
> **i** | **NOTE:** The synchronization schedule you specify does not affect account migration. Migration jobs are processed by the Directory Synchronization Agent regardless of whether the agent is allowed to process synchronization jobs or not.

For example, you may want to prohibit the agent from processing synchronization jobs during normal business hours. To specify the allowed and not allowed hours, select the area in the schedule map and select either the **Synchronization allowed** or the **Synchronization not allowed** option.

The directory synchronization jobs are processed by the agent in sessions. After each session completes, the agent is in an idle state for a time. The default sleep interval is 15 minutes. You can adjust this interval using the **Sleep duration between synchronization sessions** control.

> **i** | **NOTE:** If an agent is running a session and the time period when the agent is not allowed to process directory synchronization jobs begins, the agent will stop processing the current session and enter an idle state until such time as it is permitted to work again. At that time, the agent will continue processing the session from the point where it stopped.
>
> You can set the preferred DC and Global Catalog server, synchronization schedule, and sleep interval for each Directory Synchronization Agent independently.

Changing the Directory Synchronization Agent Credentials

During the Directory Synchronization Agent (DSA) installation, the credentials which used by DSA to connect to the ADAM or AD LDS instance are specified. If Migration Manager was installed using Express setup then DSA was installed automatically on the same computer as the Migration Manager Console and in this case DSA uses the auxiliary account to connect to ADAM or AD LDS instance.

Although it is not recommended to change DSA instance credentials, in some cases such as password expiration you may need to change them. To change the DSA instance credentials, take the following steps:

1. In the Agent Manager select DSA instances whose credentials you want to change.

   > **!** **CAUTION: The following requirements must be met on each remote computer where DSA credentials are to be changed:**
   >
   > - **The Remote Registry service must be started.**
   > - **The account under which Migration Manager console is running must have rights to edit the registry on the remote computer.**

2. Right-click the selection and then click **Change Credentials**.

   > **i** **NOTE:** To change the credentials for all DSA instances at once, you can right-click the **Agents** node in the left pane of the Agent Manager.

3. Specify the new credentials which will be used by selected DSA instances and click **OK**.

Configuring the Synchronization Job

The default synchronization job is set for a domain pair as soon as the domain pair is registered in Migration Manager. However, this job is not assigned to any Directory Synchronization Agent, and has only default settings that should be modified to suit your needs.

To configure the synchronization job, right-click the **Synchronization** node under the appropriate domain pair and select **Properties** from the shortcut menu.

Step 1. Select Synchronization Agent

Select the synchronization agent that will process the synchronization job between the domains in the domain pair. You can select any agent that is installed. If you want to use an agent that is not yet installed, you must first install it using the Agent Manager. Refer to the Directory Synchronization Agenttopic for details.
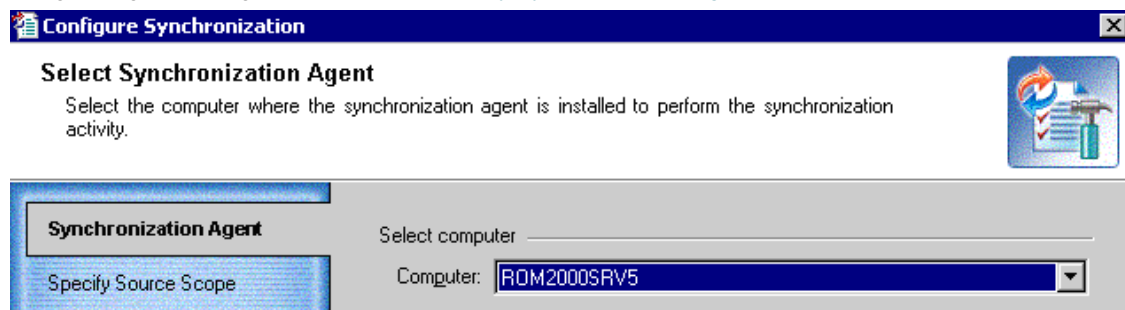
Step 2. Select Source Objects to Synchronize

This step allows you to select the source containers for which objects should be synchronized with the corresponding target objects.
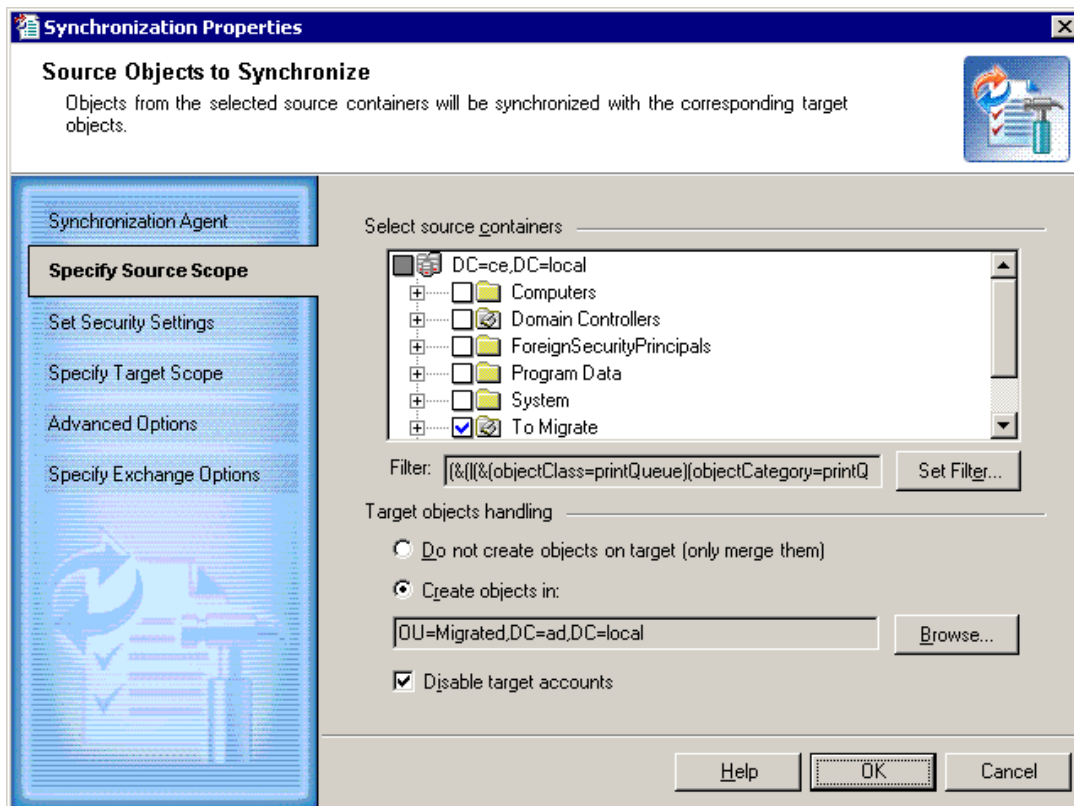
> **!** **CAUTION:** **The OU hierarchy and the movement of objects between OUs are not synchronized. Both of these operations can be performed during migration sessions.**

Select the containers in the displayed source directory tree:

- A blue check mark indicates that all objects from the selected container and all subcontainers will be synchronized.

- A white check mark indicates that only explicitly selected containers will be synchronized.

- A grey check box without any check mark indicates that the container will not be synchronized but some of its sub-containers are selected for synchronization.

Set Filter—This button allows you to filter the specific objects to be synchronized from among all objects from the selected containers. For example, you might want to synchronize only users and groups and not synchronize computers, contacts, and all other objects.

Object Class—On this tab, you can select from the list the object classes that you want to be synchronized.

Exclude List—This tab allows you to explicitly select individual objects to be excluded from synchronization. Click **Select**, and then browse for and select the objects you want to exclude. You can also import objects to be excluded from a plain-text file by clicking Import. In the text file, specify object **sAMAccountnames** or **distinguishedNames**, one per line.

Advanced—You can specify a custom LDAP filter expression here. Type the LDAP query in the window and click **OK**.
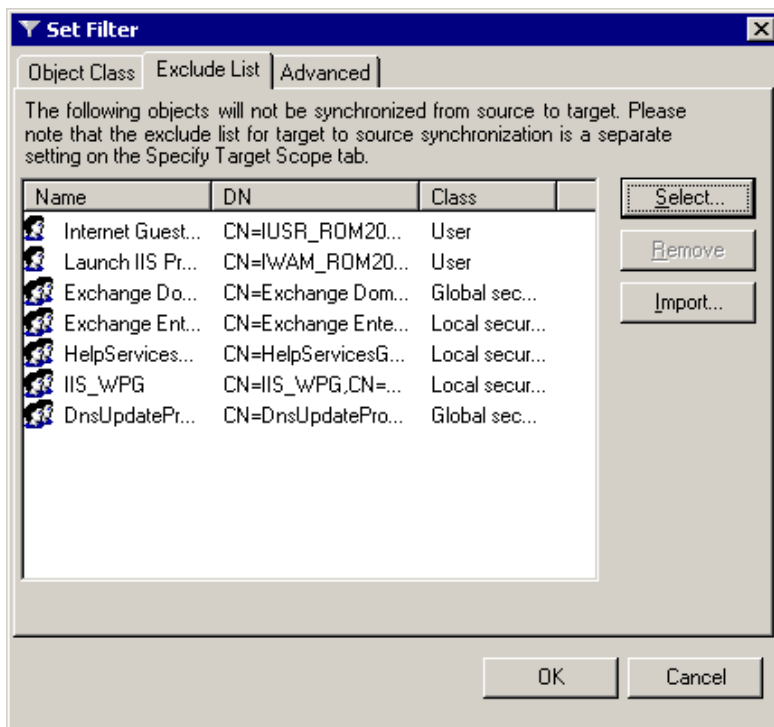
Do not create objects on target (only merge them)—If this option is selected, no new objects will be created on target during synchronization. Only objects that already exist on target will be synchronized.

Create objects in—You can select whether the agent should create the objects on the target and select the container where the objects should be created. If this option is selected, all the new objects that do not exist on target will be created in the specified container. Click **Browse** to select the container from the target directory tree.

> **ℹ NOTE:** If you select the **Create objects** in option, then for each source object for which the Directory Synchronization Agent cannot find a match in the target domain, it will create a new object in the OU you specify.
>
> If Exchange options are also configured in the directory synchronization job, the Directory Synchronization Agent may make target accounts mailbox-enabled like their corresponding source accounts. This depends on the options you set.
>
> However, if a conflict by e-mail addresses (the **proxyAddresses** attribute) arises when creating a mailbox for the newly-created target account, the agent will behave as follows:
>
> - If the conflict is not critical (for example, some e-mail addresses cannot be applied for the target account), the account will be created in the target domain and mailbox-enabled normally. However, it will not have the conflicting addresses in its proxy address list.
>
> - If the conflict is critical (for example, the agent cannot apply the target address for mail-enabled objects), the account will still be created in the target domain. It will not be mail-enabled and will be added to the failed queue.

Disable target accounts—By default the agent creates the target accounts as disabled accounts. If you want the state of the target account to be copied from the source account, clear this check box.

Step 3. Set Security Settings

This step allows you to specify the security settings for the synchronization.
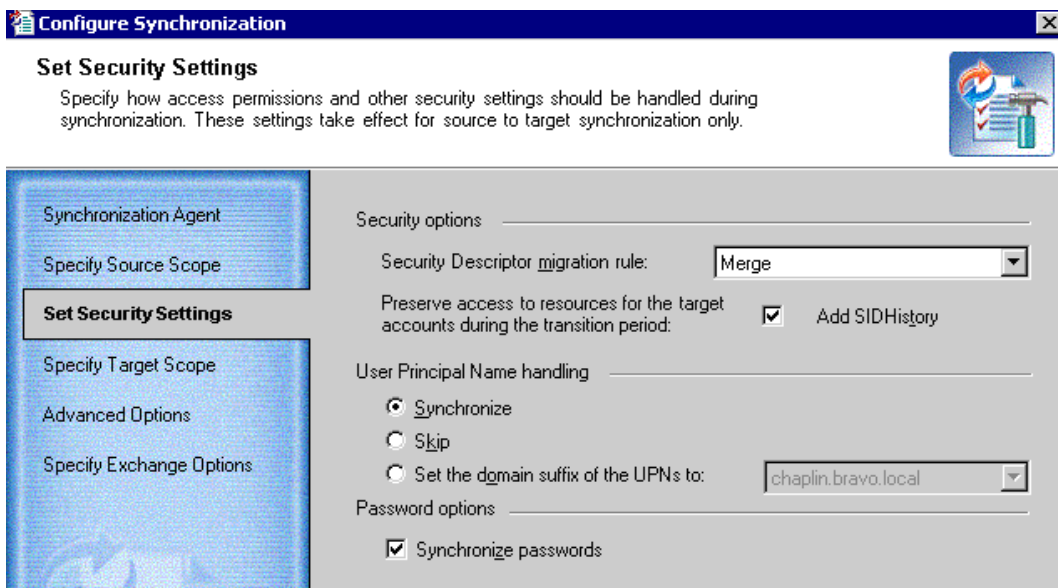
Security Descriptor migration rule—If two objects, source and target, are matched during synchronization by any matching criteria, you can select the way security descriptors of these objects will be handled.

For more information about objects matching criteria, refer to the Configuring a Domain Pair topic.

If the accounts are merged during synchronization, you have the opportunity to **Merge**, **Replace**, or **Skip** the security descriptors:

- **Merge**—The security descriptor entries of the source objects will be added to the security descriptors of the target objects.

- **Skip**—The security descriptors of the target objects will be left intact.

- **Replace**—The entries of the target object security descriptor will be replaced with those of the source objects.

Add SIDHistory—Select this checkbox if you want to allow the target accounts to access the source domain resources using SIDHistory mechanism during the coexistence period. For more information on SIDHistory adding, see Adding SID History.

The **User Principal Name handling** section allows you specify how User Principal Names (UPNs) will be processed:

- **Synchronize**—If this option is selected, source User Principal Names will be assigned to the target users.

- **Skip**—If this option is selected, the target user UPN will be left intact.

- **Set the domain suffix of the UPNs to**—Allows you to set the domain suffix of the UPNs of the target users to the value you specify.

Synchronize passwords—Select this checkbox if you want to synchronize passwords for the accounts.

If one-way directory synchronization is established, the source account's password will be applied to the target account only if it is newer than the one of the target account. No passwords will be applied from target to source accounts even if the target accounts' passwords are newer.

In the case of two-way directory synchronization, the newer password values are synchronized to the opposite directory (that is, passwords are synchronized in both directions).

These rules apply to both initial and delta synchronization.

Step 4. Select Target Objects to Synchronize

This step allows you to select the target directory objects that should be synchronized with the corresponding source objects. You can also select whether the agent should create the objects on the source and select the container where the objects should be created.
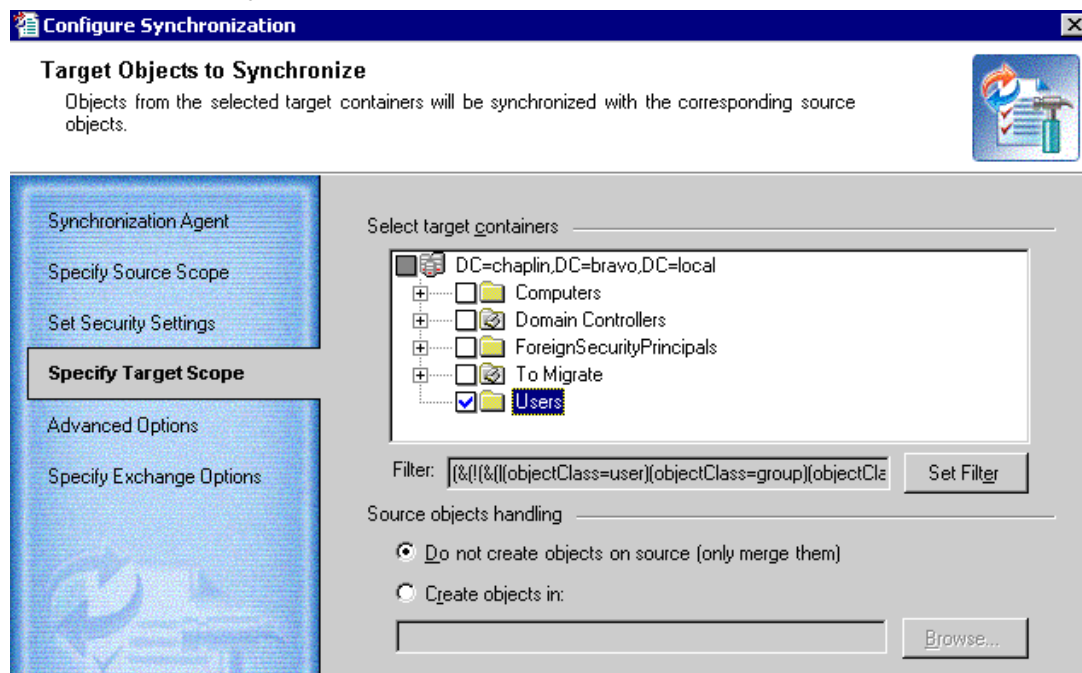
Do not create objects on source (only merge them)—If this option is selected, no new objects will be created on source during synchronization. Only the objects that already exist on source will be synchronized.

Create objects in—You can select whether the agent should create the objects on the source and select the container where the objects should be created. If this option is selected, all the new objects that do not exist on target will be created in the specified container. Click **Browse** to select the container from the source directory tree.

Step 5. Specify Advanced Options

If you are planning to use a custom add-in to process the selected objects after they have been retrieved from the source domain and before they are synchronized to the target domain, select the Use custom add-in check box and browse for the .xml custom add-in file.
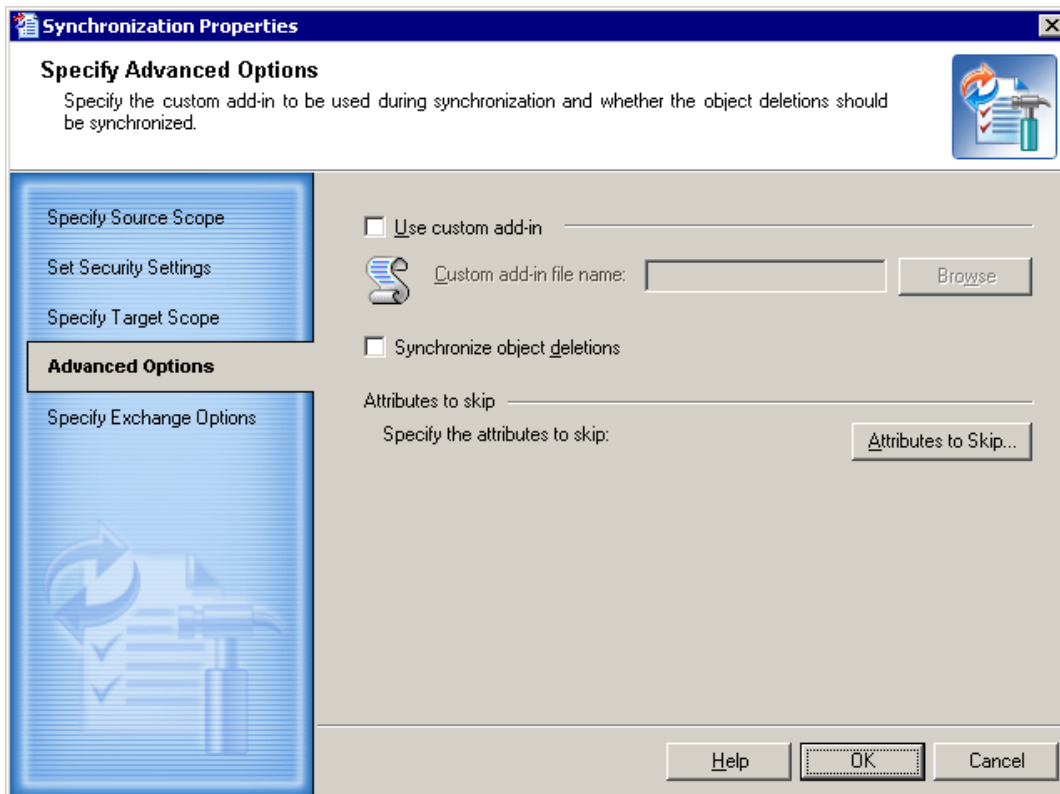
If you want object deletions to be synchronized as well, select the **Synchronize object deletions** check box. If an object in the source domain was deleted, the Directory Synchronization Agent will delete it from the target domain.

> **i** | **NOTE:**
>
> - If an object is deleted on a Windows 2000 Server domain controller in the source domain, it is not deleted from the target domain regardless of scope settings; a message like the following is written to the DSA log:
>   *For safety reasons deletion of following objects will not be synchronized, being propagated by DC running Windows 2000 Server.*
>
> - Deletion of the matched objects in the target organization is not synchronized from target to source even if the "Synchronize object deletions" option is enabled in the synchronization properties. As a result, the source objects are not affected if the matched target objects are removed.

Attributes to Skip

Select the check boxes next to the attributes you do not want to synchronize.

Use the **Save Settings** button if you want to use the same settings when you create another migration session. Click **Load Settings** to apply a set of attribute skipping options you saved earlier.
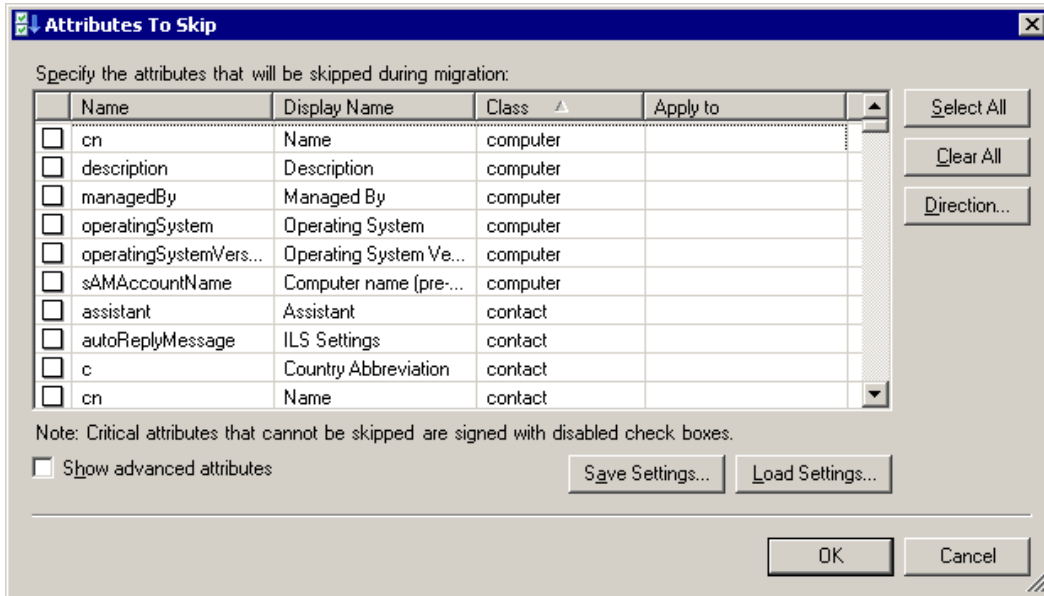
Direction—Click this button to specify the type of synchronization during which the attribute should be skipped:

- **Two-way Sync**—Select this option to skip the attribute from being applied on either side.
- **Source-to-Target Sync**—Select this option if you do not want to apply the attribute on target.
- **Target-to-Source**—Select this option if you do not want to apply the attribute on source.

Select the **Show advanced attributes** check box to choose from the complete list of attributes you can skip.

> **!** **CAUTION: The Directory Synchronization Agent service attributes used by Migration Manager should never be skipped. Otherwise, Migration Manager will not be able to migrate and synchronize objects.**

Step 6. Specify Exchange Options

If Exchange migration is planned for after directory migration, you may want to specify some options required for further Exchange messaging system synchronization.

> **i** **NOTE:** This step is displayed only when Exchange Server is installed in the source or target forest, or in both.
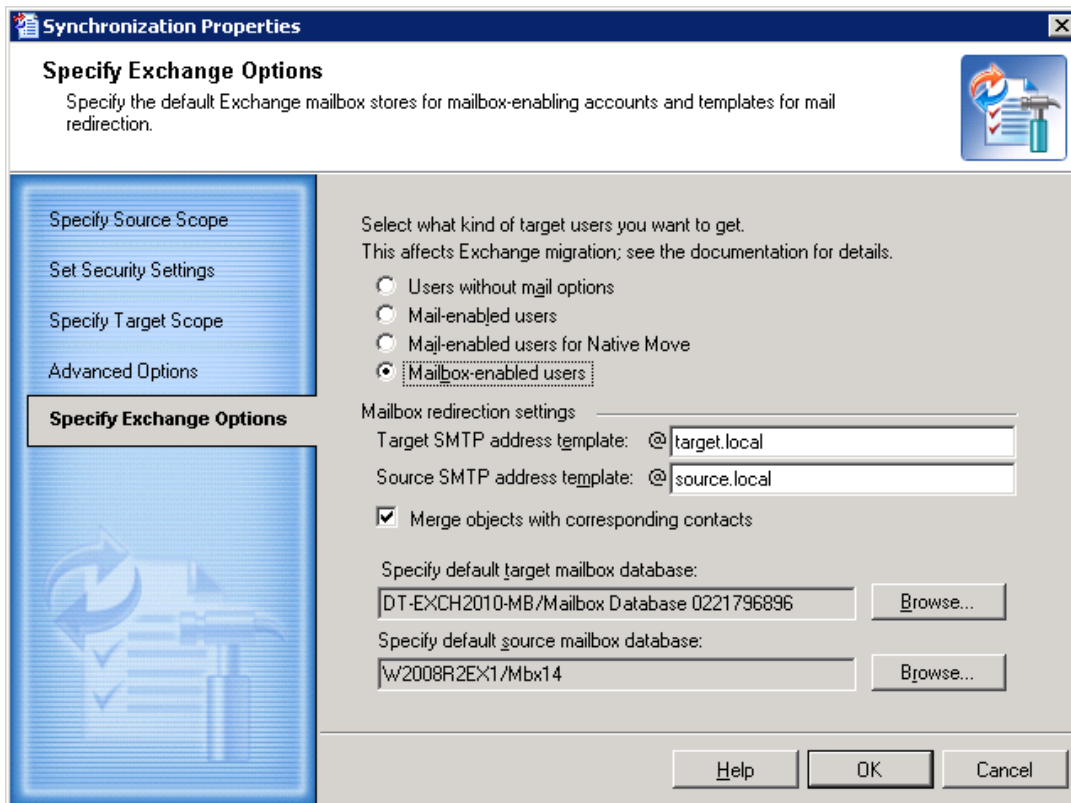
Type of Target Users

The group of options at the top lets you specify how the Directory Synchronization Agent should configure the target users' Exchange-related options:

- **Users without mail options**
  This means the target users have no attributes related to Exchange mail capabilities. This option is for situations where no Exchange migration is planned and the target users are not going to use mail. For newly-created target users, the results of this choice are as follows:

The Source User Was	The Target User Is
Without mail options	Without mail options
Mail-enabled	Without mail options
Mailbox-enabled	Without mail options

> **i** | **NOTE:** This option also affects groups and contacts.

- **Mail-enabled users**
  For newly-created target users, the results of this choice are as follows:

The Source User Was	The Target User Is
Without mail options	Without mail options
Mail-enabled	Mail-enabled
Mailbox-enabled	Mail-enabled

> **i** | **NOTE:** This option also affects groups and contacts.

- **Mail-enabled users for Native Move**
  For newly-created target users, the results of this choice are as follows:

The Source User Was	The Target User Is
Without mail options	Without mail options
Mail-enabled	Mail-enabled
Mailbox-enabled	Mail-enabled user for Native Move

> **i** | **NOTE:** This option also affects groups and contacts.

- **Mailbox-enabled users**
  For newly-created target users, the results of this choice are as follows:

The Source User Was	The Target User Is
Without mail options	Without mail options
Mail-enabled	Mail-enabled
Mailbox-enabled	Mailbox-enabled

If you are merging source users with existing target users, one of three mail option configurations can be applied to each user:

- Source user configuration
- Target user configuration
- Choice on the Specify Exchange Options step

In this case, the Directory Synchronization Agent uses the following logic to decide which configuration to apply:

- If the target configuration is the most "elevated", keep it.
- If the target configuration is the least "elevated", apply the "lesser" of the other two configurations.

> **i** | **NOTE:** If you need to convert your existing mail-enabled users to mailbox-enabled users, select the Mailbox-enabled users option. This will work only for those users that are mailbox-enabled in the source.

If Exchange migration is planned, use either the **Mail-enabled users**, **Mail-enabled users for Native Move** or the **Mailbox-enabled users** option. Your choice of option should depend on the details of the Exchange migration plans, as explained in the Exchange Migration Considerations topic.

Mail Redirection

For mail redirection purposes, the Directory Synchronization Agent adds the secondary SMTP address to the proxy addresses list of the mailbox-enabled object (the proxyAddresses attribute). You can specify the SMTP address templates to create the secondary SMTP addresses that will be applied to the source and target mailbox-enabled objects.

> ⚠ **CAUTION: You should analyze your environment for SMTP namespaces and for redirection implement SMTP address templates that are NOT being used.**

To forward mail to a recipient in another Exchange organization, the Directory Synchronization Agent populates the target address property (the **targetAddress** attribute) of either the source or target mailbox, depending on which mailbox is currently being used, with the additional SMTP address created for redirection.

Target SMTP address template—Specify the SMTP address template for the target accounts so that the target users receive their mail during the synchronization.

Source SMTP address template—Specify the SMTP address template for the source accounts so that the source users receive their mail during the synchronization.

If contacts with the same SMTP or X500 address as the synchronized objects already exist in the opposite directory, Migration Manager can merge SMTP addresses and membership for these objects and delete the corresponding contacts. To use this functionality, select the **Merge objects with corresponding contacts** check box. This option is applicable if you select the **Mail-enabled users**, **Mail-enabled users for Native Move** or the **Mailbox-enabled users** option above.

For uninterrupted user collaboration during the migration, users in each source and target Exchange organization should see other users in their Global Address Lists. That is why two mailboxes, source and target, exist for each user. However, all mail sent to the user should arrive to the mailbox he or she is currently using, no matter which mailbox it was sent to. To achieve this, mail should be automatically forwarded to the currently-used mailbox from the other mailbox.

Direct forwarding to a recipient in another Exchange organization can be done by using the target address property of the mailbox.

Mailbox Databases

Target mailbox database—Specify the mailbox database where the target users should get mailboxes. This option is required if you selected **Mail-enabled users** or **Mailbox-enabled users** above.

> ℹ **NOTE:**
>
> - Although mail-enabled users do not really have mailboxes, the mailbox database needs to be set so that Exchange-related attributes can be configured correctly.
>
> - The mailbox databases you select here can be easily overridden during Exchange migration.

Source mailbox database—Specify the mailbox database where the source users should get mailboxes. Use this option only if two-way mailbox synchronization is expected and the source users did not originally have mailboxes (that is, were mail-enabled or had no mail options). This option is required if you selected **Mailbox-enabled users** above.

Starting and Stopping Directory Synchronization

After you have configured the directory synchronization job for a pair of domains, you can start this job. To start directory synchronization between a pair of domains, complete the following steps:

1. In Migration Manager console, expand the domain pair node in the left-hand pane.

2. Right-click the **Synchronization** node and select **Start and Re-sync** in the shortcut menu. The **Job status** on the **Synchronization Statistics** screen is displayed as **Starting**.

3. Wait until the **Last operation progress** parameter for the **Start operation** reaches 100% on the **Agent Statistics** screen in the right-hand pane.

Use the **Synchronization Statistics** and **Agent Statistics** screens to track the synchronization job status and progress. For more information about viewing the directory synchronization statistics, refer to the Viewing Directory Synchronization Statistics topic.
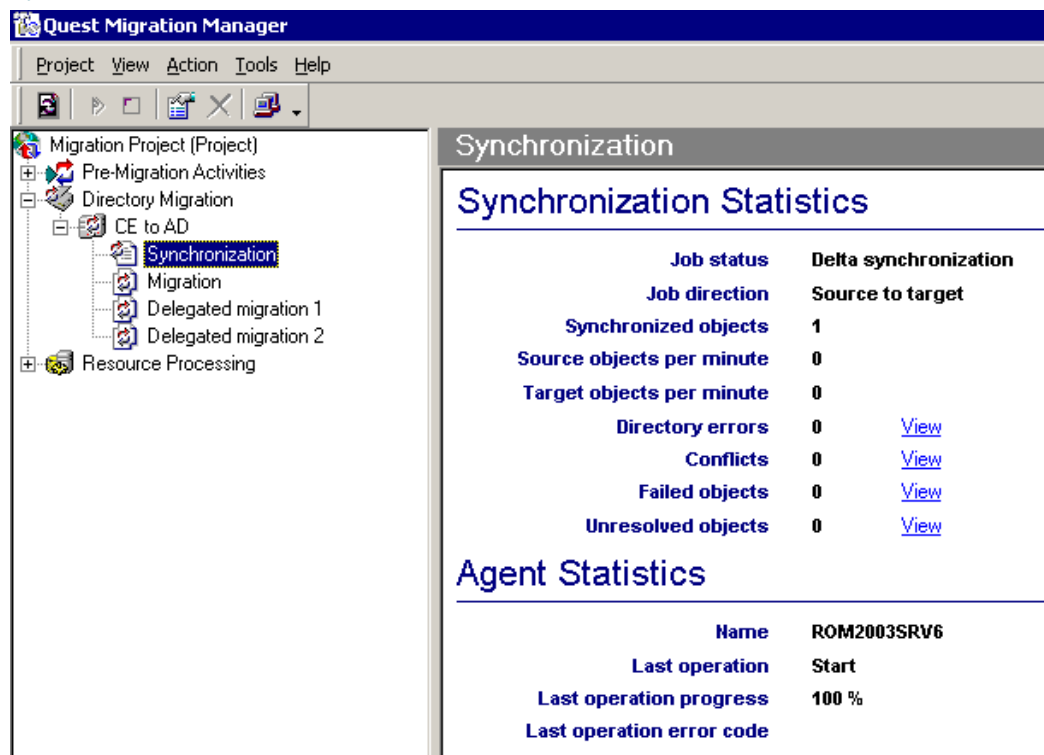
i **NOTE:** The first time you start a synchronization job, the first Directory Synchronization Agent session starts and full resynchronization takes place. The job status on the **Synchronization Statistics** screen is displayed as **Initial synchronization**. All objects in source domain and also objects in target domain if two-way synchronization is performed, are enumerated during the initial sync, but only the objects from the specified source and target scopes are synchronized. After the first synchronization session is completed, only delta changes are synchronized during the subsequent sessions.

To stop directory synchronization between a pair of domains, complete the following steps:

1. In Migration Manager console, expand the domain pair node in the left-hand pane.

2. Right-click the **Synchronization** node and select **Stop** in the shortcut menu. The **Job status** on the **Synchronization Statistics** screen is displayed as **Stopping**.

3. Wait until the **Last operation progress** parameter for the **Stop operation** reaches 100% on the **Agent Statistics** screen in the right-hand pane.

The following changes to the synchronization job require full directory re-synchronization:

- Changing the source scope (adding containers to the synchronization scope or removing them)
- Changing the target scope (adding containers to the synchronization scope or removing them)
- Changing Security options
- Changing Advanced options
- Changing Exchange options

The following changes made to the domain pair configuration also result in full directory re-synchronization:

- Changing the service attributes
- Applying conflict resolution rules

Viewing Directory Synchronization Statistics

When you select the Synchronization node under the domain pair in the migration project tree, you can see synchronization state, progress, and statistics for the selected pair of domains on the screen in the right-hand pane.

Synchronization Statistics

The following information is provided on the Synchronization Statistics screen:

Job status:

- **Initial synchronization**—Indicates that initial synchronization (full re-sync) is going on. Either this is the first synchronization session for the Directory Synchronization Agent or synchronization was re-started with Start and Re-sync option.

- **Delta synchronization**—Indicates that the initial synchronization is completed and only delta changes are synchronized.

- **Sleeping**—Indicates that the Directory Synchronization Agent is in an idle state, either because it is in a sleep interval between synchronization sessions or due to its schedule restrictions. Refer to the *Configuring the Directory Synchronization Agent* section in the Agent Manager topic for more details.

Job direction:

- **Source to target**—Indicates that synchronization job is configured to synchronize objects from source to target only (only the source scope is set).

- **Target to source**—Indicates that synchronization job is configured to synchronize objects from target to source only (only the target scope is set).

- **Two-way**—Indicates that both source and target objects will be synchronized (both source and target scopes are set).

Synchronized objects—Displays the number of object pairs that were matched and synchronized.

Source objects per minute—Shows the approximate speed of the source-to-target synchronization process

Target objects per minute—Shows the approximate speed of the target-to-source synchronization process

Directory errors—Shows the number of errors that occurred during synchronization, such as connection errors, invalid credentials errors, and server unavailability errors. Click **View** next to the counter to see the list of errors.

> **i** | **NOTE:** The Directory errors queue is not dynamic and errors put in this queue are not removed from the queue even after they are resolved.

You also can view the queues of conflicts found during synchronization, failed objects, and unresolved objects. Click the corresponding **View** link for detailed information.

Conflicts—This queue contains the object-matching conflicts.

Failed objects—This queue contains issues that occurred to the synchronized objects and their attributes, such as if an object could not be created due to insufficient rights.

Unresolved objects—All unresolved linked attributes, such as unresolved membership, are put in this queue.

> **i** | **NOTE:** The **Conflicts, Failed objects** and **Unresolved objects** queues are dynamic, so if, for example, a conflict is resolved, it is thrown out of the queue.

Agent Statistics

Name—The name of the server running the agent.

Last operation—The last operation description the agent was instructed to perform (such as start or stop).

Last operation progress—The last operation's progress (in percent).

Last operation error code—If the last operation completed successfully, this field is blank. If the operation failed, the error code is displayed.

Keeping Existing Lync Infrastructure

Migration Manager for Active Directory supports migration scenarios where Microsoft Lync Server is installed in either source or target domain and you want to provide users access to existing Lync infrastructure.

Pre-requisites

Before implementing any of the below scenarios, check the following pre-requisites:

- In case of Inter-forest migration when source and target domains reside in different forests, trusts must be established between the source and target domains.
- Active Directory schema must be extended in both source and target domains.
- Lync users' computers must trust certificate installed on the Lync Server.

Lync Server in Source Domain

If you use Lync Server in the source domain for communicating, users can still communicate with other users through that Lync Server even after they are migrated to the target domain.
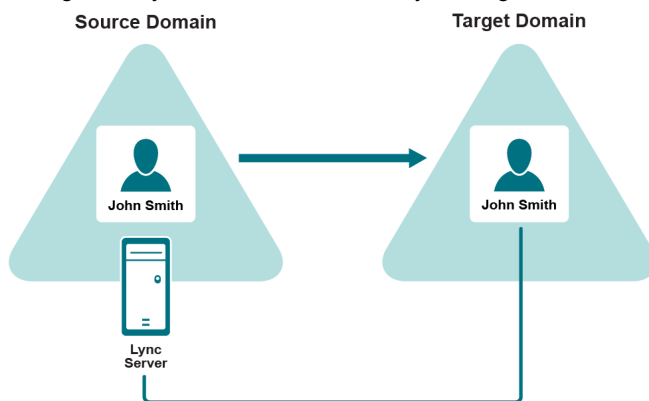
To achieve that, when configuring migration session or directory synchronization job, change settings in the wizard as follows:

- **Migration Session**: On the **Specify Object Processing Options** step select **Use custom add-in** and specify add-in located at *<Migration Manager installation folder>\Active Directory\SourceLyncSupport.xml*.
- **Directory Synchronization**: On the **Specify Advanced Options** step, select **Use custom add-in** and specify add-in located at *<Migration Manager installation folder>\Active Directory\SourceLyncSupport.xml*.

After migration session or initial synchronization completes, migrated user may log on to his or her target account and start communicating with other users through the source Lync Server.

> **!** **CAUTION: Corresponding source user must not be removed. Otherwise, target user will lose access to the source Lync Server.**

Lync Server in Target Domain

If you have Lync Server deployed in the target domain, you may want to provide users that have been previously migrated ability to communicate with other users through that Lync Server.

For that you need to take the following steps:

1. Enable target user account for using Lync Server. For information how to do that, refer to this TechNet article.

2. Configure a new migration session from the source domain to the target domain according to Creating a Migration Session with the following specific settings:

   a. On the **Select Source Objects** step specify source user accounts that correspond to the target accounts you Lync-enabled on step 1.
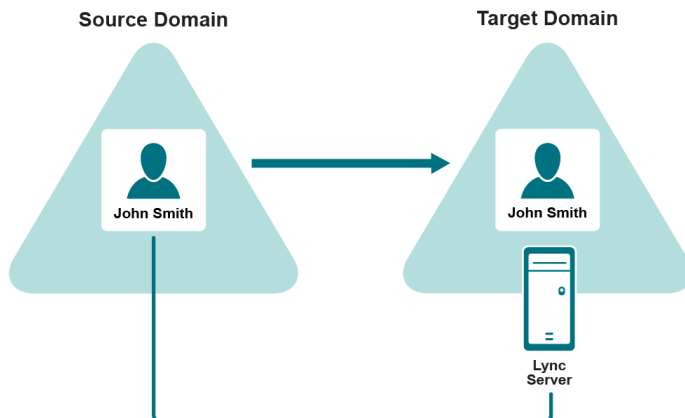
   b. On the **Specify Object Processing Options** step, select **Use custom add-in** and specify add-in located at *<Migration Manager installation folder>\Active Directory\TargetLyncSupport.xml*.

After migration session completes, migrated user may log on to the source account and start communicating with other users through the target Lync Server.

> ! **CAUTION: Corresponding target user must not be removed. Otherwise, source user will lose access to the target Lync Server.**

Adding SID History

SID History adding allows target accounts to access source domain resources during the coexistence period.

To add SID History, select **Add SIDHistory** option at **Set Security Settings** step when configuring account migration or directory synchronization.

Migration Manager for Active Directory offers two methods of SID History adding:

- Using Migration Manager for Active Directory SIDHistory Agent installed on a domain controller (default option)

- Agentless using `DsAddSidHistory` function provided by Microsoft

Both methods have their own benefits and require different permissions and preparation steps.

> **i** **TIP:** To take advantage of SID history, ensure that trusts are established and that SID filtering is turned off between the source and target domains.

Adding with SIDHistory Agent

The easiest adding method that requires no preparation steps or special configuration. SIDHistory Agent is automatically installed on the target domain controller and then controlled by the Directory Synchronization Agent.

This method requires the following:

- Target DSA account must be a member of the **Administrators** group in the target domain. If you consent to grant such permissions for the target DSA account, agent-based adding is a recommended scenario. In case you do not want to grant such permissions to the target DSA account in accordance with the least privilege principle a Preinstalled Service feature should be configured and enabled.

> **i** **TIP:** This adding method is the easiest way to go if you want to synchronize passwords.

Agentless Adding

This adding method doesn't require an agent to be installed on the domain controller and doesn't require target DSA account to be a member of the **Administrators** group in the target domain. SID history is added with native `DsAddSidHistory` function.

This method requires the following:

- Source DSA account must be a member of the **Administrators** group in the source domain.

- You should perform additional preparation steps as described in Configuring the environment for agentless adding and Enabling / reverting agentless SID history adding for the project procedures below.

> **i** **NOTE:** Agentless adding requires both source and target domains to run at Windows Server 2003 domain functional level or higher.

Configuring the environment for agentless adding:

- Ensure the **Source DSA Account** is a member of the **Administrators** group in the source domain

- Grant **Migrate SID history** permission to **Target DSA Account** on the target domain object.

i | **NOTE:** This permission must be applied to **This object only**.

- Create a local group in the source domain to support auditing. The group should be named <source_domain_name>$$$, where `source_domain_name` is the NetBIOS name of the source domain.

i | **IMPORTANT:** Do not add members to this group or adding SID history will fail.

- Enable account management auditing in the source and target domains. For SID history adding between forests under Windows Server 2008 and later, also enable directory service access auditing.

  You should turn on auditing of **Success** and **Failure** attempts for **Audit account management** and **Success** attempts for **Audit directory service access**. This can be done in **Computer Configuration | Policies | Windows Settings | Security Settings | Local Policies | Audit Policy** on the source and target domain controllers (Windows Server 2008 or later.)

i | **TIP:** For immediately applying the policy changes, open an elevated command prompt on the domain controller and execute **gpupdate /force**.

Enabling / reverting agentless SID history adding for the project

The agentless SID history adding is manually enabled for the project with `UseAgentlessSIDHistoryMigration.ps1` script provided on the Migration Manager for Active Directory installation disk. Open 32-bit (x86) version of the PowerShell prompt on the computer where Migration Manager is installed and execute the following commands:

```
cd "<Drive with product CD>:\QMM ResKit\Scripts"

.\UseAgentlessSIDHistoryMigration.ps1 $true
```

To revert to SIDHistory Agent, execute the following command:

```
.\UseAgentlessSIDHistoryMigration.ps1 $false
```

About us

Quest creates software solutions that make the benefits of new technology real in an increasingly complex IT landscape. From database and systems management, to Active Directory and Office 365 management, and cyber security resilience, Quest helps customers solve their next IT challenge now. Around the globe, more than 130,000 companies and 95% of the Fortune 500 count on Quest to deliver proactive management and monitoring for the next enterprise initiative, find the next solution for complex Microsoft challenges and stay ahead of the next threat. Quest Software. Where next meets now. 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